Contents

Quick tasks for getting started with Microsoft 365 compliance

2/18/2021 • 10 minutes to read • Edit Online

If you're new to Microsoft 365 compliance and wondering where to start, this article provides guidance on the basics and prioritizes important compliance tasks. This article will help you quickly get started with managing and monitoring your data, protecting information, and minimizing insider risks.

This article is also helpful if you're figuring out how best to manage risks, protect your data, and remain compliant with regulations and standards with a newly remote workforce. Employees are now collaborating and connecting with each other in new ways, and this means your existing compliance processes and controls may need to adapt. Identifying and managing these new compliance risks within your organization is critical to safeguarding your data and minimizing threats and risks.

After you've completed these basic compliance tasks, consider expanding compliance coverage in your organization by implementing additional Microsoft 365 compliance solutions.

Task 1: Configure compliance permissions

It's important to manage who in your organization has access to the Microsoft 365 compliance center to view content and perform management tasks. Microsoft 365 provides administrative roles specific to compliance and for using the tools included in the Microsoft 365 compliance center.

Start by assigning compliance permissions to the people in your organization so that they can perform these tasks and to prevent unauthorized people from having access to areas outside of their responsibilities. You'll want to make sure that you've assigned the proper people to the **Compliance data administrator** and the **Compliance administrator** admin roles before you start to configure and implement compliance solutions included with Microsoft 365. You'll also need to assign users to the Azure Active Directory global reader role to view data in Compliance Manager.

For step-by-step guidance to configure permissions and assign people to admin roles, see Permissions in the Security & Compliance Center.

Task 2: Know your state of compliance

It's difficult to know where to go if you don't know where you are. Meeting your compliance needs includes understanding your current level of risk and what updates may be needed in these ever changing times. Whether your organization is new to compliance requirements or has deep experience with standards and regulations that govern your industry, the single best thing you can do to improve compliance is to understand where your organization stands.

Microsoft Compliance Manager can help you understand your organization's compliance posture and highlight areas that may need improvement. Compliance Manager uses a centralized dashboard to calculate a risk-based score, measuring your progress in completing actions that help reduce risks around data protection and regulatory standards. You can also use Compliance Manager as a tool to track all your risk assessments. It provides workflow capabilities to help you efficiently complete your risk assessments through a common tool.

For step-by-step guidance to get started with Compliance Manager, see Get started with Compliance Manager.

IMPORTANT

Security and compliance are tightly integrated for most organizations. It's important that your organization addresses basic security, threat protection, and identity and access management areas to help provide a defense in-depth approach to both security and compliance.

Check your Microsoft 365 Secure Score in the Microsoft 365 security center and completing the tasks outlined in the following articles:

- Security roadmap - Top priorities for the first 30 days, 90 days, and beyond
- Top 12 tasks for security teams to support working from home

Task 3: Enable auditing for your organization

Now that you've determined your organization's current state and who can manage compliance functions, the next step is to make sure you have the data to conduct compliance investigations and generate reports for network and user activities in your organization. Enabling auditing is also an important prerequisite for compliance solutions covered later in this article.

Insights provided by the audit log are a valuable tool in helping to match your compliance requirements to solutions that can help you manage and monitor compliance areas needing improvement. Audit logging must be enabled before activities are recorded and before you can search the audit log. When enabled, user and admin activity from your organization is recorded in the audit log and retained for 90 days, and up to one year depending on the license assigned to users.

For step-by-step instructions to turn on auditing, see Turn audit log search on or off.

Task 4: Create policies to alert you about potential compliance issues

Microsoft provides several built-in alert policies that help identify admin permissions abuse, malware activity, potential external and internal threats, and information governance risks. These policies are turned on by default, but you may need to configure custom alerts to help manage compliance requirements specific to your organization.

Use alert policy and alert dashboard tools to create custom alert policies and view the alerts generated when users perform activities that match the policy conditions. Some examples could be to use alert policies to track user and admin activities affecting compliance requirements, permissions, and data loss incidents in your organization.

For step-by-step guidance to create custom alert policies, see Alert policies in the security and compliance center.

Task 5: Classify and protect sensitive data

To get their work done, people in your organization collaborate with others both inside and outside the organization. This means that content no longer stays behind a firewall—it can roam everywhere, across devices, apps, and services. And when it roams, you want it to do so in a secure, protected way that meets your organization's business and compliance policies.

Sensitivity labels let you classify and protect your organization's data, while making sure that user productivity and their ability to collaborate isn't hindered. Use sensitivity labels to enforce encryption and usage restrictions apply visual markings, and protect information across platforms and devices, on-premises and in the cloud.

For step-by-step guidance to configure and use sensitivity labels, see Get started with sensitivity labels. For sensitivity label licensing information, see Microsoft 365 licensing guidance for security & compliance.

Task 6: Configure a retention policy

A retention policy lets you proactively decide whether to retain content, delete content, or both—retain and then delete the content at the end of a specified retention period. These actions might be needed to comply with industry regulations and internal policies, as well as reduce your risk in the event of litigation or a security breach.

When content is subject to a retention policy, people can continue to edit and work with the content as if nothing's changed. The content is retained in place, in its original location. But if someone edits or deletes content that's subject to the retention policy, a copy of the original content is saved to a secure location where it's retained while the retention policy for that content is in effect.

You can quickly put a retention policy in place for multiple locations in your Microsoft 365 environment such as Exchange mail, SharePoint sites, OneDrive accounts, and Microsoft 365 groups. There are no limits to the number of mailboxes or sites this policy can automatically include. But if you need to get more selective, you can do so by configuring a retention policy for specific locations and include or exclude sites or users.

For step-by-step guidance to configure a retention policy, see Create and configure retention policies. If you're new to configuring retention in Microsoft 365, see Get started with retention policies and retention labels.

Task 7: Configure sensitive information and offensive language policies

Protecting sensitive information and detecting and acting on workplace harassment incidents is an important part of compliance with internal policies and standards. Communication compliance in Microsoft 365 helps minimize these risks by helping you quickly detect, capture, and take remediation actions for email and Microsoft Teams communications. These include inappropriate communications containing profanity, threats, and harassment and communications that share sensitive information inside and outside of your organization.

A pre-defined *Offensive language and anti-harassment* policy template allows you to scan internal and external communications for policy matches so they can be examined by designated reviewers. Reviewers can investigate scanned email, Microsoft Teams, Yammer, or third-party communications in your organization and take appropriate remediation actions to make sure they're compliant with your organization's standards.

The pre-defined *Sensitive information* policy template helps you quickly create a policy to scan email and Microsoft Teams communications containing defined sensitive information types or keywords to help make sure that important data isn't shared with people that shouldn't have access. These activities could include unauthorized communication about confidential projects or industry-specific rules on insider trading or other collusion activities.

For step-by-step guidance to plan and configure communication compliance, see Plan for communication compliance and Get started with communication compliance. For communication compliance licensing information, see Microsoft 365 licensing guidance for security & compliance.

Task 8: See what's happening with your sensitive items

Sensitivity labels, sensitive information types, retention labels and policies and trainable classifiers can be used to classify and label sensitive items across Exchange, SharePoint, and OneDrive as you've seen in the previous tasks. The last step in your quick task journey is to see which items have been labeled and what actions your users are taking on those sensitive items. content explorer and activity explorer provide this visibility.

Content explorer

Content explorer allows you to view, in their native format, all the items that have been classified as a sensitive information type or belonging to a certain classification by a trainable classifier, as well as all items that have sensitivity or retention label applied.

For step-by-step guidance to using content explorer, see Know your data - data classification overview, and Get started with content explorer.

Activity explorer

Activity explorer helps you monitor what's being done with your classified and labeled sensitive items across:

- SharePoint
- Exchange
- OneDrive

There are over 30 different filters available for use, some are:

- date range
- activity type
- location
- user
- sensitivity label
- retention label
- file path
- DLP policy

For step-by-step guidance to using activity explorer, see Get started with activity explorer.

Next steps

Now that you've configured the basics for compliance management for your organization, consider the following compliance solutions in Microsoft 365 to help you protect sensitive information and detect and act on additional insider risks.

Configure retention labels

While retention policies apply at the container level to locations such as SharePoint sites and Exchange mailboxes, retention labels allow for more specific targeting for your retention and deletion policies. For example, at the document or email message level that end users can apply manually in addition to automatic application by administrators. You can also apply a retention label to a document library, folder, or document set in SharePoint, so that all documents that are stored in that location inherit the default retention label.

Additionally, retention labels support records management to mark content as a record. When this happens, the label places additional restrictions on the content that might be needed to help your organization comply with regulatory requirements.

For step-by-step guidance to create and publish retention labels, see the following guidance:

- Create retention labels and apply them in apps
- Apply a retention label to content automatically

To get started with records management, see Get started with records management.

Identify and define sensitive information types

Define sensitive information types based on the pattern contained in information in your organization's data. Use built-in sensitive information types help identify and protect credit card numbers, bank account numbers, passport numbers, and more. Or create your own custom sensitivity information types specific to your organization.

For step-by-step guidance to define custom sensitive information types, see Create a custom sensitive information type in the Security & Compliance Center.

Prevent data loss

Data loss prevention (DLP) policies allow you to identify, monitor, and automatically protect sensitive information across your Microsoft 365 organization. Use DLP policies to identify sensitive items across Microsoft services, prevent the accidental sharing of sensitive items, and help users learn how to stay compliant without interrupting their workflow.

For step-by-step guidance to configure DLP policies, Create, test, and tune a DLP policy. For data loss management licensing information, see Microsoft 365 licensing guidance for security & compliance.

Detect and act on insider risks

More and more, employees have increasing access to create, manage, and share data across a broad spectrum of platforms and services. In most cases, organizations have limited resources and tools to identify and mitigate organization-wide risks while also meeting compliance requirements and employee privacy standards. These risks may include data theft by departing employees and data leaks of information outside your organization by accidental oversharing or malicious intent.

Insider risk management in Microsoft 365 uses the full breadth of service and 3rd-party indicators to help you quickly identify, triage, and act on risky user activity. By using logs from Microsoft 365 and Microsoft Graph, insider risk management allows you to define specific policies to identify risk indicators and to take action to mitigate these risks.

For step-by-step guidance to plan and configure insider risk management policies, see Plan for insider risk management and Get started with insider risk management. For insider risk management licensing information, see Microsoft 365 licensing guidance for security & compliance.

What's new in Microsoft 365 compliance

2/18/2021 • 12 minutes to read • Edit Online

Whether it be adding new solutions to the Microsoft 365 compliance center, updating existing features based on your feedback, or rolling out fresh and updated documentation, Microsoft 365 helps you stay on top of the ever-changing compliance landscape. Take a look below to see what's new in Microsoft 365 compliance today.

> **NOTE**
>
> Some compliance features get rolled out at different speeds to our customers. If you aren't seeing a feature yet, try adding yourself to targeted release.

> **TIP**
>
> Interested in what's going on in other admin centers? Check out these articles:
> What's new in the Microsoft 365 admin center
> What's new in the SharePoint admin center
> What's new in Microsoft 365 Defender
>
> And visit the Microsoft 365 Roadmap to learn about Microsoft 365 features that were launched, are rolling out, are in development, have been cancelled, or previously released.

January 2021

Support for card content in Teams

The following Microsoft 365 compliance solutions now support the detection of card content generated through apps in Teams messages:

- **Core and Advanced eDiscovery**. Card content can now be placed on hold or included in searches (applies to content search as well).
- **Audit**. Card activity is now recorded to the audit log.
- **Retention policies**. Can now use retention policies to retain and delete card content.

Information governance and records management

New assessment to address using information governance and records management to help meet compliance obligations for the New Zealand Public Records Act.

Sensitivity labels

- Sensitivity labels are now supported for US Government tenants (GCC and GCC-H).
- New automatic labeling support for macOS.

December 2020

Spotlight: New content for insider risk solutions

The Microsoft 365 compliance content team is hard at work creating 'content solution' docs to promote how compliance capabilities can be used together to help meet your compliance goals.

First up is content that ties together our insider risk solutions: communication compliance, insider risk management, information barriers, and privileged access management. Here's a peek at what you'll find:

- New landing page for insider risk solutions. Includes details about risks that the solutions can help mitigate, licensing requirements, deployment sequence, architecture illustrations, training resources, and more.
- New overview articles for each insider risk solution. Guidance and links to articles that help you learn about, plan, deploy, and manage each solution:
  - Communication compliance
  - Insider risk management
  - Information barriers
  - Privileged access management

More content solution docs coming soon!

Advanced eDiscovery

Improved workflow and functionality for adding custodians and non-custodial data sources to an Advanced eDiscovery case.

Data connectors

Four new Globanet connectors released: Redtail Speak, Salesforce Chatter, ServiceNow, and Yieldbroker.

Encryption

Introducing Customer Key for Microsoft 365 at the tenant level. Using keys you provide, you can create a data encryption policy (DEP) and assign it to the tenant. The DEP encrypts data across the tenant for these workloads:

- Teams chat messages (1:1 chats, group chats, meeting chats and channel conversations)
- Teams media messages (images, code snippets, videos, wiki images)
- Teams call and meeting recordings stored in Teams storage
- Teams chat notifications
- Teams chat suggestions by Cortana
- Teams status messages
- User and signal information for Exchange Online

Records management

The Records Management admin role group now grants permissions for all records management features, including disposition review.

Sensitivity labels

- Automatically label data in Azure Purview (preview). You can now create and automatically apply sensitivity labels to assets in Azure Purview, such as files in Azure Blob storage and database columns in SQL Server.
- Require users to apply a label to items. Also known as 'mandatory labeling', this new option requires users to choose and apply a sensitivity label under the specific scenarios.

November 2020

Just a reminder that we often release new and updated features in a preview state to learn how they're being used so we can hone and improve them before releasing to general availability. Your feedback is critical during preview (and beyond), so be sure to let us know what you think by opening the Feedback card at the bottom right of the compliance center.

Spotlight: Endpoint data loss prevention (DLP) released

Endpoint DLP extends the activity monitoring and protection capabilities of DLP to sensitive info on Windows 10 devices. After devices are onboarded to the Microsoft 365 compliance center, you can set up DLP policies to

protect the sensitive info on those devices.

Advanced eDiscovery

To make it easier to manage encrypted content in the eDiscovery workflow, Microsoft 365 eDiscovery tools now incorporate decryption of encrypted files that are attached to email messages and sent in Exchange. Additionally, encrypted documents stored in SharePoint and OneDrive are decrypted in Advanced eDiscovery.

Compliance Manager

- Support for Microsoft 365 Government subscriptions. Compliance Manager is now available to US Government Community (GCC) Moderate and High customers.
- Microsoft Compliance Configuration Analyzer for Compliance Manager. New PowerShell-based tool that helps you get started with Compliance Manager by scanning your organization's current configurations and validating them against Microsoft 365 recommended best practices.
- New templates. Added 56 new templates, bringing total Compliance Manager templates to over 230.

Data connectors

Five new Globanet connectors in preview. New connectors include Reuters Dealing, Reuters FX, CellTrust, XIP, generic MS SQL Database data.

Retention labels (disposition review)

To view items during a disposition review, users must now be members of the Content Explorer Content Viewer and Content Explorer List Viewer role groups. Although required to review items, these role groups aren't necessary for completing the disposition review.

Sensitivity labels

- (Preview) External sharing settings for SharePoint sites. When creating a label that will be used for groups and sites, you'll see an option to control external sharing for SharePoint sites that have the label applied. You can specify that sharing is allowed for anyone, new and existing guests, existing guests only, or just users in your organization. When the label is applied, the label settings will replace any external sharing settings configured in the SharePoint admin center.
- Remove label and encryption from a labeled document. To remove both a label and the encryption it enforces from a labeled document in SharePoint, global admins and SharePoint admins can run the new `Unlock-SPOSensitivityLabelEncryptedFile` cmdlet. This cmdlet runs even if the admin doesn't have access permissions to the site or file, or if the Azure Rights Management service is unavailable.

October 2020

Advanced eDiscovery

CJK language support. Advanced eDiscovery now supports double-byte character set languages, collectively known as CJK languages (includes Simplified Chinese, Traditional Chinese, Japanese, and Korean). These can be used in several advanced review set scenarios.

Sensitivity labels

- Label scope. When creating a sensitivity label, you'll see a new option to define the scope for the label. This option lets you configure labels just for files and emails, containers (like SharePoint sites and Teams), or both.
- Dynamic content marking. When configuring content marking for a sensitivity label, you can now use the dynamic variables such as `${Item.Label}` and `${Item.Location}` in the text string for your header, footer, or watermark.

September 2020

Spotlight: Compliance Manager

Announced at Ignite this year, Compliance Score is rebranded as Compliance Manager. This release completes

the transition from Compliance Manager's previous home in the Service Trust Portal, and introduces an end-to-end compliance management solution in the Microsoft 365 compliance center.

Watch the video below to learn how Compliance Manager can help simplify how your organization manages compliance.

Advanced Audit

- New 10-year retention of audit logs helps support long running investigations and respond to regulatory, legal, and internal obligations.
- Three new crucial events. The following new events can help you investigate possible breaches and determine the scope of compromise: Send, SearchQueryInitiatedExchange, and SearchQueryInitiatedSharePoint.

Communication compliance

- Updated role groups. Communication compliance role groups now match the role group structure available for the insider risk management solution.
- Reports dashboard. Your central location for viewing all communication compliance reports. Report widgets provide a quick view of insights most commonly needed for an overall assessment of the status of communication compliance activities.
- Power Automate flows. Set up flows to automate tasks for alerts and users, notify managers when users trigger an alerts, and more.
- 'Improve classification' remediation action. Alerts containing items that match trainable classifiers might benefit from feedback to help minimize false positives in your organization. The **Improve classification** option lets you provide feedback whether detected items match the classifier configured in the related communication compliance policy. You can even suggest other classifiers to associate with the item to improve match accuracy for future alerts.

Data connectors

- New third-party data connectors. 25 new data connectors, including 14 connectors from Globanet and 8 from Telemessage.
- Physical badging connector. Import physical badging data, such as employee's raw physical access events or any physical access alarms generated by your organization's badging system. Examples include entries to buildings, server rooms, or data centers. Physical badging data can be used by the insider risk management solution to help protect your organization from malicious activity or data theft inside your organization.

Insider risk management

- Microsoft Teams integration. When Teams integration is turned on in insider risk settings, you can coordinate and collaborate with other stakeholders in Teams on tasks like securely sharing and storing data related to individual cases, tracking and reviewing response activities from analysts and investigators, and more.
- Power Automate flows. Set up flows to automate important tasks for cases and users, such as retrieving user, alert, and case info to share with stakeholders and other apps, automating actions like posting to case notes, and more.
- Activity explorer. Available when reviewing alerts, activity explorer provides investigators and analysts with a comprehensive analytic tool for drilling down into each alert. Quickly review a timeline of detected risky activity and identify and filter all risk activities associated with alerts.

Retention policies and retention labels

- Support for Yammer. You can now use retention policies to retain and delete Yammer community messages and private messages.
- Apply labels to Teams meetings recordings. When creating an auto-labeling policy, use the keyword query

editor to identify Teams meeting recordings that are stored in users' OneDrive accounts or in SharePoint.

Records management

Support for regulatory records. Classifying a label as a regulatory record increases the restrictions placed on content to which the label is applied and limits the available management actions for the label itself. For example, after it's applied to content, nobody, not even a global admin, can remove the label. Learn more about which actions are allowed and blocked for regulatory records.

Sensitivity labels

Support for US Government customers. Sensitivity labels are now supported for GCC, GCC High, and DoD customers, only for the Azure Information Protection unified labeling client and scanner.

Trainable classifiers

New retraining and feedback capabilities helps improve accuracy and minimize false positive matches for all custom classifiers and some pre-trained classifiers. This flow lets you provide feedback on whether items match certain classifiers, suggest other classifiers to associate with items, and retrain classifiers to refine and improve match accuracy.

This new capability is included in the following features:

> **NOTE**
> For all features, if you provide at least 30 feedback responses, we'll create a retrained version of that classifier that you can review. If there's improvement, you can republish the classifier.

- Trainable classifiers. To improve the accuracy of your published classifiers, you can provide feedback on whether the detected items match the classifier.
- Communication compliance. The new **Improve classification** remediation action lets you provide feedback whether an item from a communication compliance alert matches the classifier configured in the communication compliance policy.
- Content explorer. If you set up a retention auto-labeling policy to automatically apply labels to email messages that match trainable classifiers, you can use content explorer to review the labeled items and provide feedback whether the items match the classifier.

August 2020

Spotlight: Insider risk and communication compliance updates

Several new and improved features hit public preview this month:

Insider risk management

- Check out our six new policy templates:
  - Data leaks by priority users
  - Data leaks by disgruntled users
  - General security policy violations
  - Security policy violations by departing users
  - Security policy violations by priority users
  - Security policy violations by disgruntled users
- Integration with Microsoft Defender for Endpoint allows you to import and filter Microsoft Defender for Endpoint alerts for activities detected by policies created from the new security violation policy templates. There's also a related insider risk setting where you can choose to import security alerts to insider risk management based on the Microsoft Defender for Endpoint alert triage status.

> **NOTE**
>
> To take advantage of Microsoft Defender for Endpoint integration (including the new security policy violation templates), you'll need to have Microsoft Defender for Endpoint configured in your organization. You'll also need to enable Microsoft Defender for Endpoint for insider risk management integration by configuring advanced features in Microsoft Defender for Endpoint.

- Customize indicator thresholds when creating a policy.
- Set up priority user groups to define users in your organization whose activity requires closer inspection based on factors such as their position, level of access to sensitive information, or risk history.
- Use Office 365 Management Activity APIs to export insider risk alert details to other applications your organization might use to manage or aggregate insider risk data.
- New domain settings help you define and control risk levels for activity in specific domains.

Communication compliance

- When reviewing messages in an alert, you can now remove inappropriate messages in Microsoft Teams channels, 1:1, and group chats. Removed messages and content are replaced with a policy tip that explains that it was removed due to sensitive content.
- New communication roles (these will also be included in new communication compliance role groups releasing in September).
- New communication compliance settings experience that includes settings for privacy and notice templates.
- New classifiers to help detect adult, racy, and gory images.
- New 'Pattern detected' notification that appears when reviewing messages in an alert lets you know about reoccurring instances of the same behavior by a user.

Sensitivity labels

- For US Government tenants (GCC, GCC-H, and DoD), sensitivity labels are currently supported only for the Azure Information Protection unified labeling client and scanner. For more information, see Azure Information Protection Premium Government Service Description.
- You can now use Security & Compliance Center PowerShell to create and configure all settings you see in your labeling admin center. This means that, in addition to using PowerShell for settings that aren't available in the labeling admin centers, you can now fully script the creation and maintenance of sensitivity labels and sensitivity label policies.

Records management: Content overhaul

New docs covering deployment steps, marking content as records, and record versioning:

- Get started with records management
- Declare records by using retention labels
- Use record versioning to update records stored in SharePoint or OneDrive

Retention labels & policies

Retention-related admin activity is now recorded and available to review in the audit log. For the full list, see Retention policy and retention label activities.

Advanced eDiscovery

- When adding a collection to a review set, you can now include modern attachments (also called "cloud attachments") and SharePoint document versions.
- New direct download export experience, eliminating the need to use Azure Storage Explorer to download case content.

Microsoft 365 compliance center

2/18/2021 • 6 minutes to read • Edit Online

If you're interested in your organization's compliance posture, you're going to love the Microsoft 365 compliance center. The Microsoft 365 compliance center provides easy access to the data and tools you need to manage to your organization's compliance needs.

Read this article to get acquainted with the Microsoft 365 compliance center, how to get it, frequently asked questions, and your next steps.

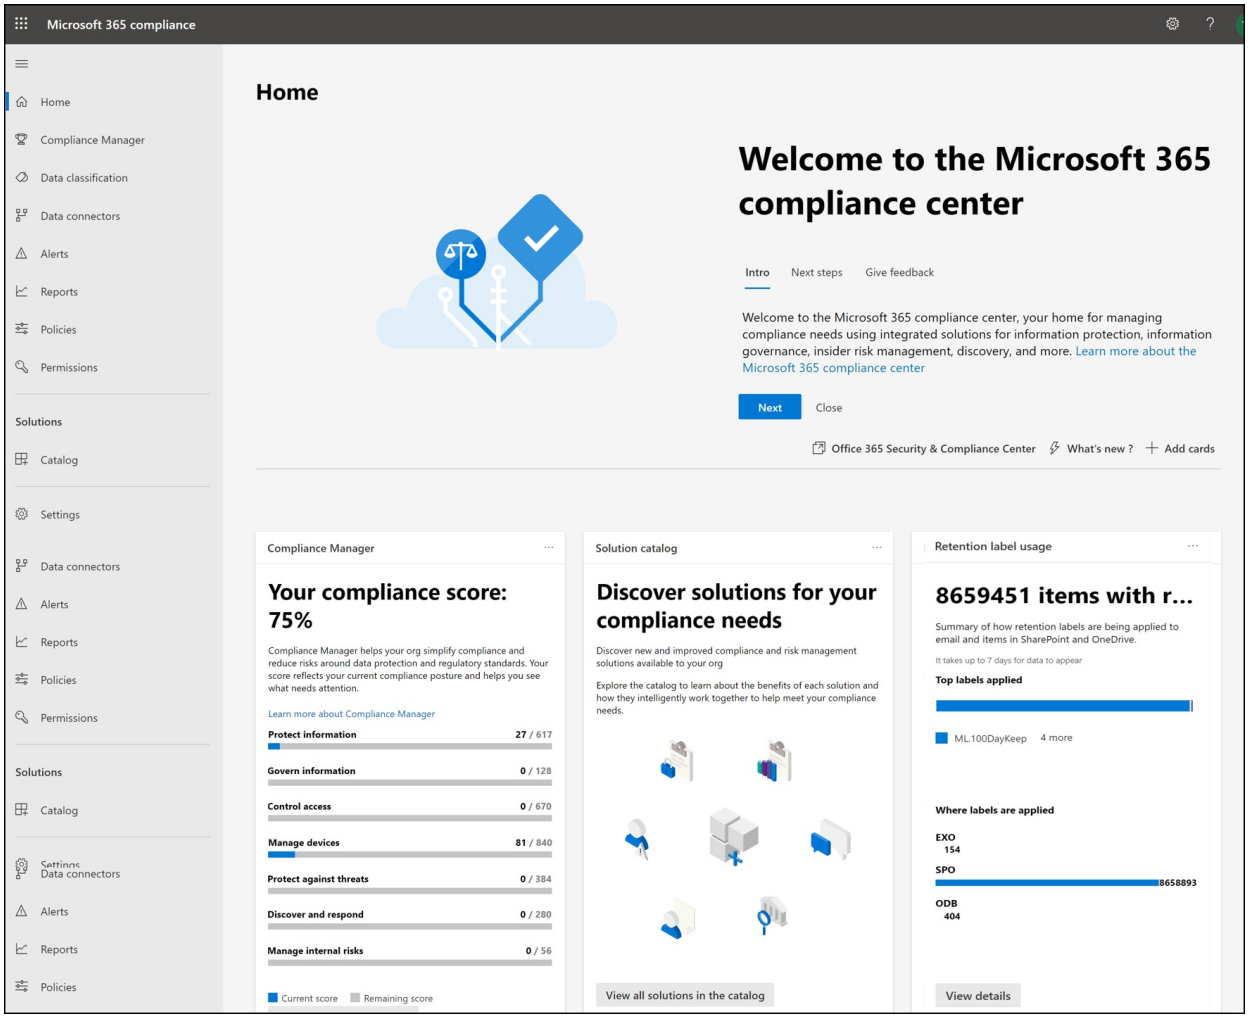

Welcome to Microsoft 365 compliance

When you go to your Microsoft 365 compliance center for the first time, you're greeted with the following welcome message:

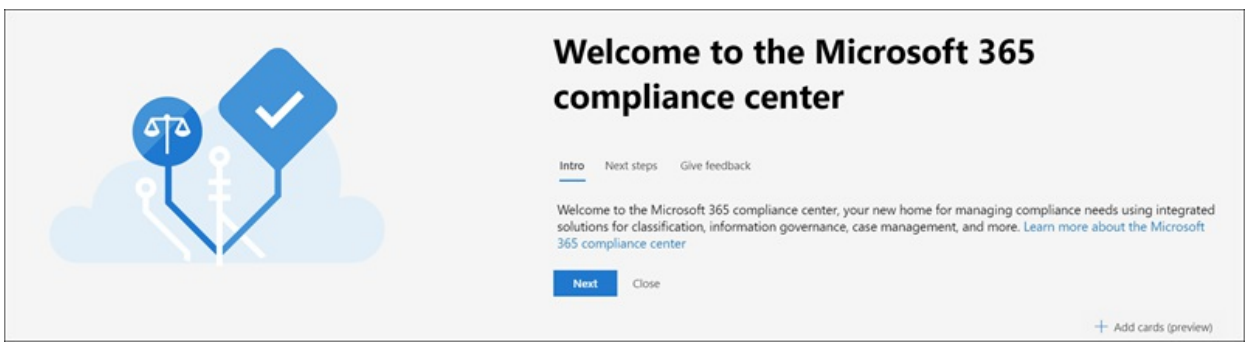

The welcome banner gives you some pointers on how to get started, with next steps, and an invitation for you to give us feedback.

Card section

When you first visit the Microsoft 365 compliance center, the card section on the home page shows you at a glance how your organization is doing with data compliance, what solutions are available for your organization, and a summary of any active alerts.

From here, you can:

- Review the **Microsoft Compliance Manager** card, which leads you to the Compliance Manager solution. Compliance Manager helps simplify the way you manage compliance. It calculates a risk-based score measuring your progress toward completing recommended actions that help reduce risks around data protection and regulatory standards. It also provides workflow capabilities and built-in control mapping to help you efficiently carry out improvement actions.

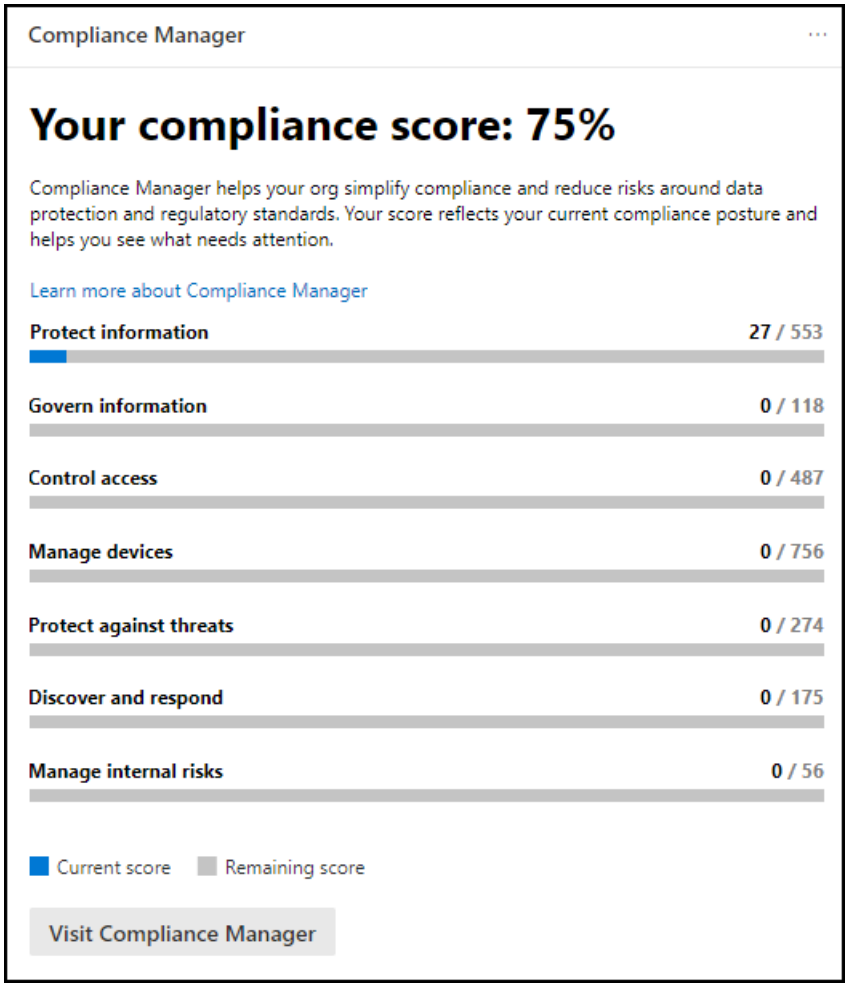

- Review the new **Solution catalog** card, which links to collections of integrated solutions you can use to help you manage end-to-end compliance scenarios. A solution's capabilities and tools might include a combination of policies, alerts, reports, and more.

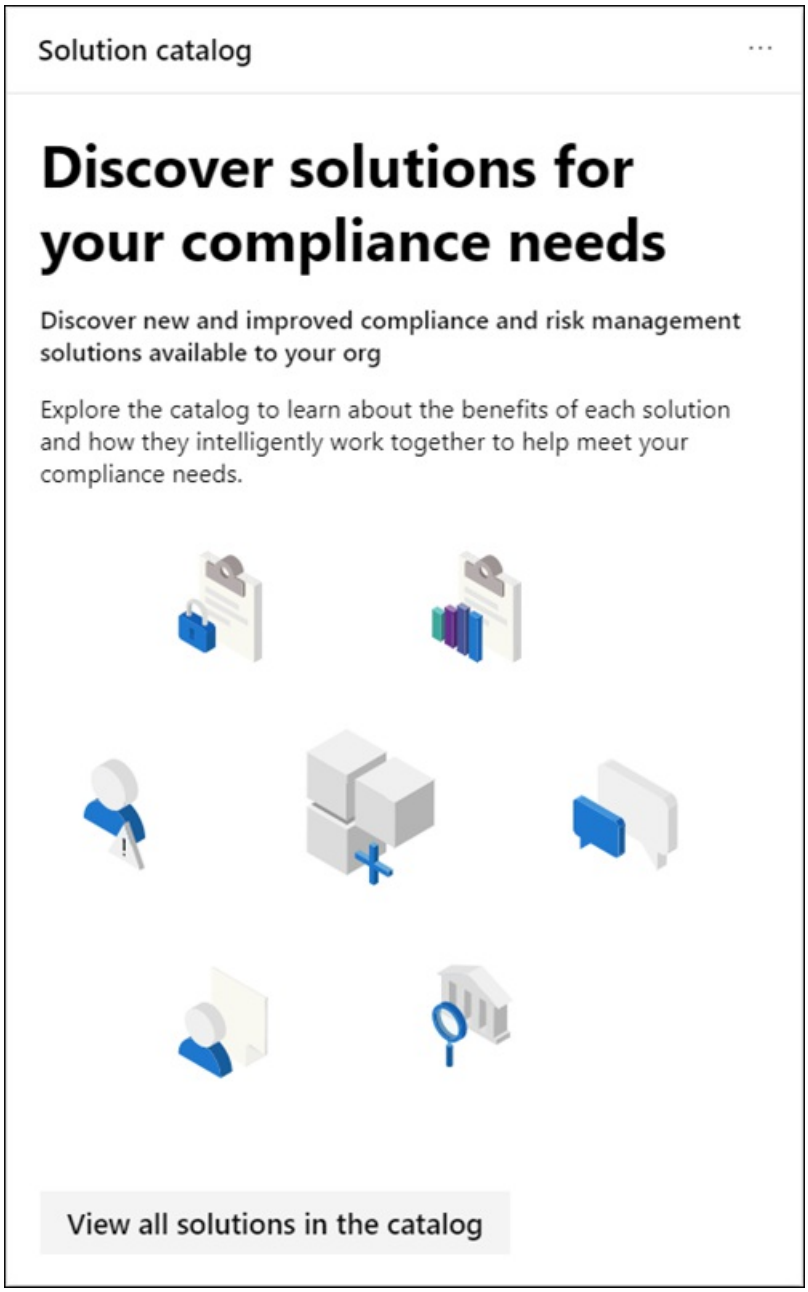

- Review the **Active alerts** card, which includes a summary of the most active alerts and includes a link where you can view more detailed information, such as Severity, Status, Category, and more.

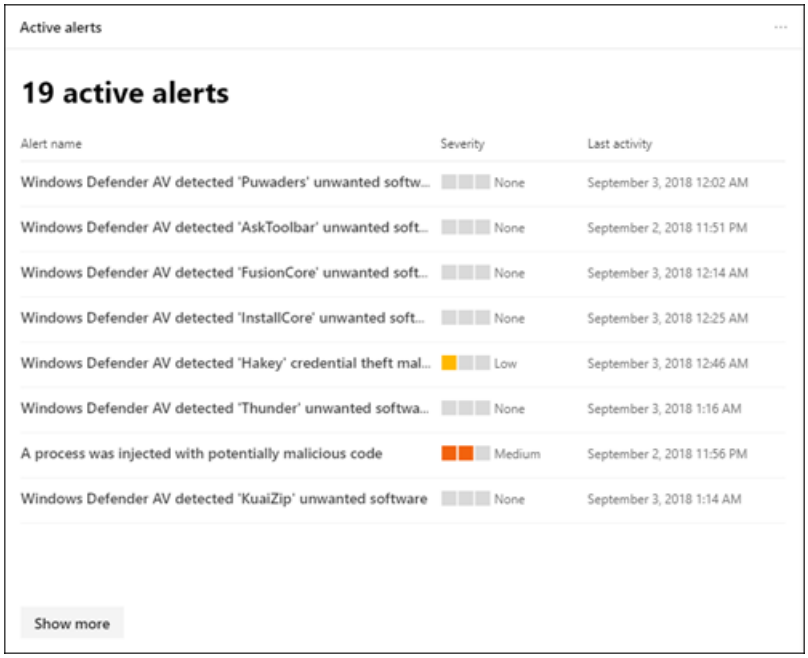

You can also use the **Add cards** feature to add additional cards, such as one showing your organization's cloud app compliance, and another showing data about users with shared files, with links to Cloud App Security or other tools where you can explore data.

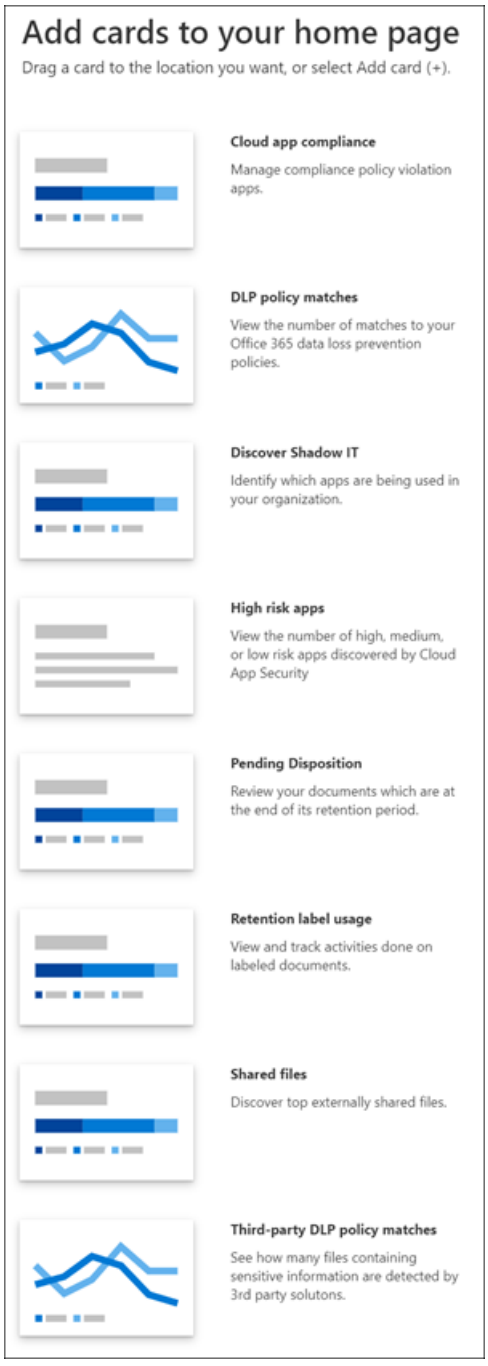

Easy navigation to more compliance features and capabilities

In addition to links in cards on the home page, you'll see a navigation pane on the left side of the screen that gives you easy access to your alerts, reports, policies, compliance solutions, and more. To add or remove options for a customized navigation pane, use the **Customize navigation** control on the navigation pane. This opens the **Customize your navigation pane** settings so you can configure which items appear in the navigation pane.

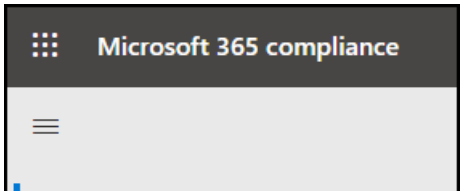	Select **Home** to return to the Microsoft 365 compliance center main page. Visit **Compliance Manager** to check your compliance score and start managing compliance for your organization.

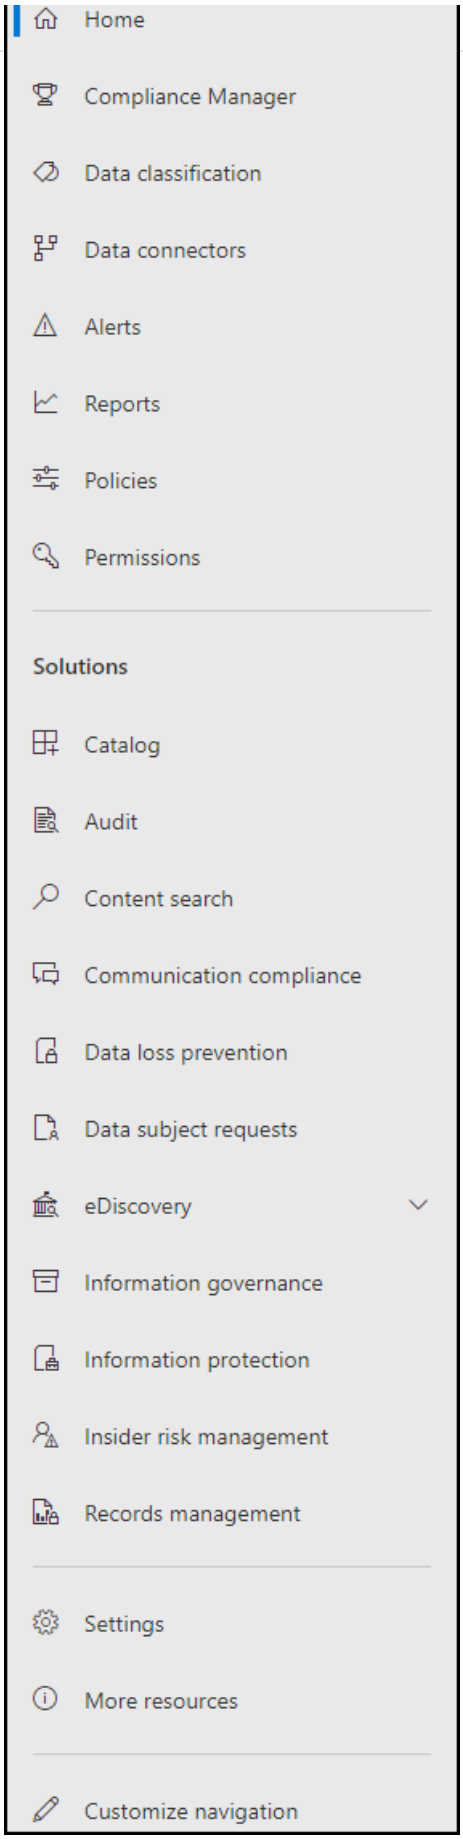

Select the **Data classification** section to access trainable classifiers, Sensitive information type entity definitions, content and activity explorers.

Select **Data connectors** to configure connectors to import and archive data in your Microsoft 365 subscription.

Go to **Alerts** to view and resolve alerts

Visit **Reports** to view data about label usage and retention, DLP policy matches and overrides, shared files, third-party apps in use, and more.

Go to **Policies** to set up policies to govern data, manage devices, and receive alerts. You can also access your DLP and retention policies.

Select **Permissions** to manage who in your organization has access to the Microsoft 365 compliance center to view content and complete tasks.

Use the links in the **Solutions** section to access your organization's compliance solutions. These include:

Catalog
Discover, learn about, and start using the intelligent compliance and risk management solutions available to your organization.

Audit
Use the Audit log to investigate common support and compliance issues.

Content search
Use Content search to quickly find email in Exchange mailboxes, documents in SharePoint sites and OneDrive locations, and instant messaging conversations in Microsoft Teams and Skype for Business.

Communication compliance
Minimize communication risks by automatically capturing inappropriate messages, investigating possible policy violations, and taking steps to remediate.

Data loss prevention
Detect sensitive content as it's used and shared throughout your organization, in the cloud and on devices, and helps prevent accidental data loss.

Data subject requests
Find and export a user's personal data to help you respond to data subject requests for the General Data Protection Regulation (GDPR).

eDiscovery
Expand this section to use the core and Advanced eDiscovery for preserving, collecting, reviewing, analyzing, and exporting content that's responsive to your organization's internal and external investigations.

Information governance
Manage your content lifecycle using features to import, store, and classify business-critical data so you can keep what you need and delete what you don't.

	Information protection Discover, classify, and protect sensitive and business-critical content throughout its lifecycle across your organization. Insider risk management Detect risky activity across your organization to help you quickly identify, investigate, and take action on insider risks and threats. Records management Automate and simplify the retention schedule for regulatory, legal and business-critical records in your organization.

How do I get the compliance center?

- If you don't have the new Microsoft 365 compliance center already, you'll have it soon. The Microsoft 365 compliance center is generally available now to Microsoft 365 SKU customers.
- To visit the Microsoft 365 compliance center, as a global administrator, compliance administrator, or compliance data administrator go to https://compliance.microsoft.com and sign in.

Frequently asked questions

Why am I taken to the Security & Compliance Center to complete some tasks, such as defining certain policies?

We're still developing the Microsoft 365 compliance center, and we add more functionality and solutions over the coming months. In the meantime, there are a few tasks that must be completed in the Security & Compliance Center (https://protection.office.com). In those cases, you'll be directed automatically to the location where you can complete the task at hand, such as creating or editing a supervision policy.

Why don't I see the new Microsoft 365 compliance center yet?

First, make sure that you have the appropriate licenses and permissions. Then, sign in at https://compliance.microsoft.com. If you don't see the new compliance center yet, you'll have it soon.

Some of my compliance features aren't available in the Microsoft 365 compliance center. What do I do?

We're still adding functionality to the Microsoft 365 compliance center. If you can't find something, such as audit log search, use the Security & Compliance Center (https://protection.office.com). Your configurations are saved in both the existing Security & Compliance Center and in the new Microsoft 365 compliance center automatically.

To go there, in the Microsoft 365 compliance center, in the navigation pane on the left side of the screen, choose **More resources**, and then, under **Office 365 Security & Compliance Center**, choose **Open**.

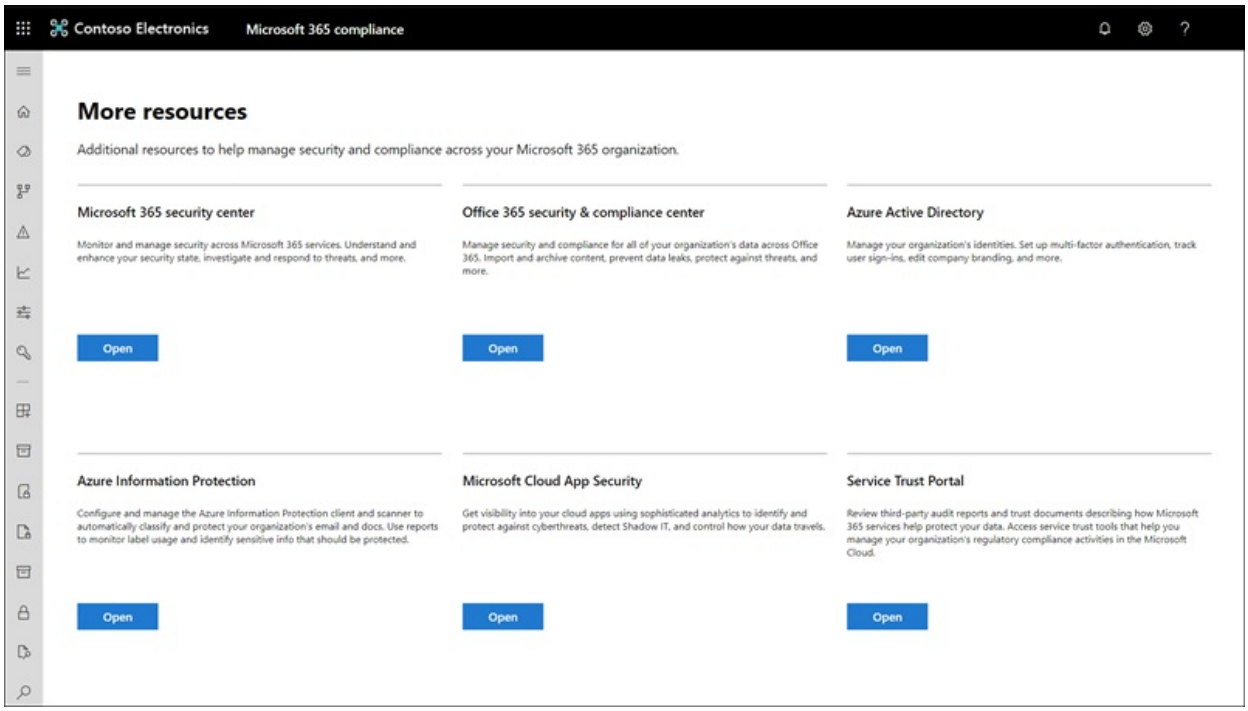

Next steps

- **Visit Microsoft Compliance Manager** to see your compliance score and start managing compliance for your organization. To learn more, see Compliance Manager.
- **Configure insider risk management policies** to help minimize internal risks and enable you to detect, investigate, and take action for risky activities in your organization. See Insider risk management.
- **Review your organization's data loss prevention policies** and make required changes as necessary. To learn more about, see Overview of data loss prevention policies.
- **Get acquainted with and set up Microsoft Cloud App Security**. See Quickstart: Get started with Microsoft Cloud App Security.
- **Learn about and create communication compliance policies** to quickly identify and remediate corporate code-of-conduct policy violations. See Communication compliance in Microsoft 365.
- **Visit your Microsoft 365 compliance center often**, and make sure to review any alerts or potential risks that arise. Go to https://compliance.microsoft.com and sign in.

Microsoft 365 solution catalog

2/18/2021 • 4 minutes to read • Edit Online

Are you looking for a way to quickly get started with compliance tasks in Microsoft 365? Check out the Microsoft 365 solution catalog to discover, learn, and quickly get started with compliance and risk management solutions.

Compliance solutions in Microsoft 365 are collections of integrated capabilities you can use to help you manage end-to-end compliance scenarios. A solution's capabilities and tools might include a combination of policies, alerts, reports, and more.

Read this article to get acquainted with the new solution catalog in the Microsoft 365 compliance center, how to get it, frequently asked questions, and your next steps.

Catalog organization

The solution catalog is organized into sections that contain information cards for each compliance solution available in your Microsoft 365 subscription. Each section contains cards for solutions grouped by compliance area.

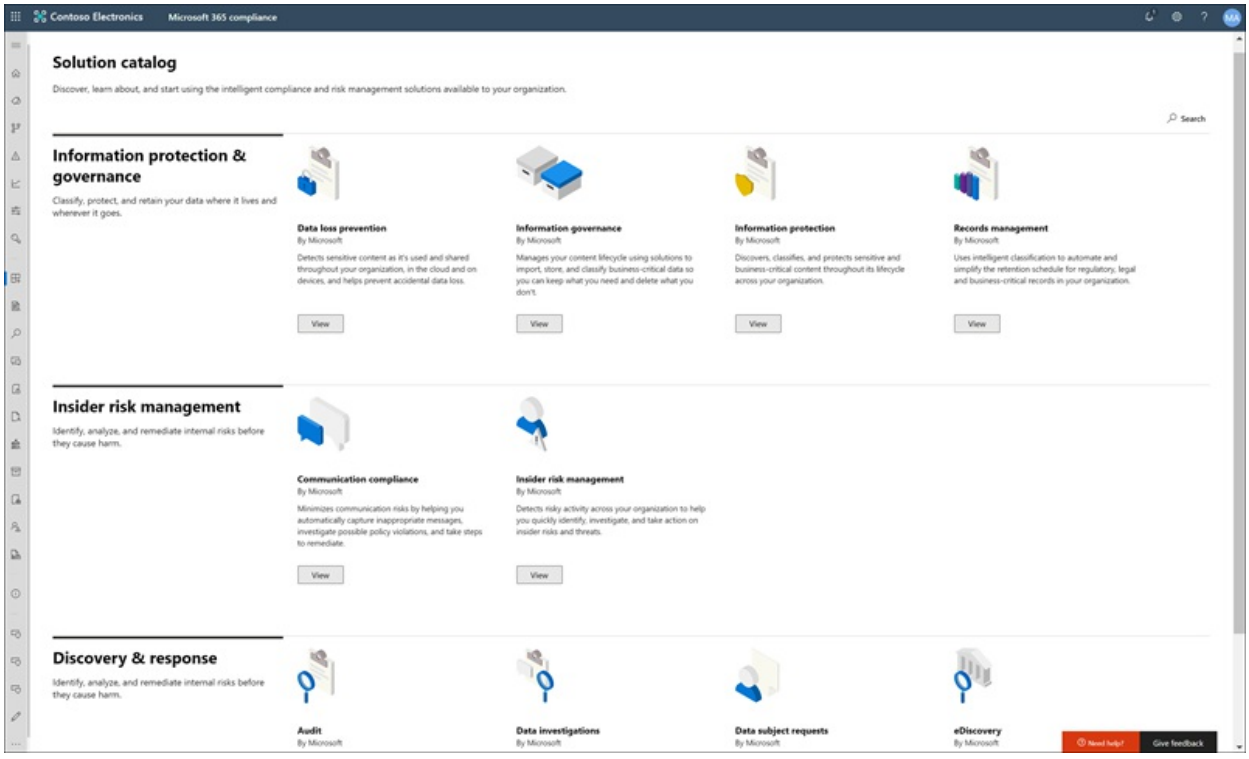

When you select **View** for a solution card, you'll see detailed information about the compliance solution and how to get started. This information includes an overview, pre-configuration requirements, learning resources, controls that allow you to pin the card to the navigation pane, and an option to share the solution as a link, email, or Microsoft Teams message.

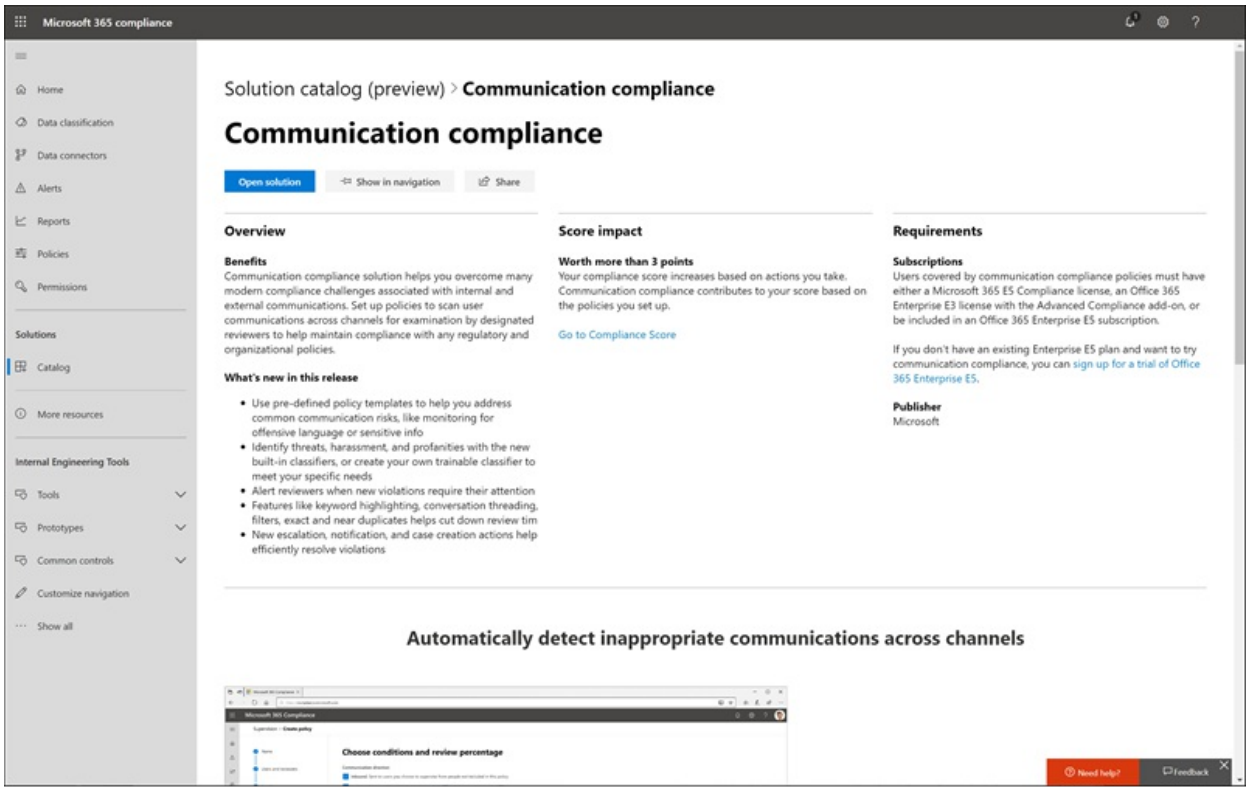

Information protection & governance section

The **Information protection & governance** section shows you at a glance how you can use Microsoft 365 compliance solutions to protect and govern data in your organization.

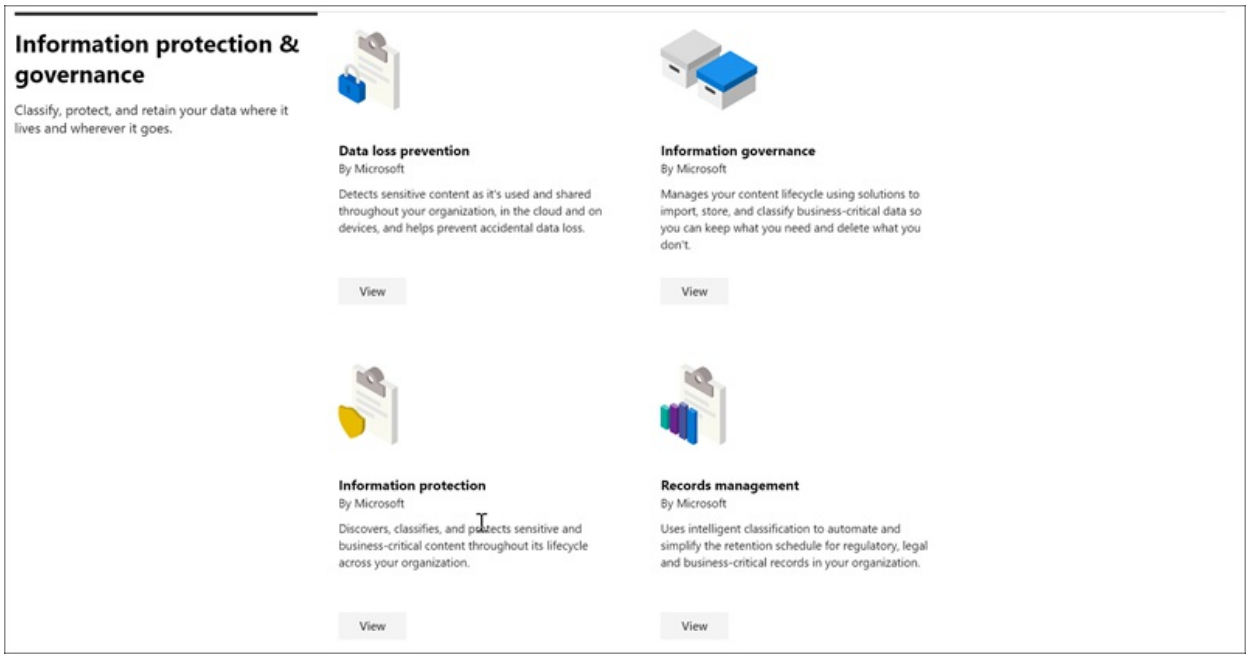

From here, you'll see cards for the following solutions:

- Data loss prevention: Detects sensitive content as it's used and shared throughout your organization, in the cloud and on devices, and helps prevent accidental data loss.
- Information governance: Manages your content lifecycle using solutions to import, store, and classify business-critical data so you can keep what you need and delete what you don't.
- Information protection: Discovers, classifies, and protects sensitive and business-critical content throughout its lifecycle across your organization.
- Records management: Uses intelligent classification to automate and simplify the retention schedule for regulatory, legal, and business-critical records in your organization.

Insider risk management section

The **Insider risk management** section on the home page shows you at a glance how your organization can identify, analyze, and take action on internal risks before they cause harm.

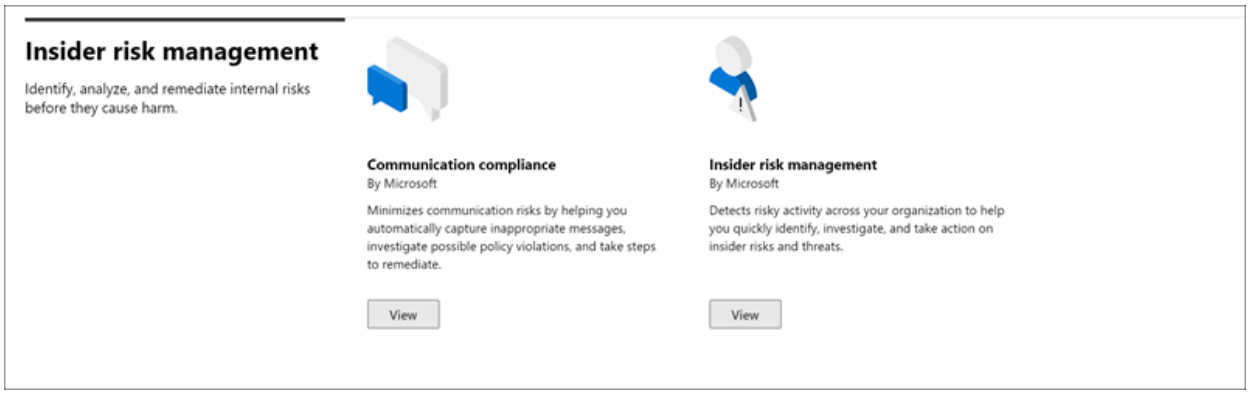

From here, you'll see cards for the following solutions:

- Communication compliance: Minimizes communication risks by helping you automatically capture inappropriate messages, investigate possible policy violations, and take steps to minimize harm.
- Insider risk management: Detect risky activity across your organization to help you quickly identify, investigate, and take action on insider risks and threats.

Discovery & response section

The **Discovery & response** section on the home page shows you at a glance how your organization can quickly find, investigate, and respond to compliance issues with relevant data.

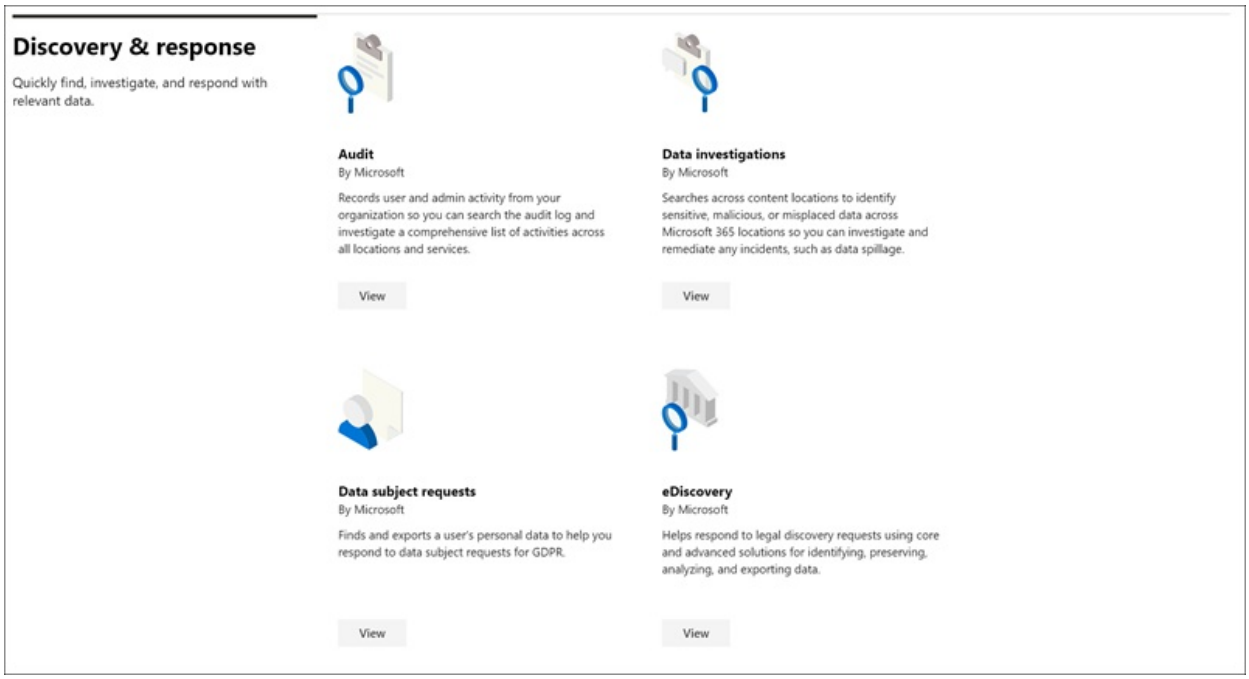

From here, you'll see cards for the following solutions:

- Audit: Records user and admin activity from your organization so you can search the audit log and investigate a comprehensive list of activities across all locations and services.
- Data subject requests: Finds and exports a user's personal data to help you respond to data subject requests for GDPR.
- eDiscovery
  - Core eDiscovery: Searches across content locations to identify, preserve, and export data in response to legal discovery requests and eDiscovery cases.

- Advanced eDiscovery: Builds on eDiscovery capabilities by providing intelligent analytics and machine learning to help you further analyze data that's relevant to discovery requests.

How do I get this?

To visit the Microsoft 365 solution catalog, go to https://compliance.microsoft.com and sign in as a global administrator, compliance administrator, or compliance data administrator. Select **Catalog** in the navigation pane on the left side of the screen to open the catalog home page.

Frequently asked questions

Why don't I see the Microsoft 365 solution catalog?

First, make sure that you have the appropriate licenses and permissions. Then, sign in at https://compliance.microsoft.com as a global administrator, compliance administrator, or compliance data administrator.

Some of the compliance features listed on the solution catalog page aren't available in the Microsoft 365 compliance center. What do I do?

We're always working to add new functionality to the Microsoft 365 compliance center and the solution catalog. If you can't find a specific solution in the navigation area, it will be accessible when the solution is available in your subscription.

If you are looking for an existing compliance solution and it's not available in the Microsoft 365 compliance center yet, you can always access solutions in the existing Security & Compliance Center by going to https://protection.office.com. Alternatively, you can click on the **More resources** tab in the left navigation of the Microsoft 365 compliance center and select the Office 365 security and compliance center card.

Next steps

- **Visit Microsoft Compliance Manager**, which helps you understand your organization's state of compliance with key standards and regulations. It provides recommended actions you can take to strengthen your overall compliance posture, and provides workflow capabilities to help you efficiently carry out those actions. To learn more, see Compliance Manager.
- **Configure insider risk management policies** to help minimize internal risks and enable you to detect, investigate, and take action for risky activities in your organization. See Insider risk management.
- **Learn about and create Communication compliance policies** to quickly identify and remediate corporate code-of-conduct policy violations. See Communication compliance.
- **Microsoft Information Protection**, learn how Microsoft 365 solutions help you discover, classify, and protect sensitive information wherever it lives or travels.
  - **Get acquainted with and set up Microsoft Cloud App Security**. See Quickstart: Get started with Microsoft Cloud App Security.
  - **Get started with classifiers**. Classifying content and then labeling it so it can be protected and handled properly is the starting place for the information protection discipline. See Learn about trainable classifiers (preview).
- **Visit your Microsoft 365 solution catalog often**, and make sure to review new solutions to help you with your compliance needs. Sign in at https://compliance.microsoft.com and then select **Catalog** in the left navigation pane.

Use your free Azure Active Directory subscription

11/2/2020 • 2 minutes to read • Edit Online

If your organization has a paid subscription to Microsoft 365, Microsoft Dynamics CRM Online, Enterprise Mobility Suite, or other Microsoft services, you have a free subscription to Microsoft Azure Active Directory. You and other admins can use Azure AD to create and manage user and group accounts. To use Azure AD, just go to the Azure portal and sign in to your account.

Open a private browsing session

Use a private browsing session (not a regular session) to access the Azure portal (in step 1 below). This prevents the credentials that you're currently logged on with from being passed to Azure. To open a private browsing session:

- In Microsoft Edge (legacy version), Internet Explorer, or Mozilla FireFox, press `CTRL+SHIFT+P`.
- In Microsoft Edge (newest version) or Google Chrome, press `CTRL+SHIFT+N`.

Access Azure Active Directory

1. Go to portal.azure.com and sign in with your work or student account.
2. In the left navigation pane in the Azure portal, click **Azure Active Directory**.

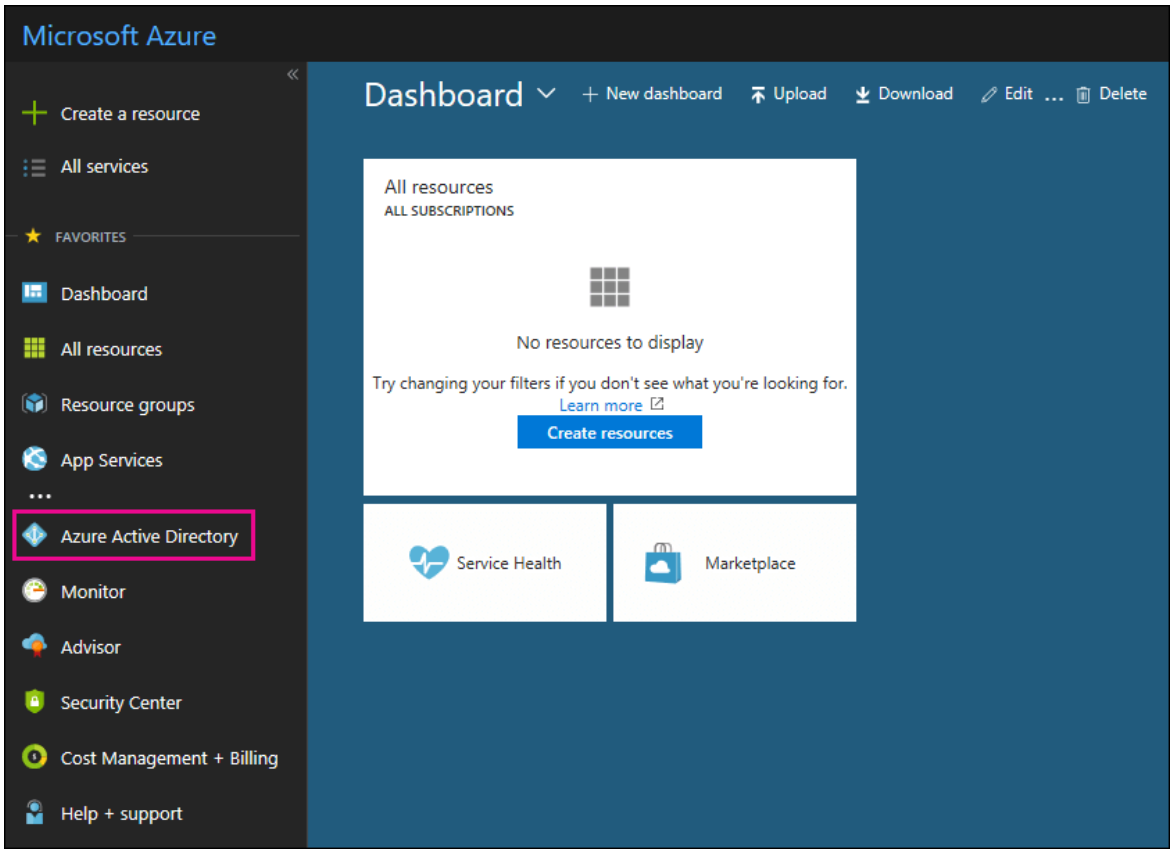

The **Azure Active Directory** admin center is displayed.

More information

- A free Azure Active Directory subscription does not include the Sign-ins activity report. To record sign-in

activity (which can be useful in a data breach), you need an Azure Active Directory Premium subscription. For more information, see How long does Azure AD store the data?.

- You can also access the **Azure Active Directory** admin center from the Microsoft 365 admin center. In the left navigation pane of the Microsoft 365 admin center, click **Admin centers** > **Azure Active Directory**.
- For information about managing users and groups and performing other directory management tasks, see Manage your Azure AD directory.

Plan for security & compliance

11/2/2020 • 2 minutes to read • Edit Online

Managing security and compliance is a partnership. You are responsible for protecting your data, identities, and devices, while Microsoft vigorously protects Microsoft 365 services. You can use Microsoft 365 and Enterprise Mobility + Security (EMS) together to help you achieve the appropriate level of protection for your organization.

Step 1: Review capabilities

Orient yourself to the information protection capabilities in the Information Protection for Office 365 poster.

Deploy information protection for data privacy regulations with Microsoft 365

Step 2: Check your Secure Score

After setting up your Microsoft 365 subscription, take note of your starting score. Secure Score provides configuration suggestions that you can take to increase your score. The goal is to be aware of opportunities that you can take to protect your environment which won't negatively affect the productivity of your users.

- Introducing the Office 365 Secure Score

Step 3: Plan access protection for identity and devices

Protecting access to your Microsoft 365 data and services is crucial to defending against cyberattacks and guarding against data loss.

- Protect access to data and services in Office 365
- Secure email policies and configurations

PDF | Visio | More languages

Step 4: Plan data protection based on data sensitivity

Review and plan for file protection capabilities organized by three levels of protection.

PDF | Visio

Step 5: Leverage the Microsoft 365 Security & Compliance Center

The Security & Compliance Center gives you a single view into the controls you will use to manage the spectrum of Microsoft 365 security, including threat management, data governance, and search and investigation.

- Go to the Security & Compliance Center
- Permissions in the Security & Compliance Center
- Give users access to the Security & Compliance Center

Step 6: Use end-to-end security scenarios as starting points

Use these recommended configurations as a starting point for enterprise scale or sophisticated access security

scenarios.

- Secure email policies and configurations
- Contoso in the Microsoft Cloud

Microsoft 365 admin centers and dashboards

Configure your security and compliance settings in these admin centers and dashboards to protect your Microsoft 365 environment

SUBSCRIPTION	MANAGEMENT URL	DASHBOARDS AND ADMIN CENTERS
Microsoft 365	`https://admin.microsoft.com`	Microsoft 365 admin center Security & Compliance Center Exchange admin center SharePoint admin center and OneDrive for Business admin center
Enterprise Mobility + Security	`https://portal.azure.com`	Azure Active Directory Microsoft Mobile Application Management Microsoft Intune
Enterprise Mobility + Security	`https://portal.cloudappsecurity.com`	Cloud App Security

Microsoft 365 compliance extensibility

11/2/2020 • 4 minutes to read • Edit Online

Microsoft 365 compliance solutions help organizations intelligently assess their compliance risks, govern and protect sensitive data, and effectively respond to regulatory requirements. Microsoft 365 compliance is rich in extensibility scenarios and enables organizations to adapt, extend, integrate, accelerate, and support their compliance solutions.

There are two key building blocks for compliance extensibility:

- **Data connectors**. Use to import and archive non-Microsoft data so you can apply Microsoft 365 protection and governance capabilities to third-party data.
- **APIs**. Enables programmatic access to Microsoft 365 compliance capabilities.

Data connectors

Microsoft provides third-party data connectors that can be configured in the Microsoft 365 compliance center. For a list of data connectors provided by Microsoft, see the Third-party data connectors table. The table of third-party data connectors also summarizes the compliance solutions that you can apply to third-party data after you import and archive data in Microsoft 365, and links to the step-by-step instructions for each connector.

To learn more about Microsoft 365 data connectors, see Archiving third-party data. If a third-party data type isn't supported by the data connectors available in the Microsoft 365 compliance center, you can work with a partner who can provide you with a custom connector. For a list of partners you can work with and the step-by-step process for this method, see Work with a partner to archive third-party data.

Prerequisites for data connectors

Many of the data connectors available in the Microsoft 365 compliance center to import and archive third-party data require that you prepare and perform configuration tasks in the third-party data source. These prerequisites are documented in detail for each third-party data connector.

For data connectors in the Microsoft 365 compliance center provided by one of Microsoft's partners, your organization will need a business relationship with the partner before you can deploy a connector.

You can find licensing requirements for third-party data connectors in the Microsoft 365 Compliance Licensing Comparison document.

APIs

Microsoft 365 compliance APIs are available in the Microsoft Information Protection SDK, Microsoft Graph API, and the Office 365 Management Activity API. Some compliance APIs are part of a new set of security and compliance APIs that enable developers for Microsoft 365 customers, independent software vendors, system integrators, and managed security service providers to build high-value security and compliance solutions.

To learn more about how to access Graph APIs, see Overview of Microsoft Graph.

Microsoft Information Protection (MIP) SDK

The MIP SDK exposes the labeling and protection services from Microsoft 365 security and compliance centers to third-party applications and services. Developers can use the SDK to build native support for applying labels and protection to files. Developers can determine which actions should be taken when specific labels are detected, and reason over MIP-encrypted information.

High-level MIP SDK use cases include:

- A line-of-business application that applies classification labels to files on export.
- A CAD/CAM design application that provides native support for MIP labeling.
- A cloud access security broker or data loss prevention solution that can encrypt data with Azure Information Protection.

To learn more about the MIP SDK, prerequisites, additional scenarios, and samples, see MIP SDK Overview.

Microsoft Graph API for Teams DLP

Data loss prevention (DLP) capabilities are widely used in Microsoft Teams particularly as organizations have shifted to remote work. Earlier this year we announced the public preview of the Microsoft Graph Change Notification API for messages in Teams. This API enables developers to build apps that can listen to Microsoft Teams messages in near-real time and then implement DLP scenarios for both customers and partners. Additionally, Microsoft Graph Patch API lets you apply DLP actions to Teams messages.

These two APIs form the Microsoft Graph API for Teams DLP. You can get started by trying out the sample app. For more information about Microsoft Teams messaging webhooks, see the documentation.

For the licensing requirements for Teams DLP, see Microsoft 365 licensing guidance for security & compliance.

Microsoft Graph API for eDiscovery (preview)

With Advanced eDiscovery, organizations can discover data where it lives, and manage more end-to-end eDiscovery workflows with intelligent machine-learning and analytics capabilities to reduce data to the relevant set – all while the data stays within the Microsoft 365 security and compliance boundary.

Graph APIs for Advanced eDiscovery can be used to create and manage cases, review sets, and review set queries in a scalable and repeatable manner. This enables customers and partners to create apps and workflows to automate common and repetitive processes such as creating cases and managing custodians and legal holds.

The first set of Graph APIs for eDiscovery are available in public preview. We plan to add more capabilities by the end of the calendar year. To learn more about these APIs and other updates for Advanced eDiscovery, see this blog.

For the licensing requirements for Advanced eDiscovery and the API, see the "eDiscovery" section in the Microsoft 365 licensing guidance for security & compliance.

Microsoft Graph API for Teams Export (preview)

Enterprise Information Archiving (EIA) for Microsoft Teams is a key scenario for our customers as it allows them to solve for regulatory requirements. In addition to our built-in capabilities for archiving content in Microsoft Teams, customers and partners can now use Teams Export APIs to solve for custom application and integration scenarios. The Teams Export APIs support bulk-export (up to 200 requests per second/per app/per tenant) of Teams messages and message attachments. Deleted messages are also accessible by the API for up to 30 days after they are deleted. For more information about these Teams Export APIs and how to use them in your applications, see Export content with the Microsoft Teams Export APIs.

For the licensing requirements for the use of the Teams Export APIs, see Microsoft 365 licensing guidance for security & compliance.

Insider risk solutions in Microsoft 365

2/18/2021 • 6 minutes to read • Edit Online

Insider risks are one of the top concerns of security and compliance professionals in the modern workplace. Industry studies have shown that insider risks are often associated with specific user events or activities. Protecting your organization against these risks can be challenging to identify and difficult to mitigate. Insider risks include vulnerabilities in a variety of areas and can cause major problems for your organization, ranging from the loss of intellectual property to workplace harassment, and more. The following figure outlines common insider risks:

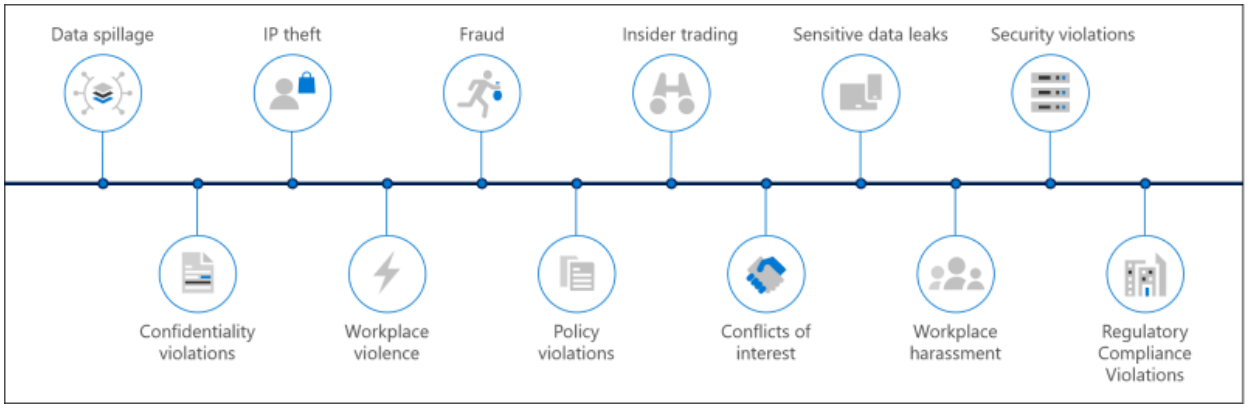

Microsoft 365 risk prevention features are designed and built-in to our insider risk products and solutions. These solutions work together and use advanced service and 3rd-party indicators to help you quickly identify, triage, and act on risk activity. Most solutions offer a comprehensive detection, alert, and remediation workflow for your data analysts and investigators to use to quickly act on and minimize these risks.

RISK ICON	RISKS	COMMUNICATION COMPLIANCE	INSIDER RISK MANAGEMENT	INFORMATION BARRIERS	PRIVILEGED ACCESS MANAGEMENT
	Data spillage	✔	✔		
	Confidentiality violations	✔	✔	✔	
	IP theft	✔	✔	✔	
	Workplace violence	✔			
	Fraud	✔	✔		
	Policy violations	✔	✔	✔	✔
	Insider trading	✔			

RISK ICON	RISKS	COMMUNICATION COMPLIANCE	INSIDER RISK MANAGEMENT	INFORMATION BARRIERS	PRIVILEGED ACCESS MANAGEMENT
	Conflicts of interest	✔		✔	
	Sensitive data leaks	✔	✔		
	Workplace harassment	✔			
	Security violations		✔		✔
	Regulatory compliance violations	✔	✔	✔	

Microsoft 365 insider risk solutions

To help protect your organization against insider risks, use these Microsoft 365 capabilities and features.

Communication compliance

Communication compliance helps minimize communication risks by helping you detect, capture, and act on inappropriate messages in your organization. Communication compliance is available in the following subscriptions:

- Microsoft 365 E5 subscription (paid or trial version)
- Microsoft 365 E3 subscription + the Microsoft 365 E5 Compliance add-on
- Microsoft 365 E3 subscription + the Microsoft 365 E5 Insider Risk Management add-on
- Microsoft 365 A5 subscription (paid or trial version)
- Microsoft 365 A3 subscription + the Microsoft 365 A5 Compliance add-on
- Microsoft 365 A3 subscription + the Microsoft 365 A5 Insider Risk Management add-on
- Microsoft 365 G5 subscription (paid or trial version)
- Microsoft 365 G5 subscription + the Microsoft 365 G5 Compliance add-on
- Microsoft 365 G5 subscription + the Microsoft 365 G5 Insider Risk Management add-on
- Office 365 Enterprise E5 subscription (paid or trial version)
- Office 365 A5 subscription (paid or trial version)
- Office 365 Enterprise E3 subscription + the Office 365 Advanced Compliance add-on (no longer available for new subscriptions)

Insider risk management

Insider risk management helps minimize internal risks by enabling you to detect, investigate, and act on malicious and inadvertent activities in your organization.

Insider risk management is available in the following subscriptions:

- Microsoft 365 E5 subscription (paid or trial version)
- Microsoft 365 E3 subscription + the Microsoft 365 E5 Compliance add-on
- Microsoft 365 E3 subscription + the Microsoft 365 E5 Insider Risk Management add-on
- Microsoft 365 A5 subscription (paid or trial version)

- Microsoft 365 A3 subscription + the Microsoft 365 A5 Compliance add-on
- Microsoft 365 A3 subscription + the Microsoft 365 A5 Insider Risk Management add-on
- Microsoft 365 G5 subscription (paid or trial version)
- Microsoft 365 G3 subscription + the Microsoft 365 G5 Compliance add-on
- Microsoft 365 G3 subscription + the Microsoft 365 G5 Insider Risk Management add-on
- Office 365 E3 subscription + Enterprise Mobility and Security E3 + the Microsoft 365 E5 Compliance add-on

Information barriers

Information barriers allow you to restrict communication and collaboration between two internal groups to avoid a conflict of interest from occurring in your organization.

Information barriers are available in the following subscriptions:

- Microsoft 365 E5 subscription (paid or trial version)
- Microsoft 365 A5 subscription (paid or trial version)
- Office 365 Enterprise E5 subscription (paid or trial version)
- Office 365 A5 subscription (paid or trial version)
- Office 365 Advanced Compliance add-on (no longer available for new subscriptions)
- Microsoft 365 E3 subscription + the Microsoft 365 E5 Compliance add-on
- Microsoft 365 E3 subscription + the Microsoft 365 E5 Insider Risk Management add-on
- Microsoft 365 A3 subscription + the Microsoft 365 A5 Compliance add-on
- Microsoft 365 A3 subscription + the Microsoft 365 A5 Insider Risk Management add-on

Privileged access management

Privileged access management allows granular access control over privileged Exchange Online admin tasks in Office 365. It can help protect your organization from breaches that use existing privileged admin accounts with standing access to sensitive data or access to critical configuration settings.

Privileged access management is available in the following subscriptions:

- Microsoft 365 E5 subscription (paid or trial version)
- Microsoft 365 A5 subscription (paid or trial version)
- Office 365 Enterprise E5 subscription (paid or trial version)
- Office 365 A5 subscription (paid or trial version)
- Microsoft 365 E3 subscription + the Microsoft 365 E5 Compliance add-on
- Microsoft 365 E3 subscription + the Microsoft 365 E5 Information Protection and Governance add-on
- Microsoft 365 A3 subscription + the Microsoft 365 A5 Compliance add-on
- Microsoft 365 A3 subscription + the Microsoft 365 A5 Information Protection and Governance add-on

Deploy Microsoft 365 insider risk solutions

To help protect your organization against insider risks, set up and deploy the following Microsoft 365 solutions:

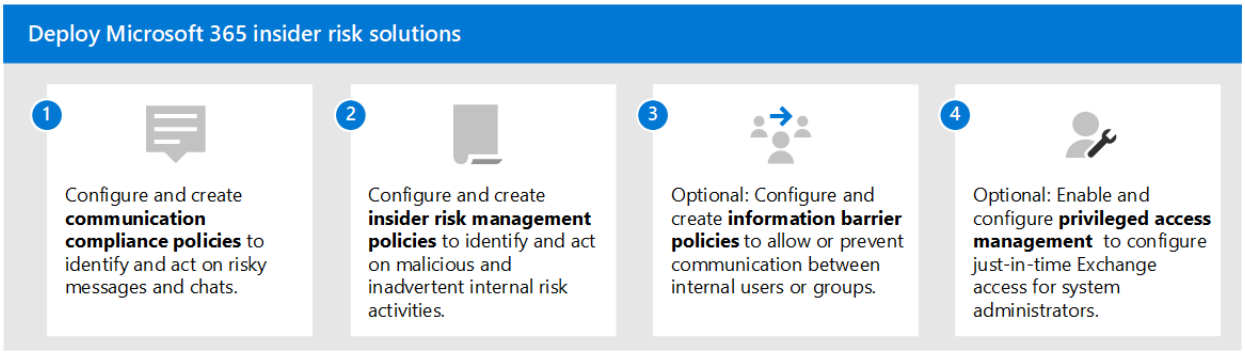

1. Configure and create communication compliance policies.
2. Configure and create insider risk management policies.
3. Optional: Configure and create information barrier policies.
4. Optional: Enable and configure privileged access management.

Illustrations with examples

To help you plan an integrated strategy for implementing Microsoft 365 insider risk capabilities, download the *Microsoft 365 information protection and compliance capabilities* set of illustrations. For insider risk capabilities, see the architecture illustration pages 5-7. Feel free to adapt these illustrations for your own use.

ITEM	DESCRIPTION
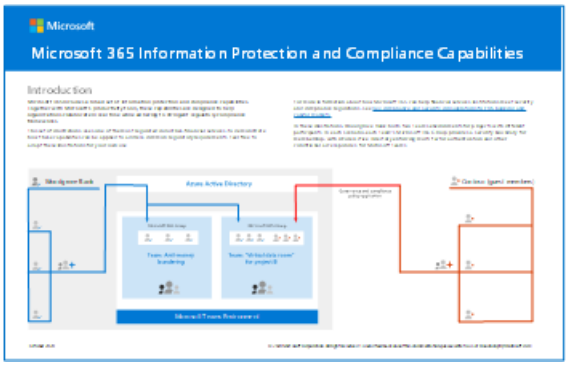 Download as a PDF | Download as a Visio Updated October 2020	Includes: • Microsoft information protection and data loss prevention • Retention policies and retention labels • Information barriers • Communication compliance • Insider risk management • Third-party data ingestion

Training

Training your administrators and compliance team in the basics for each insider risk solution can help your organization get started more quickly with your deployment and implementation efforts.

Microsoft 365 provides the following resources to help inform and train these users in your organization:

SOLUTION/AREA	RESOURCES
Manage insider risk in Microsoft 365	Complete learning path This learning path includes all the individual solution modules for communication compliance, insider risk management, information barriers, and privileged access management. Select this learning path to complete all the modules.
Communication compliance	Learning module: Prepare communication compliance in Microsoft 365 This module helps you learn the basics on how to identify and remediate code-of-conduct policy violations with communication compliance, cover the prerequisites needed before creating communication compliance policies, and learn about the types of built-in, pre-defined policy templates in communication compliance.

SOLUTION/AREA	RESOURCES
Insider risk management	Learning module: Insider risk management in Microsoft 365 This module helps you learn how insider risk management in Microsoft 365 can help prevent, detect, and contain internal risks in an organization, learn about the types of built-in, pre-defined policy templates, understand the basic prerequisites needed before creating insider risk policies, and explains the types of actions you can take on insider risk management cases.
Information barriers	Learning module: Plan for information barriers This module helps you learn how information barrier policies can help your organization maintain compliance with relevant industry standards and regulations, lists the types of situations when information barriers would be applicable, helps explain the process of creating an information barrier policy, and helps explain how to troubleshoot unexpected issues after information barriers are in place.
Privileged access management	Learning module: Implement privileged access management This module helps you understand the difference between privileged access management and privileged identity management, understand the privileged access management process flow, and understand the basics of how to configure and enable privileged access management.

Communication compliance in Microsoft 365

2/18/2021 • 2 minutes to read • Edit Online

Protecting sensitive information and detecting and acting on workplace harassment incidents is an important part of compliance with internal policies and standards. Communication compliance in Microsoft 365 helps minimize these risks by helping you quickly detect, capture, and take remediation actions for email and Microsoft Teams communications. These include inappropriate communications containing profanity, threats, and harassment and communications that share sensitive information inside and outside of your organization.

Configure communication compliance for Microsoft 365

Use the following steps to configure communication compliance for your organization:

1. Learn about communication compliance in Microsoft 365
2. Plan for communication compliance and verify licensing
3. Review feature reference and solution capabilities
4. Configure prerequisites and permissions
5. Create and configure communication compliance policies

More information about communication compliance

- Investigate and remediate alerts
- Case study - Contoso quickly configures an offensive language policy for Microsoft Teams, Exchange, and Yammer communications

Learn about communication compliance in Microsoft 365

2/18/2021 • 11 minutes to read • Edit Online

Communication compliance is an insider risk solution in Microsoft 365 that helps minimize communication risks by helping you detect, capture, and act on inappropriate messages in your organization. Pre-defined and custom policies allow you to scan internal and external communications for policy matches so they can be examined by designated reviewers. Reviewers can investigate scanned email, Microsoft Teams, Yammer, or third-party communications in your organization and take appropriate actions to make sure they're compliant with your organization's message standards.

Communication compliance policies in Microsoft 365 help you overcome many modern challenges associated with compliance and internal and external communications, including:

- Scanning increasing types of communication channels
- The increasing volume of message data
- Regulatory enforcement and the risk of fines

Additionally, there may be a separation of duties between your IT admins and your compliance management team. Communication compliance supports the separation between configuration of policies and the investigation and review of messages. For example, the IT group for your organization may be responsible for setting up communication compliance role permissions, groups, and policies and investigators and reviewers may be responsible for message triage, review, and mitigation actions.

For a quick overview of communication compliance, see the Detect workplace harassment and respond with Communication Compliance in Microsoft 365 video on the Microsoft Mechanics channel.

Scenarios for communication compliance

Communication compliance policies can assist with reviewing messages in your organization in several important compliance areas:

- **Corporate policies**

  Users must comply with acceptable use, ethical standards, and other corporate policies in all their business-related communications. Communication compliance policies can detect policy matches and help you take corrective actions to help mitigate these types of incidents. For example, you could scan user communications in your organization for potential human resources concerns such as harassment or the use of inappropriate or offensive language.

- **Risk management**

  Organizations are responsible to all communications distributed throughout their infrastructure and corporate network systems. Using communication compliance policies to help identify and manage potential legal exposure and risk can help minimize risks before they can damage corporate operations. For example, you could scan messages in your organization for unauthorized communications and conflicts of interest about confidential projects such as upcoming acquisitions, mergers, earnings disclosures, reorganizations, or leadership team changes.

- **Regulatory compliance**

  Most organizations must comply with some type of regulatory compliance standards as part of their

normal operating procedures. These regulations often require organizations to implement some type of supervisory or oversight process for messaging that is appropriate for their industry. The Financial Industry Regulatory Authority (FINRA) Rule 3110 is a good example of a requirement for organizations to have supervisory procedures in place to scan user communications and the types of businesses in which it engages. Another example may be a need to review broker-dealer communications in your organization to safeguard against potential money laundering, insider trading, collusion, or bribery activities. Communication compliance policies can help your organization meet these requirements by providing a process to both scan and report on corporate communications. For more information on support for financial organizations, see Key compliance and security considerations for US banking and capital markets.

Key feature areas

Communication compliance in Microsoft 365 offers several important features to help address compliance concerns on your messaging platforms:

- Intelligent customizable templates
- Flexible remediation workflows
- Actionable insights

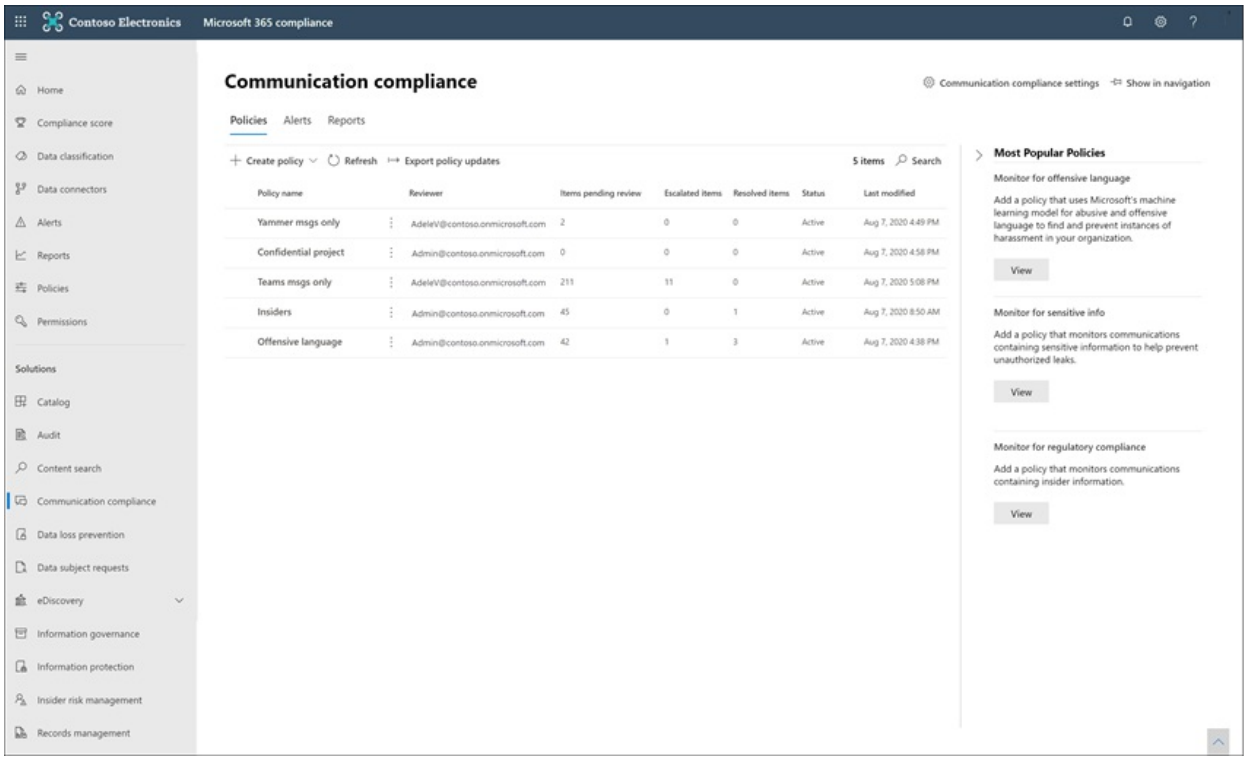

Intelligent customizable templates

Intelligent customizable templates in communication compliance allow you to apply machine learning to intelligently detect communication violations in your organization.

- **Customizable pre-configured templates**: New policy templates help address the most common communications risks. Initial policy creation and follow-on updating are now quicker with pre-defined anti-harassment and offensive language, sensitive information, conflict of interest, and regulatory compliance templates.
- **New machine learning support**: Built-in threat, harassment, profanity, and image classifiers help reduce false positives in scanned messages, saving reviewers time during the investigation and remediation process.
- **Improved condition builder**: Configuring policy conditions is now streamlined into a single, integrated experience in the policy wizard, reducing confusion in how conditions are applied for policies.

Flexible remediation workflows

Built-in remediation workflows allow you to quickly identify and take action on messages with policy matches in your organization. The following new features increase efficiency for investigation and remediation activities:

- **Flexible remediation workflow**: New remediation workflow helps you quickly take action on policy matches, including new options to escalate messages to other reviewers and to send email notifications to users with policy matches.
- **Conversation threading**: Messages are now visually grouped by original message and all associated reply messages, giving you better context during investigation and remediation actions.
- **Keyword highlighting**: Terms matching policy conditions are highlighted in the message text view to help reviewers quickly locate and remediate policy alerts.
- **Exact and near duplicate detection**: In addition to scanning for exact terms matching communication compliance policies, near duplicate detection groups textually similar terms and messages together to help speed up your review process.
- **New filters**: Investigate and remediate policy alerts faster with message filters for several fields, including sender, recipient, date, domains, and many more.
- **Improved message views**: Investigation and remediation actions are now quicker with new message source, text, and annotation views. Message attachments are now viewable to provide complete context when taking remediation actions.
- **User history view**: Historical view of all user message remediation activities, such as past notifications and escalations for policy matches, now provides reviewers with more context during the remediation workflow process. First-time or repeat instances of policy matches for users are now archived and easily viewable.
- **Pattern detected notification**: Many harassing and bullying actions take place over time and involve reoccurring instances of the same behavior by a user. The new Pattern detected notification displayed in alert details helps raise attention to these alerts and this type of behavior.
- **Show Translate view**: Quickly investigate message details in other languages using translate support in the remediation workflow. Messages in other languages are automatically converted to the display language of the reviewer.

Actionable insights

New interactive dashboards for alerts, policy matches, actions, and trends help you quickly view the status of pending and resolved alerts in your organization.

- **Proactive intelligent alerts**: Alerts for policy matches requiring immediate attention include new dashboards for pending items sorted by severity and new automatic email notifications sent to designated reviewers.
- **Interactive dashboards**: New dashboards display policy matches, pending and resolved actions, and trends by users and policy.
- **Auditing support**: A full log of policy and review activities is easily exported from the Microsoft 365 compliance center to help support audit review requests.

Integration with Microsoft 365 services

Communication compliance policies scan and capture messages across several communication channels to help you quickly review and remediate compliance issues:

- **Microsoft Teams**: Chat communications for public and private Microsoft Teams channels and individual chats are supported in communication compliance as a standalone channel source or with other Microsoft 365 services. You'll need to manually add individual users, distribution groups, or specific Microsoft Teams channels when you select users and groups to supervise in a communication compliance policy.
- **Exchange Online**: All mailboxes hosted on Exchange Online in your Microsoft 365 organization are eligible for scanning. Emails and attachments matching communication compliance policy conditions are instantly available for monitoring and in compliance reports. Exchange Online is now an optional source channel and

is no longer required in communication compliance policies.

- **Yammer**: Private messages and public community conversations in Yammer are supported in communication compliance policies. Yammer is an optional channel and must be in native mode to support scanning of messages and attachments.
- **Skype for Business Online**: Communication compliance policies support scanning chat communications and associated attachments in Skype for Business Online.
- **Third-party sources**: You can scan messages from third-party sources for data imported into mailboxes in your Microsoft 365 organization. Communication compliance supports connections to several popular platforms, including Instant Bloomberg and others.

To learn more about messaging channel support in communication compliance policies, see supported communication types.

Workflow

Communication compliance helps you address common pain points associated with complying with internal policies and regulatory compliance requirements. With focused policy templates and a flexible workflow, you can use actionable insights to quickly resolve detected compliance issues.

Identifying and resolving compliance issues with communication compliance in Microsoft 365 uses the following workflow:

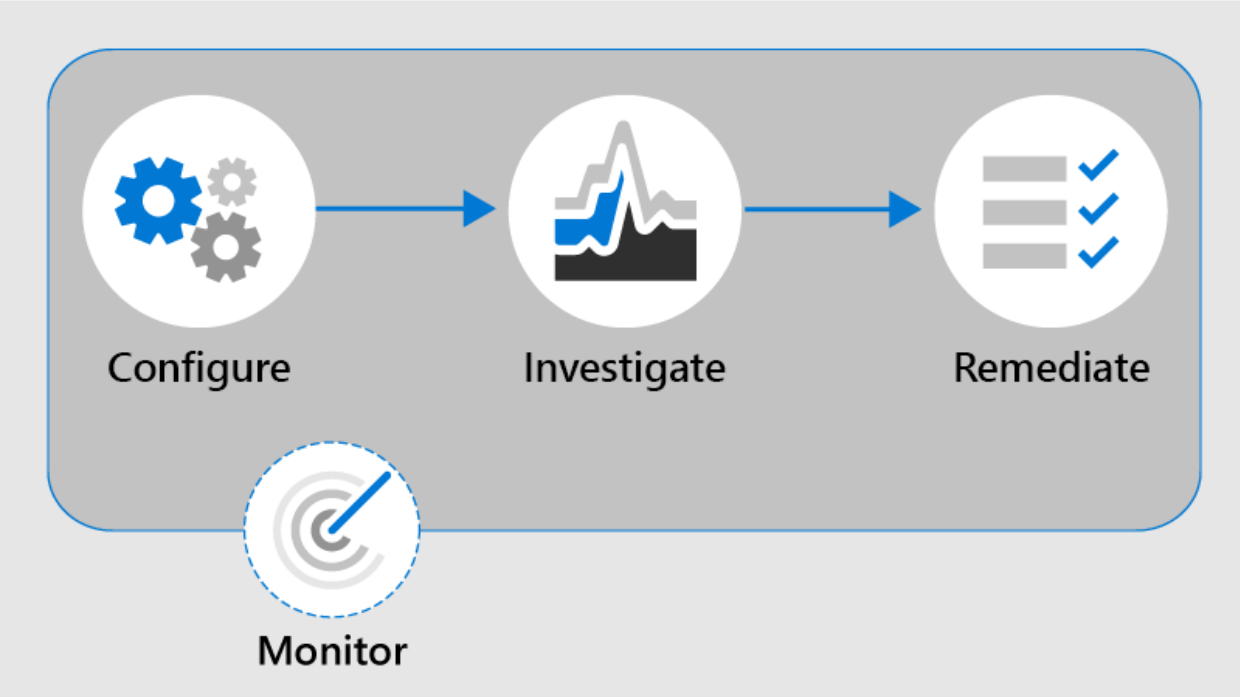

Configure

In this workflow step, you identify your compliance requirements and configure applicable communication compliance policies. Policy templates are a great way to not only quickly configure a new compliance policy, but to also quickly modify and update policies as your requirements change. For example, you may want to quickly test a policy for offensive language and anti-harassment on communications for a small group of users before configuring a policy for all users in your organization.

> **IMPORTANT**
> By default, Global Administrators do not have access to communication compliance features. To enable permissions for communication compliance features, see Make communication compliance available in your organization.

You can choose from the following policy templates in the Microsoft 365 compliance center:

- **Offensive or threatening language**: Use this template to quickly create a policy that uses built-in classifiers to automatically detect content that may be considered abusive or offensive.
- **Sensitive information**: Use this template to quickly create a policy to scan communications containing defined sensitive information types or keywords to help make sure that important data isn't shared with people that shouldn't have access.
- **Regulatory compliance**: Use this template to quickly create a policy to scan communications for references to standard financial terms associated with regulatory standards.
- **Conflict of interest**: Use this template to quickly create a policy to monitor communications between two groups or two users to help avoid conflicts of interest.
- **Custom policy**: Use this template to configure specific communication channels, individual detection conditions, and the amount of content to monitor and review in your organization.

Investigate

In this step, you look deeper into the issues detected as matching your communication compliance policies. This step includes the following actions available in the Microsoft 365 compliance center:

- **Alerts**: When a message matches a policy condition, an alert is automatically generated. For each alert, you can see the status, the severity, the time detected, and if an Advanced eDiscovery case is assigned and its status. New alerts are displayed on the communication compliance home page and the **Alerts** page and are listed in order of severity.
- **Issue management**: For each alert, you can take investigative actions to help remediate the issue detected in the message.
- **Document review**: During the investigation of an issue, you can use several views of the message to help properly evaluate the detected issue. The views include a conversation summary, text-only, annotated, and detail views of the communication conversation.
- **Reviewing user activity history**: View the history of user message activities and remediation actions, such as past notifications and escalations, for policy matches.
- **Filters**: Use filters such as sender, recipient, date, and subject to quickly narrow down the message alerts that you want to review.

Remediate

The next step is to remediate communication compliance issues you've investigated using the following options:

- **Resolve**: After reviewing an issue, you can remediate by resolving the alert. Resolving an alert removes it from the pending alert queue, and the action is preserved as an entry in the Resolved queue for the matching policy. Alerts are automatically resolved after marking the alert as a false positive, sending a notice to a user about the alert, or opening a new case for the alert.
- **Tag a message**: As part of the resolution of an issue, you can tag the detected message as compliant, non-compliant, or as questionable as it relates to the policies and standards for your organization. Tagging can help you micro-filter policy alerts for escalations or as part of other internal review processes.
- **Notify the user**: Often, users accidentally or inadvertently violate a communication compliance policy. You can use the notify feature to provide a warning notice to the user and to resolve the issue.
- **Escalate to another reviewer**: Sometimes, the initial reviewer of an issue needs input from other reviewers to help resolve the incident. You can easily escalate message issues to reviewers in other areas of your organization as part of the resolution process.
- **Mark as a false positive**: Messages incorrectly detected as matches of compliance policies will occasionally slip through to the review process. You can mark these types of alerts as false positives and automatically resolve the issue.
- **Remove message in Teams (preview)**: Inappropriate messages may be removed from displaying in Microsoft Teams channels or personal and group chat messages. Inappropriate messages that are removed are replaced with a notification that the message has been removed for a policy violation.

- **Escalate for investigation**: In the most serious situations, you may need to share communication compliance information with other reviewers in your organization. Communication compliance is tightly integrated with other Microsoft 365 compliance features to help you with end-to-end risk resolution. Escalating a case for investigation allows you to transfer data and management of the case to Advanced eDiscovery in Microsoft 365. Advanced eDiscovery provides an end-to-end workflow to preserve, collect, review, analyze, and export content that's responsive to your organization's internal and external investigations. It allows legal teams to manage the entire legal hold notification workflow. To learn more about Advanced eDiscovery cases, see Overview of Advanced eDiscovery in Microsoft 365.

Monitor

Keeping track and managing compliance issues identified by communication compliance policies spans the entire workflow process. As alerts are generated and investigation and remediation actions are implemented, existing policies may need review and updates, and new policies may need to be created.

- **Monitor and report**: Use communication compliance dashboard widgets, export logs, and events recorded in the unified audit logs to continually evaluate and improve your compliance posture.

Ready to get started?

- For planning information, see Plan for communication compliance.
- Check out the case study for Contoso and see how they quickly configured a communication compliance policy to monitor for offensive language in Microsoft Teams, Exchange Online, and Yammer communications.
- To configure communication compliance for your Microsoft 365 organization, see Configure communication compliance for Microsoft 365.

Plan for communication compliance

11/2/2020 • 4 minutes to read • Edit Online

Before getting started with communication compliance in your organization, there are important planning activities and considerations that should be reviewed by your information technology and compliance management teams. Thoroughly understanding and planning for deployment in the following areas will help ensure that your implementation and use of communication compliance features goes smoothly and is aligned with the best practices for the solution.

Work with stakeholders in your organization

Identify the appropriate stakeholders in your organization to collaborate for taking actions on communication compliance alerts. Some recommended stakeholders to consider including in initial planning and the end-to-end communication compliance workflow are people from the following areas of your organization:

- Information technology
- Compliance
- Privacy
- Security
- Human resources
- Legal

Plan for the investigation and remediation workflow

Select dedicated stakeholders to monitor and review the alerts and cases on a regular cadence in the Microsoft 365 compliance center. Make sure understand how you will assign users and stakeholders to different communication compliance role groups in your organization.

Depending on how you wish to manage communication policies and alerts, you'll need to assign users to one or more role groups for administrators, reviewers, and investigators. You have the option to assign users to specific role groups to manage different sets of communication compliance features. Or you may decide to assign all the communication compliance users to the Communication Compliance role group. Use a single role group or multiple groups to best fit your compliance management requirements.

Plan to choose from these role group options when configuring communication compliance:

ROLE	ROLE PERMISSIONS
Communication Compliance	Use this role group to manage communication compliance for your organization in a single group. By adding all user accounts for designated administrators, analysts, investigators, and viewers, you can configure communication compliance permissions in a single group. This role group contains all the communication compliance permission roles. This configuration is the easiest way to quickly get started with communication compliance and is a good fit for organizations that do not need separate permissions defined for separate groups of users.

ROLE	ROLE PERMISSIONS
Communication Compliance Admin	Use this role group to initially configure communication compliance and later to segregate communication compliance administrators into a defined group. Users assigned to this role group can create, read, update, and delete communication compliance policies, global settings, and role group assignments. Users assigned to this role group cannot view message alerts.
Communication Compliance Analyst	Use this group to assign permissions to users that will act as communication compliance analysts. Users assigned to this role group can view policies where they are assigned as Reviewers, view message metadata (not message content), escalate to additional reviewers, or send notifications to users. Analysts cannot resolve pending alerts.
Communication Compliance Investigator	Use this group to assign permissions to users that will act as communication compliance investigators. Users assigned to this role group can view message metadata and content, escalate to additional reviewers, escalate to an Advanced eDiscovery case, send notifications to users, and resolve the alert.
Communication Compliance Viewer	Use this group to assign permissions to users that will manage communication reports. Users assigned to this role group can access all reporting widgets on the communication compliance home page and can view all communication compliance reports.

Plan for policies

Creating communication compliance policies is quick and easy with the pre-defined templates for offensive language, sensitive information, and regulatory compliance. Custom communication compliance policies allow the flexibility for detecting and investigation issues specific to your organization and requirements.

When planning for communication compliance policies, consider the following areas:

- Consider adding all users in your organization as in-scope for your communication compliance policies. Identifying specific users as in-scope for individual policies are useful in some circumstances, however most organizations should include all users in communication compliance policies optimized for harassment or discrimination detection.
- To simplify your setup, consider creating groups for people who need their communications reviewed. If you're using groups; you might need several. For example, if you want to scan communications between two distinct groups of people, or if you want to specify a group that isn't supervised.
- Configure the percentage of communications to review at 100% to ensure that policies are catching all issues of concern in communications for your organization.
- You can scan communications from third-party sources for data imported into mailboxes in your Microsoft 365 organization. To include review of communications in these platforms, you'll need to configure a connector to these services before messages meeting policy conditions are monitored by communication policy.
- Policies can support monitoring languages other than English in custom communication compliance policies. Build a custom keyword dictionary of offensive words in the language of your choice or build your own machine-learning model using trainable classifiers in Microsoft 365.
- All organizations have different communication standards and policy needs. Monitor for specific keywords using communication compliance policy conditions or monitor for specific types of information with custom

sensitive information types.

Ready to get started?

To configure communication compliance for your Microsoft 365 organization, see Configure communication compliance for Microsoft 365 or check out the case study for Contoso and how they quickly configured a communication compliance policy to monitor for offensive language in Microsoft Teams, Exchange Online, and Yammer communications.

Get started with communication compliance

2/18/2021 • 15 minutes to read • Edit Online

Use communication compliance policies to identify user communications for examination by internal or external reviewers. For more information about how communication compliance policies can help you monitor communications in your organization, see communication compliance policies in Microsoft 365. If you'd like to review how Contoso quickly configured a communication compliance policy to monitor for offensive language in Microsoft Teams, Exchange Online, and Yammer communications, check out this case study.

Subscriptions and licensing

Before you get started with communication compliance, you should confirm your Microsoft 365 subscription and any add-ons. To access and use communication compliance, your organization must have one of the following subscriptions or add-ons:

- Microsoft 365 E5 subscription (paid or trial version)
- Microsoft 365 E3 subscription + the Microsoft 365 E5 Compliance add-on
- Microsoft 365 E3 subscription + the Microsoft 365 E5 Insider Risk Management add-on
- Microsoft 365 A5 subscription (paid or trial version)
- Microsoft 365 A3 subscription + the Microsoft 365 A5 Compliance add-on
- Microsoft 365 A3 subscription + the Microsoft 365 A5 Insider Risk Management add-on
- Microsoft 365 G5 subscription (paid or trial version)
- Microsoft 365 G5 subscription + the Microsoft 365 G5 Compliance add-on
- Microsoft 365 G5 subscription + the Microsoft 365 G5 Insider Risk Management add-on
- Office 365 Enterprise E5 subscription (paid or trial version)
- Office 365 A5 subscription (paid or trial version)
- Office 365 Enterprise E3 subscription + the Office 365 Advanced Compliance add-on (no longer available for new subscriptions, see note)

Users included in communication compliance policies must be assigned one of the licenses above.

> **IMPORTANT**
> Office 365 Advanced Compliance is no longer sold as a standalone subscription. When current subscriptions expire, customers should transition to one of the subscriptions above, which contain the same or additional compliance features.

If you don't have an existing Office 365 Enterprise E5 plan and want to try communication compliance, you can add Microsoft 365 to your existing subscription or sign up for a trial of Office 365 Enterprise E5.

Step 1 (required): Enable permissions for communication compliance

> **IMPORTANT**
> By default, Global Administrators do not have access to communication compliance features. The roles assigned in this step are required before any communication compliance features will be accessible. After configuring your role groups, it may take up to 30 minutes for the role group permissions to apply to assigned users across your organization.

There are five role groups used to configure permissions to manage communication compliance features. To

make **Communication compliance** available as a menu option in Microsoft 365 compliance center and to continue with these configuration steps, you must be assigned to the *Communication Compliance* or *Communication Compliance Admin* role groups. To access and manage communication compliance features after initial configuration, users must be a member of at least one communication compliance role group.

Depending on how you wish to manage communication policies and alerts, you'll need to assign users to specific role groups. You have the option to assign users with different compliance responsibilities to specific role groups to manage different areas of communication compliance features. Or you may decide to assign all user accounts for designated administrators, analysts, investigators, and viewers to the *Communication Compliance* role group. Use a single role group or multiple role groups to best fit your compliance management requirements.

Choose from these role group options when configuring communication compliance:

ROLE	ROLE PERMISSIONS
Communication Compliance	Use this role group to manage communication compliance for your organization in a single group. By adding all user accounts for designated administrators, analysts, investigators, and viewers, you can configure communication compliance permissions in a single group. This role group contains all the communication compliance permission roles. This configuration is the easiest way to quickly get started with communication compliance and is a good fit for organizations that do not need separate permissions defined for separate groups of users.
Communication Compliance Admin	Use this role group to initially configure communication compliance and later to segregate communication compliance administrators into a defined group. Users assigned to this role group can create, read, update, and delete communication compliance policies, global settings, and role group assignments. Users assigned to this role group cannot view message alerts.
Communication Compliance Analyst	Use this group to assign permissions to users that will act as communication compliance analysts. Users assigned to this role group can view policies where they are assigned as Reviewers, view message metadata (not message content), escalate to additional reviewers, or send notifications to users. Analysts cannot resolve pending alerts.
Communication Compliance Investigator	Use this group to assign permissions to users that will act as communication compliance investigators. Users assigned to this role group can view message metadata and content, escalate to additional reviewers, escalate to an Advanced eDiscovery case, send notifications to users, and resolve the alert.
Communication Compliance Viewer	Use this group to assign permissions to users that will manage communication reports. Users assigned to this role group can access all reporting widgets on the communication compliance home page and can view all communication compliance reports.

Option 1: Assign all compliance users to the Communication Compliance role group

1. Sign into https://protection.office.com/permissions using credentials for an admin account in your Microsoft 365 organization.

2. In the Security & Compliance Center, go to **Permissions**. Select the link to view and manage roles in Office 365.
3. Select the *Communication Compliance* role group, then select **Edit role group**.
4. Select **Choose members** from the left navigation pane, then select **Edit**.
5. Select **Add** and then select the checkbox for all users you want to add to the *Communication Compliance* role group.
6. Select **Add**, then select **Done**.
7. Select **Save** to add the users to the role group. Select **Close** to complete the steps

Option 2: Assign users to specific communication compliance role groups

Use this option to assign users to specific role groups to segment communication compliance access and responsibilities among different users in your organization.

1. Sign into https://protection.office.com/permissions using credentials for an admin account in your Microsoft 365 organization.
2. In the Security & Compliance Center, go to **Permissions**. Select the link to view and manage roles in Office 365.
3. Select one of the communication compliance role groups, then select **Edit role group**.
4. Select **Choose members** from the left navigation pane, then select **Edit**.
5. Select **Add** and then select the checkbox for all users you want to add to the role group.
6. Select **Add**, then select **Done**.
7. Select **Save** to add the users to the role group.
8. Select the next communication compliance role group, then repeat steps 4-7 for each required role group.
9. Select **Close** to complete the steps.

For more information about role groups and permissions, see Permissions in the Compliance Center.

Step 2 (required): Enable the audit log

Communication compliance requires audit logs to show alerts and track remediation actions taken by reviewers. The audit logs are a summary of all activities associated with a defined organizational policy or anytime a communication compliance policy changes.

For step-by-step instructions to turn on auditing, see Turn audit log search on or off. After you turn on auditing, a message is displayed that says the audit log is being prepared and that you can run a search in a couple of hours after the preparation is complete. You only have to do this action once. For more information about the using the audit log, see Search the audit log.

Step 3 (optional): Set up groups for communication compliance

When you create a communication compliance policy, you define who has their communications reviewed and who performs reviews. In the policy, you'll use email addresses to identify individuals or groups of people. To simplify your setup, you can create groups for people who have their communication reviewed and groups for people who review those communications. If you're using groups, you may need several. For example, if you want to monitor communications between two distinct groups of people or if you want to specify a group that isn't going to be supervised.

Use the following chart to help you configure groups in your organization for communication compliance policies:

POLICY MEMBER	SUPPORTED GROUPS	UNSUPPORTED GROUPS
Supervised users Non-supervised users	Distribution groups Microsoft 365 Groups	Dynamic distribution groups Nested distribution groups Mail-enabled security groups
Reviewers	None	Distribution groups Dynamic distribution groups Nested distribution groups Mail-enabled security groups

When you assign a distribution group in the policy, the policy monitors all emails and Teams chats from each user in distribution group. When you assign a Microsoft 365 group in the policy, the policy monitors all emails and Teams chats sent to that group, not the individual emails and chats received by each group member.

If you're an organization with an Exchange on-premises deployment or an external email provider and you want to monitor Microsoft Teams chats for your users, you must create a distribution group for the users with on-premises or external mailboxes to monitor. Later in these steps, you'll assign this distribution group as the **Supervised users and groups** selection in the policy wizard.

> **IMPORTANT**
> You must file a request with Microsoft Support to enable your organization to use the graphical user interface in the Security & Compliance Center to search for Teams chat data for on-premises users. For more information, see Searching cloud-based mailboxes for on-premises users.

To manage supervised users in large enterprise organizations, you may need to monitor all users across large groups. You can use PowerShell to configure a distribution group for a global communication compliance policy for the assigned group. This enables you to monitor thousands of users with a single policy and keep the communication compliance policy updated as new employees join your organization.

1. Create a dedicated distribution group for your global communication compliance policy with the following properties: Make sure that this distribution group isn't used for other purposes or other Office 365 services.

   - **MemberDepartRestriction = Closed**. Ensures that users cannot remove themselves from the distribution group.
   - **MemberJoinRestriction = Closed**. Ensures that users cannot add themselves to the distribution group.
   - **ModerationEnabled = True**. Ensures that all messages sent to this group are subject to approval and that the group is not being used to communicate outside of the communication compliance policy configuration.

   ```
   New-DistributionGroup -Name <your group name> -Alias <your group alias> -MemberDepartRestriction
   'Closed' -MemberJoinRestriction 'Closed' -ModerationEnabled $true
   ```

2. Select an unused Exchange custom attribute to track users added to the communication compliance policy in your organization.

3. Run the following PowerShell script on a recurring schedule to add users to the communication compliance policy:

```
$Mbx = (Get-Mailbox -RecipientTypeDetails UserMailbox -ResultSize Unlimited -Filter {CustomAttribute9
-eq $Null})
$i = 0
ForEach ($M in $Mbx)
{
  Write-Host "Adding" $M.DisplayName
  Add-DistributionGroupMember -Identity <your group name> -Member $M.DistinguishedName -ErrorAction
SilentlyContinue
  Set-Mailbox -Identity $M.Alias -<your custom attribute name> SRAdded
  $i++
}
Write-Host $i "Mailboxes added to supervisory review distribution group."
```

For more information about setting up groups, see:

- Create and manage distribution groups
- Overview of Microsoft 365 Groups

Step 4 (optional): Verify your Yammer tenant is in Native Mode

In Native Mode, all Yammer users are in Azure Active Directory (Azure AD), all groups are Office 365 Groups, and all files are stored in SharePoint Online. Your Yammer tenant must be in Native Mode for communication compliance policies to scan and identify risky conversations in private messages and community conversations in Yammer.

For more information about configuring Yammer in Native Mode, see:

- Overview of Yammer Native Mode in Microsoft 365
- Configure your Yammer network for Native Mode for Microsoft 365

Step 5 (required): Create a communication compliance policy

> **IMPORTANT**
> Using PowerShell to create and manage communication compliance policies is not supported. To create and manage these policies, you must use the policy management controls in the Microsoft 365 communication compliance solution.

1. Sign into https://compliance.microsoft.com using credentials for an admin account in your Microsoft 365 organization.
2. In the Microsoft 365 compliance center, select **Communication compliance**.
3. Select the **Policies** tab.
4. Select **Create policy** to create and configure a new policy from a template or to create and configure a custom policy.

   If you choose a policy template to create a policy, you will:

   - Confirm or update the policy name. Policy names cannot be changed once the policy is created.
   - Choose the users or groups to supervise, including choosing users or groups you'd like to exclude. When using the conflict of interest template, you'll select two groups or two users to monitor for internal communications.
   - Choose the reviewers for the policy. Reviewers are individual users and all reviewers must have mailboxes hosted on Exchange Online. Reviewers added here are the reviewers that you can choose from when escalating an alert in the investigation and remediation workflow. When

reviewers are added to a policy, they automatically receive an email message that notifies them of the assignment to the policy and provides links to information about the review process.

- Choose a limited condition field, usually a sensitive info type or keyword dictionary to apply to the policy.

If you choose to use the policy wizard to create a custom policy, you will:

- Give the policy a name and description. Policy names can't be changed once the policy is created.
- Choose the users or groups to supervise, including all users in your organization, specific users and groups, or other users and groups you'd like to exclude.
- Choose the reviewers for the policy. Reviewers are individual users and all reviewers must have mailboxes hosted on Exchange Online. Reviewers added here are the reviewers that you can choose from when escalating an alert in the investigation and remediation workflow. When reviewers are added to a policy, they automatically receive an email message that notifies them of the assignment to the policy and provides links to information about the review process.
- Choose the communication channels to scan, including Exchange, Microsoft Teams, Yammer, or Skype for Business. You'll also choose to scan third-party sources if you've configured a connector in Microsoft 365.
- Choose the communication direction to monitor, including inbound, outbound, or internal communications.
- Define the communication compliance policy conditions. You can choose from message address, keyword, file types, and size match conditions.
- Choose if you'd like to include sensitive information types. This step is where you can select default and custom sensitive info types. Pick from existing custom sensitive information types or custom keyword dictionaries in the communication compliance policy wizard. You can create these items before running the wizard if needed. You can also create new sensitive information types from within the communication compliance policy wizard.
- Choose if you'd like to enable classifiers. Classifiers can detect inappropriate language and images sent or received in the body of email messages or other types of text. You can choose the following built-in classifiers: *Threat*, *Profanity*, *Targeted harassment*, *Adult images*, *Racy images*, and *Gory images*.

  Caution

  We are deprecating the **Offensive Language** built-in classifier because it has been producing a high number of false positives. Don't use it and if you are currently using it, you should move your business processes off of it. We recommend using the **Threat**, **Profanity**, and **Targeted harassment** built-in classifiers instead.
- Define the percentage of communications to review.
- Review your policy selections and create the policy.

5. Select **Create policy** when using the templates or **Submit** when using the custom policy wizard.
6. The **Your policy was created** page is displayed with guidelines on when policy will be activated and which communications will be captured.

Step 6 (optional): Create notice templates and configure user anonymization

If you want to have the option of responding to a policy alert by sending a reminder notice to the associated

user, you'll need to create at least one notice template in your organization. The notice template fields are editable before they're sent as part of the alert remediation process, and creating a customized notice template for each communication compliance policy is recommended.

You can also choose to enable anonymization for displayed usernames when investigating policy matches and taking action on messages.

1. Sign into https://compliance.microsoft.com using credentials for an admin account in your Microsoft 365 organization.
2. In the Microsoft 365 compliance center, go to **Communication compliance**.
3. To configure anonymization for usernames, select the **Privacy** tab.
4. To enable anonymization, select **Show anonymized versions of usernames**.
5. Select **Save**.
6. Navigate to the **Notice templates** tab and then select **Create notice template**.
7. On the **Modify a notice template** page, complete the following fields:
   - Template name (required)
   - Send from (required)
   - Cc and Bcc (optional)
   - Subject (required)
   - Message body (required)
8. Select **Save** to create and save the notice template.

Step 7 (optional): Test your communication compliance policy

After you create a communication compliance policy, it's a good idea to test it to make sure that the conditions you defined are being properly enforced by the policy. You may also want to test your data loss prevention (DLP) policies if your communication compliance policies include sensitive information types. Make sure you give your policies time to activate so that the communications you want to test are captured.

Follow these steps to test your communication compliance policy:

1. Open an email client, Microsoft Teams, or Yammer while signed in as a supervised user defined in the policy you want to test.
2. Send an email, Microsoft Teams chat, or Yammer message that meets the criteria you've defined in the communication compliance policy. This test can be a keyword, attachment size, domain, etc. Make sure you determine if your configured conditional settings in the policy are too restrictive or too lenient.

   > **NOTE**
   > Email messages can take up to 24 hours to fully process in a policy. Communications in Microsoft Teams, Yammer, and third-party platforms can take up to 48 hours to fully process in a policy.

3. Sign in to Microsoft 365 as a reviewer designated in the communication compliance policy. Navigate to **Communication compliance** > **Alerts** to view the alerts for your policies.
4. Remediate the alert using the remediation controls and verify that the alert is properly resolved.

Next steps

After you've completed these steps to create your first communication compliance policy, you'll start to receive alerts from activity indicators after 24-48 hours. Configure additional policies as needed using the guidance in Step 5 of this article.

To learn more about investigating communication compliance alerts, see Investigate and remediate communication compliance alerts.

Investigate and remediate communication compliance alerts

2/18/2021 • 11 minutes to read • Edit Online

After you've configured your communication compliance policies, you'll start to receive alerts in the Microsoft 365 compliance center for message issues that match your policy conditions. Follow the workflow instructions here to investigate and remediate alert issues.

Investigate alerts

The first step to investigate issues detected by your policies is to review communication compliance alerts in the Microsoft 365 compliance center. There are several areas in the communication compliance solution area to help you to quickly investigate alerts, depending on how you prefer to view alert grouping:

- **Communication compliance policy page**: When you sign in to https://compliance.microsoft.com using credentials for an admin account in your Microsoft 365 organization, select **Communication compliance** to display the communication compliance **Policy** page. This page displays communication compliance policies configured for your Microsoft 365 organization and links to recommended policy templates. Each policy listed includes the count of alerts that need review, the number of escalated and resolved items, status of the policy, and the date and time of the last policy scan. Selecting a policy displays all the pending alerts for matches to the policy, select a specific alert to launch the policy details page and to start remediation actions.
- **Alerts**: Navigate to **Communication compliance** > **Alerts** to display the last 30 days of alerts grouped by policy matches. This view allows you to quickly see which communication compliance policies are generating the most alerts ordered by severity. To start remediation actions, select the policy associated with the alert to launch the **Policy details** page. From the **Policy details** page, you can review a summary of the activities on the **Overview** page, review and act on alert messages on the **Pending** page, or review the history of closed alerts on the **Resolved** page.
- **Reports**: Navigate to **Communication compliance** > **Reports** to display communication compliance report widgets. Each widget provides an overview of communication compliance activities and statuses, including access to deeper insights about policy matches and remediation actions.

Using filters

The next step is to sort the messages so that it's easier for you to investigate alerts. From the **Policy details** page, communication compliance supports multi-level filtering for several message fields to help you quickly investigate and review messages with policy matches. Filtering is available for pending and resolved items for each configured policy. You can configure filter queries for a policy or configure and save custom and default filter queries for use in each specific policy. After configuring fields for a filter, you'll see the filter fields displayed on the top of the alert message queue that you can configure for specific filter values.

For a complete list of filters and field details, see Filters in the feature reference article.

To configure a filter

1. Sign into https://compliance.microsoft.com using credentials for an admin account in your Microsoft 365 organization.
2. In the Microsoft 365 compliance center, go to **Communication compliance**.
3. Select the **Policies** tab and then select a policy for investigation, double-click to open the **Policy** page.
4. On the **Policy** page, select either the **Pending** or **Resolved** tab to display the items for filtering.

5. Select the **Filters** control to open the **Filters** details page.
6. Select one or more checkboxes to enable filters for these alerts. You can choose from numerous filters, including *Date*, *Sender*, *Subject/Title*, *Classifiers*, and more.
7. If you'd like to save the filter selected as the default filter, select **Save as default**. If you want to use this filter as a saved filter, select **Done**.
8. If you'd like to save the selected filters as a filter query, select **Save the query** control after you've configured at least one filter value. Enter a name for the filter query and select **Save**. This filter is available to use for only this policy and is listed in the **Saved filter queries** section of the **Filters** details page.

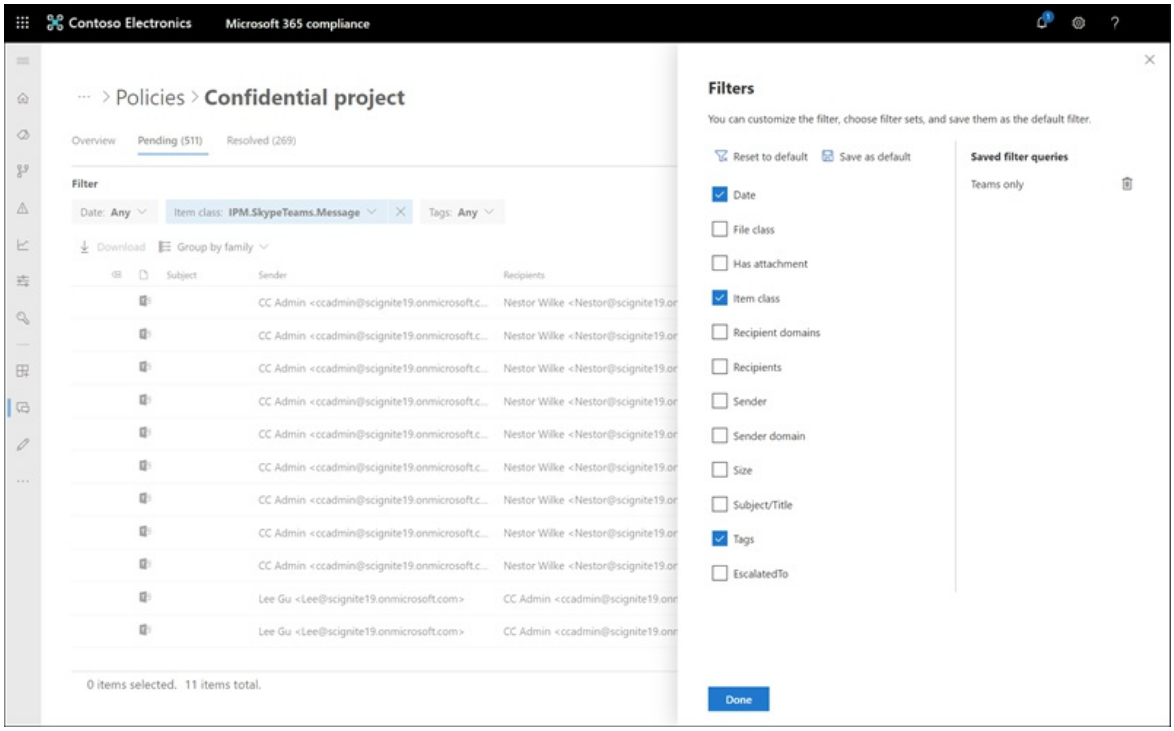

Using near and exact duplicate analysis

Communication compliance policies automatically scan and pre-group near and exact message duplicates without any additional configuration steps. This view allows you to quickly act on similar messages one-by-one or as a group, reducing the message investigation burden for reviewers. As duplicates are detected, the **Near Duplicates** and/or the **Exact Duplicates** controls are displayed in the remediation action toolbar. This view isn't available if near or exact duplicates aren't found.

To remediate duplicates

1. Sign into https://compliance.microsoft.com using credentials for an admin account in your Microsoft 365 organization.
2. In the Microsoft 365 compliance center, go to **Communication compliance**.
3. Select the **Policies** tab and then select a policy for investigation, double-click to open the **Policy** page.
4. On the **Policy** page, select either the **Pending** or **Resolved** tab to display duplicate messages.
5. Select the **Near Duplicates** or **Exact Duplicates** controls to open the duplicates details page.
6. Select one or more messages to remediation action controls for these messages.
7. Select **Resolve**, **Notify**, **Escalate**, or **Download** to apply the action to the selected duplicate messages as the default filter.
8. Select **Close** after completing the remediation actions on the messages.

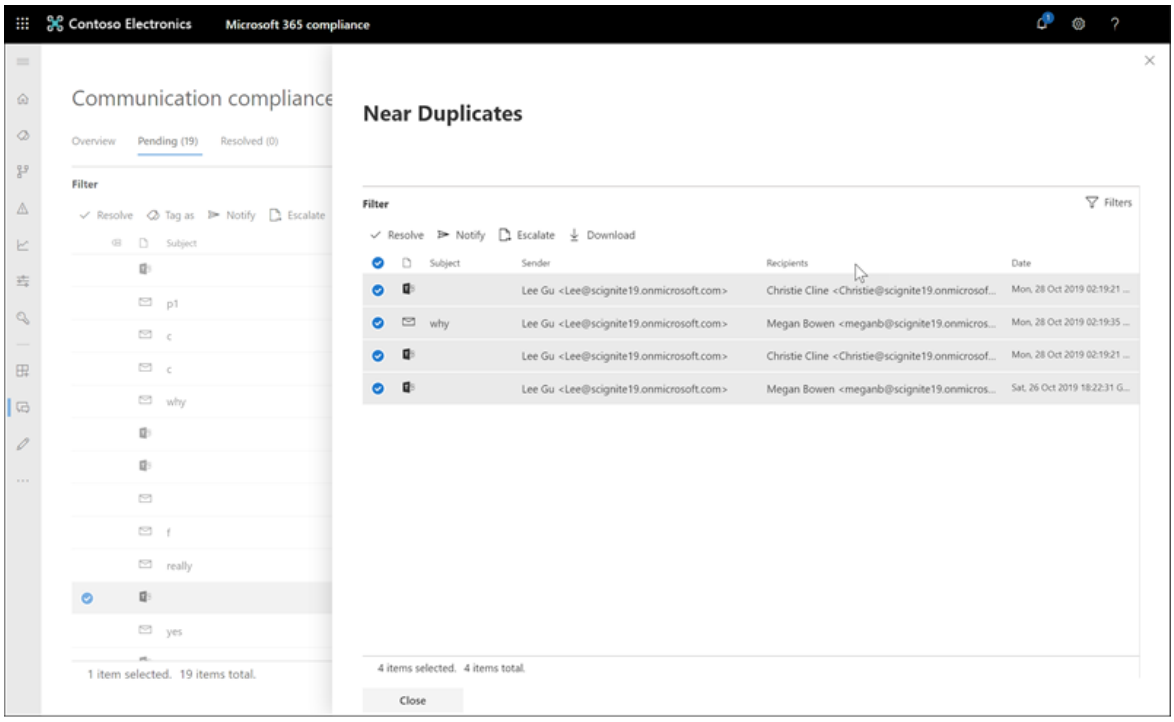

Remediate alerts

No matter where you start to review alerts or the filtering you configure, the next step is to take action to remediate the alert. Start your alert remediation using the following workflow on the **Policy** or **Alerts** pages.

Step 1: Examine the message basics

Sometimes it's obvious from the source or subject that a message can be immediately remediated. It may be that the message is spurious or incorrectly matched to a policy and it should be resolved as a false positive. Select the **False Positive** control to immediately resolve the alert and remove from the pending alert queue. From the source or sender information, you may already know how the message should be routed or handled in these circumstances. Consider using the **Tag as** or **Escalate** controls to assign a tag to applicable messages or to send messages to a designated reviewer.

Step 2: Examine the message details

After reviewing the message basics, it's time to open a message to examine the details and to determine further remediation actions. Select a message to view the complete message header and body information. Several different views are available to help you decide the proper course of action:

- **Source view**: This view is the standard message view commonly seen in most web-based messaging platforms. The header information is formatted in the normal style and the message body supports imbedded graphic files and word-wrapped text.
- **Text view**: Text view displays a line-numbered text-only view of the message and includes keyword highlighting in messages and attachments for terms matched in the associated communication compliance policy. Keyword highlighting can help you quickly scan long messages and attachments for the area of interest. In some cases, highlighted text may be only in attachments for messages matching policy conditions. Embedded files aren't displayed and the line numbering this view is helpful for referencing pertinent details among multiple reviewers.

- **Annotate view**: This view allows reviewers to add annotations directly on the message that are saved to the view of the message.
- **User history**: User history view displays all other alerts generated by any communication compliance policy for the user sending the message.
- **Message detail view**: Advanced view of message metadata and configuration information.
- **Pattern detected notification**: Many harassing and bullying actions over time and involve reoccurring instances of the same behavior by a user. The *Pattern detected* notification is displayed in the alert details and raises attention to the alert. Detection of patterns is on a per-policy basis and evaluates behavior over the last 30 days when at least two messages are sent to the same recipient by a sender. Investigators and reviewers can use this notification to identify repeated behavior to evaluate the alert as appropriate.
- **Show Translate view**: This view automatically converts alert message text to the language configured in the *Displayed language* setting in the Microsoft 365 subscription for each reviewer. The Translate view helps broaden investigative support for organizations with multilingual users and eliminates the need for additional translation services outside of the communication compliance review process. Using Microsoft Translate services, the Translate view can be turned on and off as needed and supports a wide range of languages. For a complete list of supported languages, see Microsoft Translator Languages. Languages listed in the *Translator Language List* are supported in the Translate view.

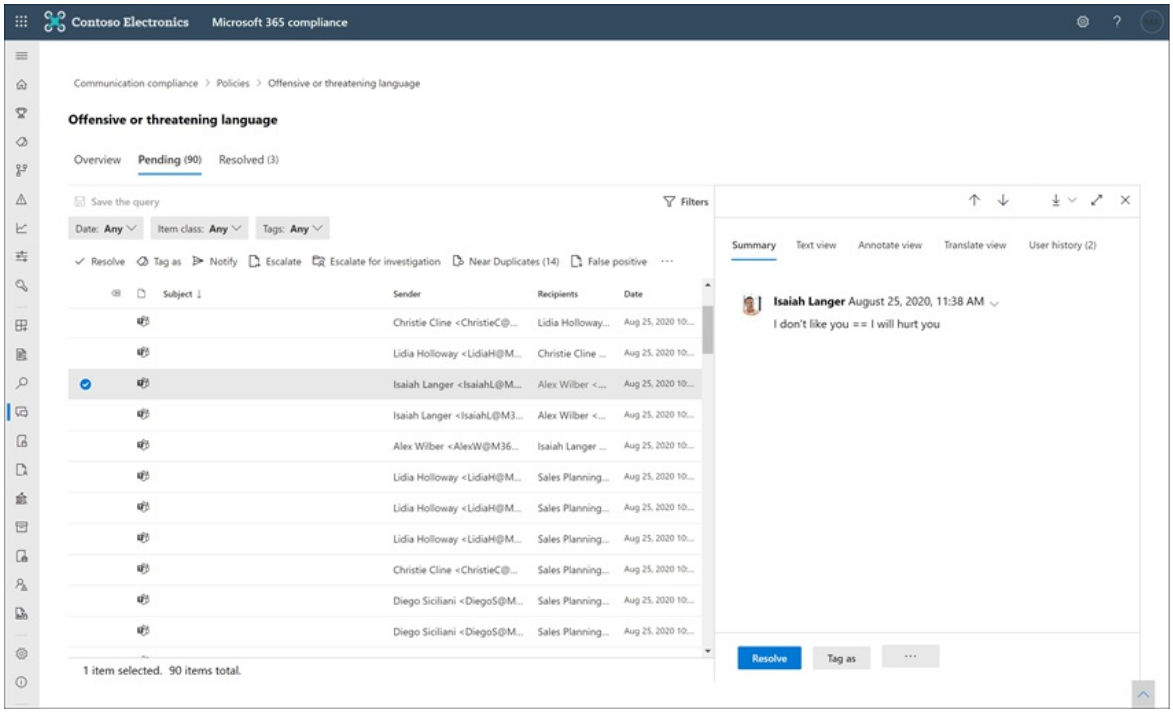

Step 3: Decide on a remediation action

Now that you've reviewed the details of the message for the alert, you can choose several remediation actions:

- **Resolve**: Selecting the **Resolve** control immediately removes the message from the **Pending alerts** queue and no further action can be taken on the message. By selecting **Resolve**, you've essentially closed the alert without further classification and it can't be reopened for further actions. All resolved messages are displayed in the **Resolved** tab.
- **False Positive**: You can always resolve a message as a false positive at any point during the message review workflow. False positive signifies that the alert was non-actionable or that the alert was incorrectly generated by the alerting process. The message cannot be reopened and all false positive messages are displayed in the **Resolved** tab.
- **Power Automate (preview)**: Use a Power Automate flow to automate process tasks for an alert message. By default, communication compliance includes the *Notify manager when a user has a*

communication compliance alert flow template that reviewers can use to automate the notification process for users with message alerts. For more information about creating and managing Power Automate flows in communication compliance, see the Communication compliance feature reference article.

- **Tag as**: Tag the message as *compliant*, *non-compliant*, or as *questionable* as it relates to the policies and standards for your organization. Adding tags and tagging comments helps you micro-filter policy alerts for escalations or as part of other internal review processes. After tagging is complete, you can also choose to resolve the message to move it out of the pending review queue.
- **Notify**: You can use the **Notify** control to assign a custom notice template to the alert and to send a warning notice to the user. Choose the appropriate notice template configured in the **Communication compliance settings** area and select **Send** to email a reminder to the user that sent the message and to resolve the issue.
- **Escalate**: Using the **Escalate** control, you can choose who else in your organization should review the message. Choose from a list of reviewers configured in the communication compliance policy to send an email notification requesting additional review of the message alert. The selected reviewer can use a link in the email notification to go directly to items escalated to them for review.
- **Escalate for investigation**: Using the **Escalate for investigation** control, you can create a new Advanced eDiscovery case for single or multiple messages. You'll provide a name and notes for the new case, and user who sent the message matching the policy is automatically assigned as the case custodian. You don't need any additional permissions to manage the case. Creating a case does not resolve or create a new tag for the message. You can select a total of 100 messages when creating an Advanced eDiscovery case during the remediation process. Messages in all communication channels monitored by communication compliance are supported. For example, you could select 50 Microsoft Teams chats, 25 Exchange Online email messages, and 25 Yammer messages when you open a new Advanced eDiscovery case for a user.
- **Improve classification (preview)**: Alerts created from classifier type matches may need feedback to help minimize false positives in your organization. Use the **Improve classification** control to provide feedback on if the communication compliance classification is valid or to suggest other trainable classifiers for this type of match. You can confirm that the classifiers are either a *Match* or *Not a match*, or suggest other trainable classifiers to associate with this type of alert activity in the future.

  1. Select a message from the alert list.
  2. Choose the ellipsis and select **Improve classification**.
  3. In the **Detailed classifier feedback** pane, if the item is a true positive, choose **Match**. If the item was incorrectly included in the category as a false positive, choose **Not a match**.
  4. If there is another classifier that would be more appropriate for the item, choose it from the **Suggest other trainable classifiers** list. This feedback triggers the other classifier to evaluate the item.

  > **TIP**
  > You can provide feedback on multiple items simultaneously by choosing them all and then choosing **Provide detailed feedback** in the command bar.

  5. Choose **Send feedback** to send your evaluation of the **Match** and **Not a match** classifications and suggest other trainable classifiers. When you've provided 30 instances of feedback to a classifier, it automatically retrains. Retraining may take 1-4 hours to complete. Classifiers can only be retrained twice per day.

> **IMPORTANT**
>
> This information goes to the classifier in your tenant, **it does not go back to Microsoft**.

To learn more about retraining classifier for communication compliance, see the How to retrain a classifier in communications compliance article.

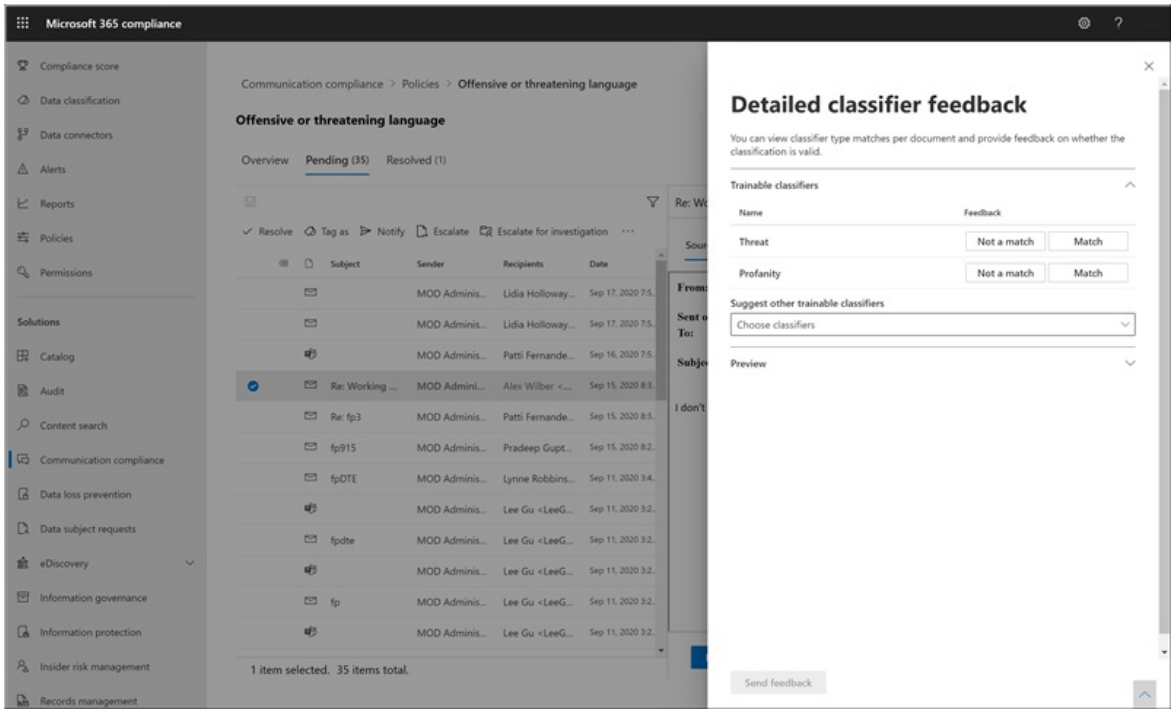

- **Remove message in Teams**: Using the **Remove message in Teams** control, you can block inappropriate messages and content identified in alerts from Microsoft Teams channels and 1:1 and group chats. Removed messages and content are replaced with a policy tip that explains that it is blocked and the policy that applies to its removal from view. Recipients are provided a link in the policy tip to learn more about the applicable policy and the review process. The sender receives a policy tip for the blocked message and content but can review the details of the blocked message and content for context regarding the removal.

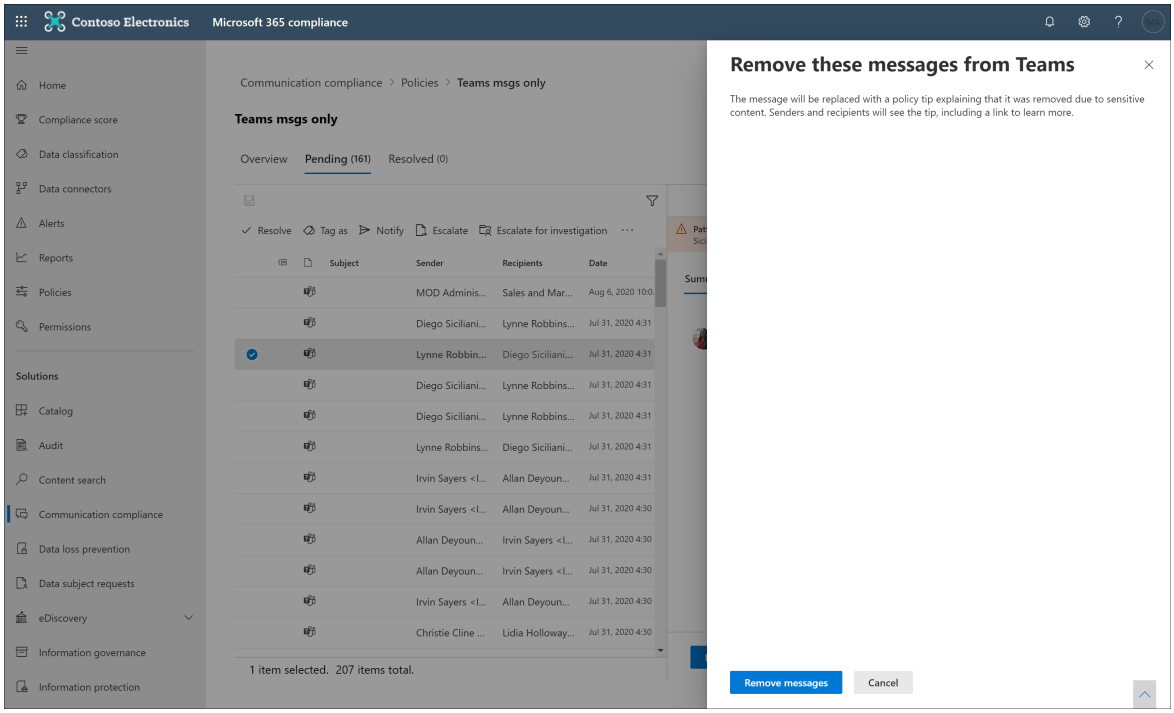

Step 4: Determine if message details should be archived outside of communication compliance

Message details can be exported or downloaded if you need to archive the messages in a separate storage solution. Selecting the **Download** control automatically adds selected messages to a .ZIP file that can be saved to storage outside of Microsoft 365.

Communication compliance feature reference

2/18/2021 • 38 minutes to read • Edit Online

Policies

> **IMPORTANT**
>
> Using PowerShell to create and manage communication compliance policies is not supported. To create and manage these policies, you must use the policy management controls in the Microsoft 365 communication compliance solution.

You create communication compliance policies for Microsoft 365 organizations in the Microsoft 365 compliance center. Communication compliance policies define which communications and users are subject to review in your organization, define which custom conditions the communications must meet, and specify who should do reviews. Users assigned the *Communication Compliance Admin* role can set up policies, and anyone who has this role assigned can access the **Communication compliance** page and global settings in the Microsoft 365 compliance center. If needed, you can export the history of modifications to a policy to a .csv file that also includes the status of alerts pending review, escalated items, and resolved items. Policies can't be renamed and can be deleted when no longer needed.

> **NOTE**
>
> Supervision policies created in the Security & Compliance Center for Office 365 subscriptions cannot migrate to Microsoft 365. If you're migrating from an Office 365 subscription to a Microsoft 365 subscription, you'll need to create new communication compliance polices to replace existing Supervision policies.

Policy templates

Policy templates are pre-defined policy settings that you can use to quickly create policies to address common compliance scenarios. Each of these templates has differences in conditions and scope, and all templates use the same types of scanning signals. You can choose from the following policy templates:

AREA	POLICY TEMPLATE	DETAILS
Offensive language and anti-harassment	Monitor communications for offensive language	- Locations: Exchange Online, Microsoft Teams, Yammer, Skype for Business - Direction: Inbound, Outbound, Internal - Review Percentage: 100% - Conditions: Offensive language classifier

AREA	POLICY TEMPLATE	DETAILS
Sensitive information	Monitor communications for sensitive information	- Locations: Exchange Online, Microsoft Teams, Yammer, Skype for Business - Direction: Inbound, Outbound, Internal - Review Percentage: 10% - Conditions: Sensitive information, out-of-the-box content patterns, and types, custom dictionary option, attachments larger than 1 MB
Regulatory compliance	Monitor communications for info related to financial regulatory compliance	- Locations: Exchange Online, Microsoft Teams, Yammer, Skype for Business - Direction: Inbound, Outbound - Review Percentage: 10% - Conditions: custom dictionary option, attachments larger than 1 MB
Conflict of interest	Monitor communications between two groups or two users to help avoid conflicts of interest	- Locations: Exchange Online, Microsoft Teams, Yammer, Skype for Business - Direction: Internal - Review Percentage: 100% - Conditions: None

Communications are scanned every 24 hours from the time policies are created. For example, if you create an offensive language policy at 11:00 AM, the policy will gather communication compliance signals every 24 hours at 11:00 AM daily. Editing a policy doesn't change this time. To view the last scan date and time for a policy, navigate to the *Last policy scan* column on the **Policy** page. After creating a new policy, it may take up to 24 hours to view the first policy scan date and time. The date and time of the last scan will be converted to the time zone of your local system.

Permissions

> **IMPORTANT**
> By default, Global Administrators do not have access to communication compliance features. The roles assigned in this step are required before any communication compliance features will be accessible.

There are five role groups used to configure permissions to manage communication compliance features. To make **Communication compliance** available as a menu option in Microsoft 365 compliance center and to continue with these configuration steps, you must be assigned to the *Communication Compliance* or *Communication Compliance Admin* role groups. To access and manage communication compliance features after initial configuration, users must be a member of at least one communication compliance role group.

Depending on how you wish to manage communication policies and alerts, you'll need to assign users to specific role groups. You have the option to assign users with different compliance responsibilities to specific role groups to manage different areas of communication compliance features. Or you may decide to assign all user accounts for designated administrators, analysts, investigators, and viewers to the *Communication Compliance* role group. Use a single role group or multiple role groups to best fit your compliance management requirements.

Choose from these role group options when configuring communication compliance:

ROLE GROUP	ROLE GROUP PERMISSIONS
Communication Compliance	Use this role group to manage communication compliance for your organization in a single group. By adding all user accounts for designated administrators, analysts, investigators, and viewers, you can configure communication compliance permissions in a single group. This role group contains all the communication compliance permission roles. This configuration is the easiest way to quickly get started with communication compliance and is a good fit for organizations that do not need separate permissions defined for separate groups of users.
Communication Compliance Admin	Use this role group to initially configure communication compliance and later to segregate communication compliance administrators into a defined group. Users assigned to this role group can create, read, update, and delete communication compliance policies, global settings, and role group assignments. Users assigned to this role group cannot view message alerts.
Communication Compliance Analyst	Use this group to assign permissions to users that will act as communication compliance analysts. Users assigned to this role group can view policies where they are assigned as Reviewers, view message metadata (not message content), escalate to other reviewers, or send notifications to users. Analysts cannot resolve pending alerts.
Communication Compliance Investigator	Use this group to assign permissions to users that will act as communication compliance investigators. Users assigned to this role group can view message metadata and content, escalate to other reviewers, escalate to an Advanced eDiscovery case, send notifications to users, and resolve the alert.
Communication Compliance Viewer	Use this group to assign permissions to users that will manage communication reports. Users assigned to this role group can access all reporting widgets on the communication compliance home page and can view all communication compliance reports.

For organizations using the original permissions and role groups

The new role group structure replaces initial role group structure for communication compliance. For organizations already using communication compliance, you needed to be assigned the Supervisory Review Administrator role to get started with communication compliance in the Microsoft 365 compliance center. Additionally, you had to create a new role group for reviewers with the Supervisory Review Administrator, Case Management, Compliance Administrator, and Review roles to investigate and remediate messages with policy matches. Essentially, all admins and reviewers were in a single role group and everyone had the same access and management permissions. With the latest updates to communication compliance, you should plan to migrate from the previous role group structure to the new role group structure. Support for the previous role group structure will be phased out.

To help your migration planning, consider the following example. You currently have three types of users in your organization, IT admins, triage, and reviewers. These three types of users are in the previous role group structure and are all members of a single role group with the following roles assigned:

- Supervisory Review Administrator
- Case Management

- Compliance Administrator
- Review

To update the roles for these users for the new role group structure, and to separate the access and management permissions for the users, you may consider three new groups and the associated new role group assignments:

- **IT Admins**: Assigned to the new *Communication Compliance Admin* role group.
- **Triage**: Assigned to the *Communication Compliance Analyst* role group.
- **Reviewers**: Assigned to the new *Communication Compliance Investigator* role group.

Supervised users

Before you start using communication compliance, you must determine who needs their communications reviewed. In the policy, user email addresses identify individuals or groups of people to supervise. Some examples of these groups are Microsoft 365 Groups, Exchange-based distribution lists, Yammer communities, and Microsoft Teams channels. You also can exclude specific users or groups from scanning with a specific exclusion group or a list of groups.

> **IMPORTANT**
> Users covered by communication compliance policies must have either a Microsoft 365 E5 Compliance license, an Office 365 Enterprise E3 license with the Advanced Compliance add-on, or be included in an Office 365 Enterprise E5 subscription. If you don't have an existing Enterprise E5 plan and want to try communication compliance, you can sign up for a trial of Office 365 Enterprise E5.

Reviewers

When you create a communication compliance policy, you must determine who reviews the messages of the supervised users. In the policy, user email addresses identify individuals or groups of people to review supervised communications. All reviewers must have mailboxes hosted on Exchange Online and must be assigned to either the *Communication Compliance Analysis* or *Communication Compliance Investigation* roles. Reviewers (either analysts or investigators) must also have the *Communication Compliance Case Management* role assigned. When reviewers are added to a policy, they automatically receive an email message that notifies them of the assignment to the policy and provides links to information about the review process.

Groups for supervised users and reviewers

To simplify your setup, create groups for people who need their communications reviewed and groups for people who review those communications. If you're using groups, you might need several. For example, if you want to scan communications between two distinct groups of people, or if you want to specify a group that isn't supervised.

When you assign a Distribution group in the policy, the policy monitors all emails from each user in Distribution group. When you assign a Microsoft 365 group in the policy, the policy monitors all emails sent to that group, not the individual emails received by each group member.

Adding groups and distribution lists to communication compliance policies are part of the overall conditions and rules set, so the maximum number of groups and distribution lists that a policy supports varies depending on the number of conditions also added to the policy. Each policy should support approximately 20 groups or distribution lists, depending on the number of additional conditions present in the policy.

Supported communication types

With communication compliance policies, you can choose to scan messages in one or more of the following communication platforms as a group or as standalone sources. Communications captured across these platforms are retained for seven years for each policy by default, even if users leave your organization and their mailboxes are deleted.

- **Microsoft Teams**: Chat communications in both public and private Microsoft Teams channels and individual chats can be scanned. When users are assigned to a communication compliance policy with Microsoft Teams coverage selected, chat communications for the users are automatically monitored across all Microsoft Teams where the users are a member. Microsoft Teams coverage is automatically included for pre-defined policy templates and is selected by default in the custom policy template. Teams chats matching communication compliance policy conditions may take up to 48 hours to process. Use the following group management configurations to supervise individual user chats and channel communications in Teams:
  - **For Teams chat communications:** Assign individual users or assign a distribution group to the communication compliance policy. This setting is for one-to-one or one-to-many user/chat relationships.
  - **For Teams Channel communications:** Assign every Microsoft Teams channel or Microsoft 365 group you want to scan that contains a specific user to the communication compliance policy. If you add the same user to other Microsoft Teams channels or Microsoft 365 groups, be sure to add these new channels and groups to the communication compliance policy.
  - **For Teams chat communications with hybrid email environments**: Communication compliance can monitor chat messages for users for organizations with an Exchange on-premises deployment or an external email provider that have enabled Microsoft Teams. You must create a distribution group for the users with on-premises or external mailboxes to monitor. When creating a communication compliance policy, you'll assign this distribution group as the **Supervised users and groups** selection in the policy wizard.

> **IMPORTANT**
> You must file a request with Microsoft Support to enable your organization to use the graphical user interface in the Security & Compliance Center to search for Teams chat data for on-premises users. For more information, see Searching cloud-based mailboxes for on-premises users.

You must file a request with Microsoft Support to enable your organization to use the graphical user interface in the Security & Compliance Center to search for Teams chat data in the cloud-based mailboxes for on-premises users.

- **Exchange email**: Mailboxes hosted on Exchange Online as part of your Microsoft 365 or Office 365 subscription are all eligible for message scanning. Exchange email messages and attachments matching communication compliance policy conditions may take up to 24 hours to process. Supported attachment types for communication compliance are the same as the file types supported for Exchange mail flow rule content inspections.
- **Yammer**: Private messages and public conversations and associated attachments in Yammer communities can be scanned. When a user is added to communication compliance policy that includes Yammer as a defined channel, communications across all Yammer communities that the user is a member of are included in the scanning process. Yammer chats and attachments matching communication compliance policy conditions may take up to 24 hours to process. Yammer must be in Native Mode for communication compliance policies to monitor Yammer communications and attachments. In Native Mode, all Yammer users are in Azure Active Directory (AAD), all groups are Office 365 Groups, and all files are stored in SharePoint Online.
- **Skype for Business Online**: Chat communications and associated attachments in Skype for Business

Online can be supervised. Skype for Business Online chats matching communication compliance policy conditions may take up to 24 hours to process. Supervised chat conversations are sourced from previous conversations saved in Skype for Business Online. Use the following group management configuration to supervise user chat communications in Skype for Business Online:

  - **For Skype for Business Online chat communications**: Assign individual users or assign a distribution group to the communication compliance policy. This setting is for one-to-one or one-to-many user/chat relationships.

- **Third-party sources**: You can scan communications for data imported into mailboxes in your Microsoft 365 organization from third-party sources like Instant Bloomberg, Slack, Zoom, SMS, and many others. For a full list of connectors supported in communication compliance, see Archive third-party data.

  You must configure a third-party connector for your Microsoft 365 organization before you can assign the connector to a communication compliance policy. The **Third-Party Sources** section of the communication compliance policy wizard only displays currently configured third-party connectors.

Transitioning from Supervision in Office 365

Organizations using supervision policies in Office 365 and planning to transition to communication compliance policies in Microsoft 365 need to understand these important points:

- Both solutions may be used side by side in your organization, but policies used in each solution must have unique policy names. Groups and custom keyword dictionaries can be shared between solutions during a transition period.
- Messages saved in supervision in Office 365 policy matches cannot be moved or shared into communication compliance in Microsoft 365.
- The supervision solution in Office 365 will be fully replaced by the communication compliance solution in Microsoft 365. We recommend creating new policies in communication compliance that have the same settings as existing supervision policies to use the new investigation and remediation improvements. When transitioning to communication compliance in Microsoft 365, you should plan to export reporting data from supervision in Office 365 if you have internal compliance retention policy requirements.

For retirement information for supervision in Office 365, see the Microsoft 365 Roadmap for details.

Policy settings

Users

You have the option to select **All users** or to define specific users in a communication compliance policy. Selecting **All users** applies the policy to all users and all groups that any user is included in as a member. Defining specific users applies the policy to the defined users and any groups the defined users are included in as a member.

Direction

By default, the **Direction is** condition is displayed and can't be removed. Communication direction settings in a policy are chosen individually or together:

- **Inbound**: You can choose **Inbound** to review communications sent **to** the people you chose to supervise.
- **Outbound**: You can choose **Outbound** if you want to review communications sent **from** the people you chose to supervise.
- **Internal**: You can choose **Internal** to review communications sent **between** the people you identified in the policy.

Sensitive information types

You have the option of including sensitive information types as part of your communication compliance policy.

Sensitive information types are either pre-defined or custom data types that can help identify and protect credit card numbers, bank account numbers, passport numbers, and more. As part of data loss prevention (DLP), the sensitive information configuration can use patterns, character proximity, confidence levels, and even custom data types to help identify and flag content that may be sensitive. The default sensitive information types are:

- Financial
- Medical and health
- Privacy
- Custom information type

To learn more about sensitive information details and the patterns included in the default types, see Sensitive information type entity definitions.

Custom keyword dictionaries

Configure custom keyword dictionaries (or lexicons) to provide simple management of keywords specific to your organization or industry. Keyword dictionaries support up to 100 KB of terms (post-compression) in the dictionary and support any language. The tenant limit is also 100 KB after compression. If needed, you can apply multiple custom keyword dictionaries to a single policy or have a single keyword dictionary per policy. These dictionaries are assigned in a communication compliance policy and can be sourced from a file (such as a .csv or .txt list), or from a list you can Import in the Compliance center. Use custom dictionaries when you need to support terms or languages specific to your organization and policies.

Classifiers

Built-in trainable and global classifiers scan sent or received messages across all communication channels in your organization for different types of compliance issues. Classifiers use a combination of artificial intelligence and keywords to identify language in messages likely to violate anti-harassment policies. Built-in classifiers currently support only English keywords in messages.

Communication compliance built-in trainable and global classifiers scan communications for terms, images, and sentiment for the following types of language and content:

- **Threat**: Scans for threats to commit violence or physical harm to a person or property.
- **Targeted harassment**: Scans for offensive conduct targeting people regarding race, color, religion, national origin.
- **Profanity**: Scans for profane expressions that embarrass most people.
- **Adult images**: Scans for images that are sexually explicit in nature.
- **Racy images**: Scans for images that are sexually suggestive in nature, but contain less explicit content than images deemed Adult.
- **Gory images**: Scans for images that depict violence and gore.

The *Adult*, *Racy*, and *Gory* image classifiers scan files in .JPEG, .PNG, .GIF, and .BMP formats. The size for image files must be less than 4 megabytes (MB) and the dimensions of the images must be greater than 50x50 pixels and greater than 50 kilobytes (KB) for the image to qualify for evaluation. Image identification is supported for Exchange Online email messages and Microsoft Teams channels and chats.

The built-in trainable and global classifiers don't provide an exhaustive list of terms or images across these areas. Further, language and cultural standards continually change, and in light of these realities, Microsoft reserves the right to update classifiers at its discretion. While classifiers may assist your organization in monitoring these areas, classifiers aren't intended to provide your organization's sole means of monitoring or addressing such language or imagery. Your organization, not Microsoft, remains responsible for all decisions related to monitoring, scanning, and blocking language and images in these areas, including compliance with local privacy and other applicable laws. Microsoft encourages consulting with legal counsel before deployment and use.

NOTE

Policies using classifiers will inspect and evaluate messages with a word count of six or greater. Messages containing less than six words aren't evaluated in policies using classifiers. To identify and take action on shorter messages containing inappropriate content, we recommend including a custom keyword dictionary to communication compliance policies monitoring for this type of content.

For information about trainable classifiers in Microsoft 365, see Getting started with trainable classifiers.

Conditional settings

The conditions you choose for the policy apply to communications from both email and third-party sources in your organization (like from Instant Bloomberg).

The following table explains more about each condition.

CONDITION	HOW TO USE THIS CONDITION
Content matches any of these classifiers	Apply to the policy when any classifiers are included or excluded in a message. Some classifiers are pre-defined in your tenant, and custom classifiers must be configured separately before they're available for this condition. Only one classifier can be defined as a condition in a policy. For more information about configuring classifiers, see Learn about trainable classifiers (preview).
Content contains any of these sensitive info types	Apply to the policy when any sensitive information types are included or excluded in a message. Some classifiers are pre-defined in your tenant, and custom classifiers can be configured separately or as part of the condition assignment process. Each sensitive information type you choose is applied separately and only one of these sensitive information types must apply for the policy to apply to the message. For more information about custom sensitive information types, see Learn about sensitive information types.
Message is received from any of these domains **Message is not received from any of these domains**	Apply the policy to include or exclude specific domains or email addresses in received messages. Enter each domain or email address and separate multiple domains or email addresses with a comma. Each domain or email address entered is applied separately, only one domain or email address must apply for the policy to apply to the message. If you want to scan all email from a specific domain, but want to exclude messages that don't need review (newsletters, announcements, and so on), you must configure a **Message is not received from any of these domains** condition that excludes the email address (example "newsletter@contoso.com").

CONDITION	HOW TO USE THIS CONDITION
Message is sent to any of these domains **Message is not sent to any of these domains**	Apply the policy to include or exclude specific domains or email addresses in sent messages. Enter each domain or email address and separate multiple domains or email addresses with a comma. Each domain or email address is applied separately, only one domain or email address must apply for the policy to apply to the message. If you want to scan all email sent to a specific domain, but want to exclude sent messages that don't need review, you must configure two conditions: - A **Message is sent to any of these domains** condition that defines the domain ("contoso.com"), AND - A **Message is not sent to any of these domains** condition that excludes the email address ("subscriptions@contoso.com").
Message is classified with any of these labels **Message is not classified with any of these labels**	To apply the policy when certain retention labels are included or excluded in a message. Retention labels must be configured separately and configured labels are chosen as part of this condition. Each label you choose is applied separately (only one of these labels must apply for the policy to apply to the message). For more information about retention labels, see Learn about retention policies and retention labels.
Message contains any of these words **Message contains none of these words**	To apply the policy when certain words or phrases are included or excluded in a message, enter each word separated with a comma. For phrases of two words or more, use quotation marks around the phrase. Each word or phrase you enter is applied separately (only one word must apply for the policy to apply to the message). For more information about entering words or phrases, see the next section Matching words and phrases to emails or attachments.
Attachment contains any of these words **Attachment contains none of these words**	To apply the policy when certain words or phrases are included or excluded in a message attachment (such as a Word document), enter each word separated with a comma. For phrases of two words or more, use quotation marks around the phrase. Each word or phrase you enter is applied separately (only one word must apply for the policy to apply to the attachment). For more information about entering words or phrases, see the next section Matching words and phrases to emails or attachments.
Attachment is any of these file types **Attachment is none of these file types**	To supervise communications that include or exclude specific types of attachments, enter the file extensions (such as .exe or .pdf). If you want to include or exclude multiple file extensions, enter these on separate lines. Only one attachment extension must match for the policy to apply.
Message size is larger than **Message size is not larger than**	To review messages based on a certain size, use these conditions to specify the maximum or minimum size a message can be before it's subject to review. For example, if you specify **Message size is larger than** > **1.0 MB**, all messages that are 1.01 MB and larger are subject to review. You can choose bytes, kilobytes, megabytes, or gigabytes for this condition.

CONDITION	HOW TO USE THIS CONDITION
Attachment is larger than **Attachment is not larger than**	To review messages based on the size of their attachments, specify the maximum or minimum size an attachment can be before the message and its attachments are subject to review. For example, if you specify **Attachment is larger than** > **2.0 MB**, all messages with attachments 2.01 MB and over are subject to review. You can choose bytes, kilobytes, megabytes, or gigabytes for this condition.

Matching words and phrases to emails or attachments

Each word you enter and separate with a comma is applied separately (only one word must apply for the policy condition to apply to the email or attachment). For example, let's use the condition, **Message contains any of these words**, with the keywords "banker", "confidential", and "insider trading" separated by a comma (banker, confidential,"insider trading"). The policy applies to any messages that includes the word "banker", "confidential", or the phrase "insider trading". Only one of these words or phrases must occur for this policy condition to apply. Words in the message or attachment must exactly match what you enter.

> **IMPORTANT**
> When importing a custom dictionary file, each word or phrase must be separated with a carriage return and on a separate line.
> For example:
>
> *banker*
> *confidential*
> *insider trading*

To scan both email messages and attachments for the same keywords, create a data loss prevention policy with a custom keyword dictionary for the terms you wish to scan in messages. This policy configuration identifies defined keywords that appear in either the email message **OR** in the email attachment. Using the standard conditional policy settings (*Message contains any of these words* and *Attachment contains any of these words*) to identify terms in messages and in attachments requires the terms to be present in **BOTH** the message and the attachment.

Enter multiple conditions

If you enter multiple conditions, Microsoft 365 uses all the conditions together to determine when to apply the communication compliance policy to communication items. When you set up multiple conditions, all conditions must be met for the policy to apply, unless you enter an exception. For example, you need a policy that applies if a message contains the word "trade", and is larger than 2 MB. However, if the message also contains the words "Approved by Contoso financial", the policy shouldn't apply. In this example, the three conditions would be defined as follows:

- **Message contains any of these words**, with the keyword "trade"
- **Message size is larger than**, with the value 2 MB
- **Message contains none of these words**, with the keywords "Approved by Contoso financial team"

Review percentage

If you want to reduce the amount of content to review, you can specify a percentage of all the communications governed by a communication compliance policy. A real-time, random sample of content is selected from the total percentage of content that matches chosen policy conditions. If you want reviewers to review all items, you can configure **100%** in a communication compliance policy.

Privacy

Protecting the privacy of users that have policy matches is important and can help promote objectivity in data investigation and analysis reviews for communication compliance alerts. This setting applies only to user names displayed the communication compliance solution. It does not affect how names are displayed in other compliance solutions or admin center.

For users with a communication compliance match, you can choose one of the following settings in **Communication compliance settings**:

- **Show anonymized versions of usernames**: User names are anonymized to prevent users in *Communication Compliance Analyst* role group from seeing who is associated with policy alerts. Users in the *Communication Compliance Investigator* role group will always see user names, not the anonymized versions. For example, a user 'Grace Taylor' would appear with a randomized pseudonym such as 'AnonIS8-988' in all areas of the communication compliance experience. Choosing this setting anonymizes all users with current and past policy matches and applies to all policies. User profile information in the communication compliance alert details will not be available when this option is chosen. However, user names are displayed when adding new users to existing policies or when assigning users to new policies. If you choose to turn off this setting, user names are displayed for all users that have current or past policy matches.
- **Do not show anonymized versions of usernames**: User names are displayed for all current and past policy matches for communication compliance alerts. User profile information (the name, title, alias, and organization or department) is displayed for the user for all communication compliance alerts.

Notice templates

You can create notice templates if you want to send users an email reminder notice for policy matches as part of the issue resolution process. Notices can only be sent to the user email address associated with the policy match that generated the specific alert for remediation. When selecting a notice template to apply to a policy violation as part of the remediation workflow, you can choose to accept the field values defined in the template or overwrite the fields as needed.

Notices templates are custom email templates where you can define the following message fields in the **Communication compliance settings** area:

FIELD	REQUIRED	DETAILS
Template name	Yes	Friendly name for the notice template that you'll select in the notify workflow during remediation, supports text characters.
Sender address	Yes	The address of one or more users or groups that send the message to the user with a policy match, selected from the Active Directory for your subscription.
CC and BCC addresses	No	Optional users or groups to be notified of the policy match, selected from the Active Directory for your subscription.
Subject	Yes	Information that appears in the subject line of the message, supports text characters.

FIELD	REQUIRED	DETAILS
Message body	Yes	Information that appears in the message body, supports text or HTML values.

HTML for notices

If you'd like to create more than a simple text-based email message for notifications, you can create a more detailed message by using HTML in the message body field of a notice template. The following example provides the message body format for a basic HTML-based email notification template:

```
<!DOCTYPE html>
<html>
    <body>
        <h2>Action Required: Contoso Employee Code of Conduct Policy Training</h2>
        <p>A recent message you've sent has generated a policy alert for the Contoso Employee <a
href='https://www.contoso.com'>Code of Conduct Policy</a>.</p>
        <p>You are required to attend the Contoso Employee Code of Conduct <a
href='https://www.contoso.com'>training</a> within the next 14 days. Please contact <a
href='mailto:hr@contoso.com'>Human Resources</a> with any questions about this training request.</p>
        <p>Thank you,</p>
        <p><em>Human Resources</em></p>
    </body>
</html>
```

> **NOTE**
> HTML href attribute implementation in the communication compliance notification templates currently support only single quotation marks instead of double quotation marks for URL references.

Filters

Communication compliance filters allow you to filter and sort alert messages for quicker investigation and remediation actions. Filtering is available on the **Pending** and **Resolved** tabs for each policy. To save a filter or filter set as a saved filter query, one or more values must be configured as filter selections. The following table outlines filter details:

FILTER	DETAILS
Date	The date the message was sent or received by a user in your organization. To filter for a single day, select a date range that starts with the day you want results for and end with the following day. For example, if you wanted to filter results for 9/20/2020, you would choose a filter date range of 9/20/2020-9/21/2020.
File class	The class of the message based on the message type, either *message* or *attachment*.
Has attachment	The attachment presence in the message.
Item class	The source of the message based on the message type, email, Microsoft Team chat, Bloomberg, etc. For more information on common Item Types and Message Classes, see Item Types and Message Classes.

FILTER	DETAILS
Recipient domains	The domain to which the message was sent. This domain is normally your Microsoft 365 subscription domain by default.
Recipient	The user to which the message was sent.
Sender	The person who sent the message.
Sender domain	The domain that sent the message.
Size	The size of the message in KB.
Subject/Title	The message subject or chat title.
Tags	The tags assigned to a message, either *Questionable*, *Compliant*, or *Non-compliant*.
Escalated To	The user name of the person included as part of a message escalation action.
Classifiers	The name of built-in and custom classifiers that apply to the message. Some examples include *Offensive Language*, *Targeted Harassment*, *Profanity*, *Threat*, and more.

Alert policies

After you configure a policy, a corresponding alert policy is automatically created and alerts are generated for messages that match conditions defined in the policy. By default, all policy matches alert triggers are assigned a severity level of medium in the associated alert policy. Alerts are generated for a communication compliance policy once the aggregation trigger threshold level is met in the associated alert policy.

For communication compliance policies, the following alert policy values are configured by default:

ALERT POLICY TRIGGER	DEFAULT VALUE
Aggregation	Simple aggregation
Threshold	4 activities
Window	60 minutes

> **NOTE**
> The alert policy threshold trigger settings for activities supports a minimum value of 3 or higher for communication compliance policies.

You can change the default settings for triggers on number of activities, period for the activities, and for specific users in alert policies on the **Alert policies** page in the Security & Compliance Center.

Change the severity level for an alert policy

If you'd like to change the severity level assigned in an alert policy for a specific communication compliance policy, complete the following steps:

1. Sign into https://compliance.microsoft.com using credentials for an admin account in your Microsoft 365 organization.
2. In the Microsoft 365 compliance center, go to **Policies**.
3. Select **Office 365 alert** on the **Policies** page to open the **Alerts policies** page in the **Office 365 Security & Compliance center**.
4. Select the checkbox for the communication compliance policy you want to update, then select **Edit policy**.
5. On the **Description** tab, select the **Severity** dropdown to configure the policy alert level.
6. Select **Save** to apply the new severity level to the policy.
7. Select **Close** to exit the alert policy details page.

Power Automate flows

Microsoft Power Automate is a workflow service that automates actions across applications and services. By using flows from templates or created manually, you can automate common tasks associated with these applications and services. When you enable Power Automate flows for communication compliance, you can automate important tasks for alerts and users. You can configure Power Automate flows to notify managers when users have communication compliance alerts and other applications.

Customers with Microsoft 365 subscriptions that include communication compliance do not need additional Power Automate licenses to use the recommended default communication compliance Power Automate template. The default template can be customized to support your organization and cover core communication compliance scenarios. If you choose to use premium Power Automate features in these templates, create a custom template using the Microsoft 365 compliance connector, or use Power Automate templates for other compliance areas in Microsoft 365, you may need additional Power Automate licenses.

> **IMPORTANT**
>
> Are you receiving prompts for additional license validation when testing Power Automate flows? Your organization may not have received service updates for this preview feature yet. Updates are being deployed and all organizations with Microsoft 365 subscriptions that include communication compliance should have license support for flows created from the recommended Power Automate templates by October 30, 2020.

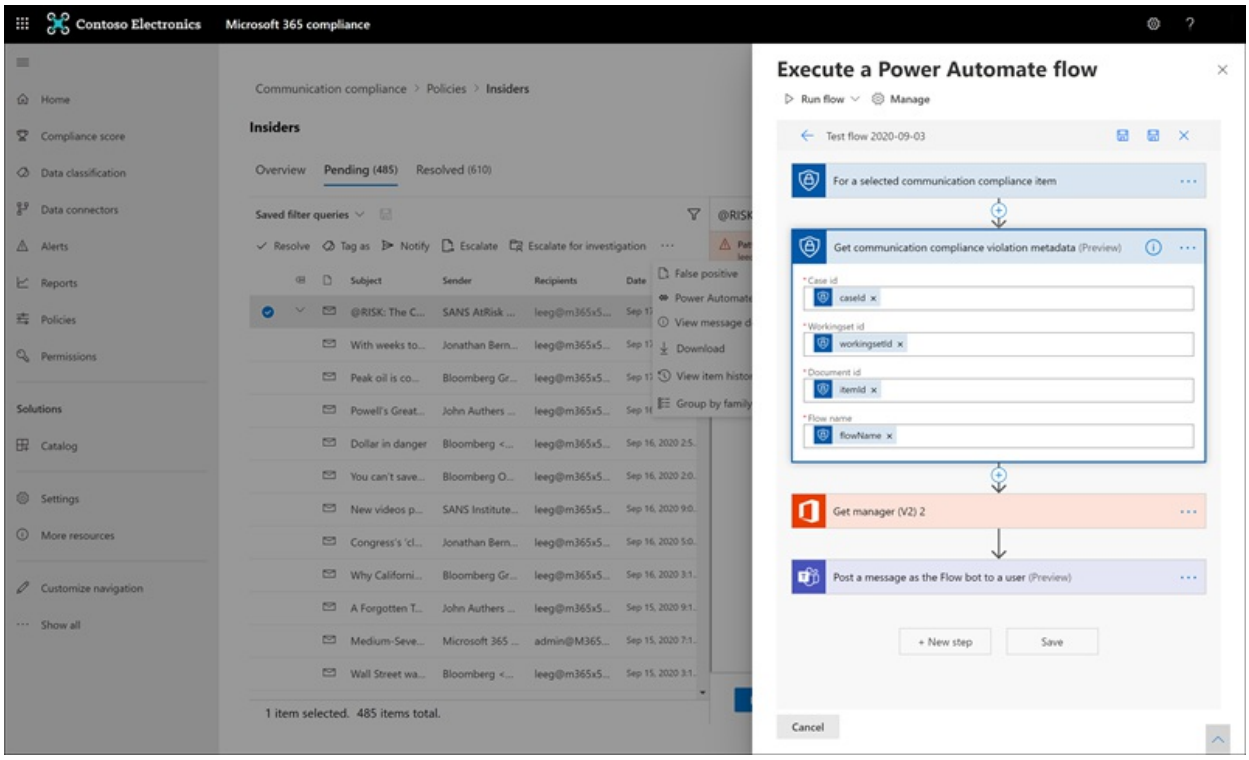

The following Power Automate template is provided to customers to support process automation for communication compliance alerts:

- **Notify manager when a user has a communication compliance alert**: Some organizations may need to have immediate management notification when a user has a communication compliance alert. When this flow is configured and selected, the manager for the case user is sent an email message with the following information about all alerts:
  - Applicable policy for the alert
  - Date/Time of the alert
  - Severity level of the alert

Create a Power Automate flow

To create a Power Automate flow from a recommended default template, you'll use the **Manage Power Automate flows** option from the **Automate** control when working directly in an alert. To create a Power Automate flow with **Manage Power Automate flows**, you must be a member of at least one communication compliance role group.

Complete the following steps to create a Power Automate flow from a default template:

1. In the Microsoft 365 compliance center, go to **Communication compliance** > **Policies** and select the policy with the alert you want review.
2. From the policy, select the **Pending** tab and select a pending alert.
3. Select **Power Automate** from the alert action menu.
4. On the **Power Automate** page, select a default template from the **Communication compliance templates you may like** section on the page.
5. The flow will list the embedded connections needed for the flow and will display if the connection statuses are available. If needed, update any connections that aren't displayed as available. Select **Continue**.
6. By default, the recommended flows are pre-configured with the recommended communication compliance and Microsoft 365 service data fields required to complete the assigned task for the flow. If needed, customize the flow components by using the **Show advanced options** control and configuring the available properties for the flow component.
7. If needed, add any additional steps to the flow by selecting the **New step** button. In most cases, this change should not be needed for the recommended default templates.

8. Select **Save draft** to save the flow for further configuration later, or select **Save** to complete the configuration for the flow.
9. Select **Close** to return to the Power Automate flow page. The new template will be listed as a flow on the **My flows** tab and is automatically available from the Power Automate control for the user that created the flow when working with communication compliance alerts.

Share a Power Automate flow

By default, Power Automate flows created by a user are only available to that user. For other communication compliance users to have access and use a flow, the flow must be shared by the flow creator. To share a flow, you'll use the **Power Automate** control when working directly in an alert.

To share a Power Automate flow, you must be a member of at least one communication compliance role group. Complete the following steps to share a Power Automate flow:

1. In the Microsoft 365 compliance center, go to **Communication compliance** > **Policies** and select the policy with the alert you want review.
2. From the policy, select the **Pending** tab and select a pending alert.
3. Select **Power Automate** from the alert action menu.
4. On the **Power Automate flows** page, select the **My flows** or **Team flows** tab.
5. Select the flow to share, then select **Share** from the flow options menu.
6. On the flow sharing page, enter the name of the user or group you want to add as an owner for the flow.
7. On the **Connection Used** dialog, select **OK** to acknowledge that the added user or group will have full access to the flow.

Edit a Power Automate flow

If you need to edit a flow, you'll use the **Power Automate** control when working directly in an alert. To edit a Power Automate flow, you must be a member of at least one communication compliance role group.

Complete the following steps to edit a Power Automate flow:

1. In the Microsoft 365 compliance center, go to **Communication compliance** > **Policies** and select the policy with the alert you want review.
2. From the policy, select the **Pending** tab and select a pending alert.
3. Select **Power Automate** from the alert action menu.
4. On the **Power Automate flows** page, select flow to edit. Select **Edit** from the flow control menu.
5. Select the **ellipsis** > **Settings** to change a flow component setting or **ellipsis** > **Delete** to delete a flow component.
6. Select **Save** and then **Close** to complete editing the flow.

Delete a Power Automate flow

If you need to delete a flow, you'll use the **Power Automate** control when working directly in an alert. To delete a Power Automate flow, you must be a member of at least one communication compliance role group.

Complete the following steps to delete a Power Automate flow:

1. In the Microsoft 365 compliance center, go to **Communication compliance** > **Policies** and select the policy with the alert you want review.
2. From the policy, select the **Pending** tab and select a pending alert.
3. Select **Power Automate** from the alert action menu.
4. On the **Power Automate flows** page, select flow to delete. Select **Delete** from the flow control menu.
5. On the deletion confirmation dialog, select **Delete** to remove the flow or select **Cancel** to exit the deletion action.

Reports (preview)

The new **Reports** dashboard is the central location for viewing all communication compliance reports. Report widgets provide a quick view of insights most commonly needed for an overall assessment of the status of communication compliance activities. Information contained in the report widgets is not exportable. Detailed reports provide in-depth information related to specific communication compliance areas and offer the ability to filter, group, sort, and export information while reviewing.

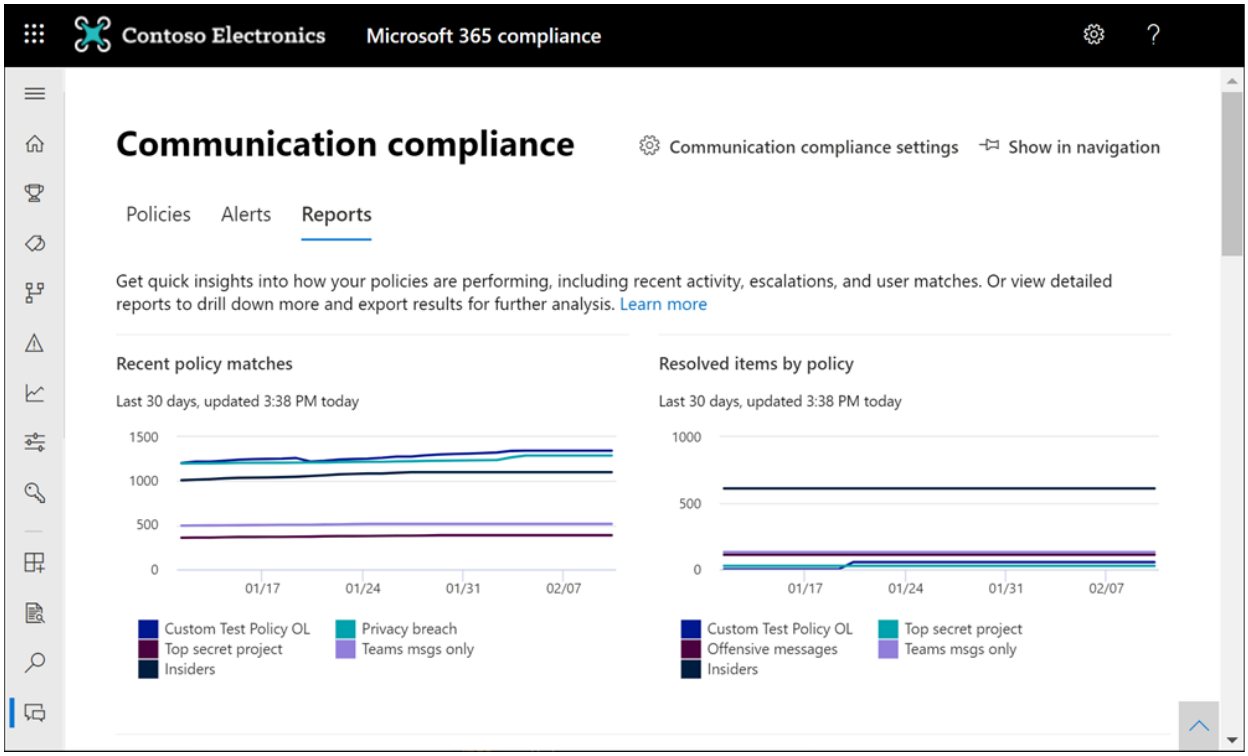

The **Reports dashboard** contains the following report widgets and detailed reports links:

- **Recent policy matches** widget: displays the number of matches by active policy over time.
- **Resolved items by policy** widget: displays the number of policy match alerts resolved by policy over time.
- **Users with most policy match** widget: displays the users (or anonymized usernames) and number of policy matches for a given period.
- **Policy with most matches** widget: displays the policies and the number of matches for a given period, ranked highest to lowest for matches.
- **Escalations by policy** widget: displays the number of escalations per policy over a given time.
- **Policy settings and status** detailed report: provides a detailed look at policy configuration and settings, as well as the general status for each of the policy (matches and actions) on messages. Includes policy information and how policies are associated with users and groups, locations, review percentages, reviewers, status, and when the policy was last modified. Use the *Export* option to create a .CSV file containing the report details.
- **Items and actions per policy** detailed report: Review and export matching items and remediation actions per policy. Includes policy information and how policies are associated with:
  - Items matched
  - Escalated items
  - Resolved items
  - Tagged as compliant

- Tagged as non-compliant
- Tagged as questionable
- Items pending review
- User notified
- Case created

Use the *Export* option to create a .csv file containing the report details.

- **Item and actions per location** detailed report: Review and export matching items and remediation actions per Microsoft 365 location. Includes information about how workload platforms are associated with:

  - Items matched
  - Escalated items
  - Resolved items
  - Tagged as compliant
  - Tagged as non-compliant
  - Tagged as questionable
  - Items pending review
  - User notified
  - Case created

  Use the *Export* option to create a .csv file containing the report details.

- **Activity by user** detailed report: Review and export matching items and remediation actions per user. Includes information about how users are associated with:

  - Items matched
  - Escalated items
  - Resolved items
  - Tagged as compliant
  - Tagged as non-compliant
  - Tagged as questionable
  - Items pending review
  - User notified
  - Case created

  Use the *Export* option to create a .csv file containing the report details.

Audit

In some instances, you must provide information to regulatory or compliance auditors to prove supervision of user activities and communications. This information may be a summary of all activities associated with a defined organizational policy or anytime a communication compliance policy changes. Communication compliance policies have built-in audit trails for complete readiness for internal or external audits. Detailed audit histories of every create, edit, and delete action are captured by your communication policies to provide proof of supervisory procedures.

> **IMPORTANT**
> Auditing must be enabled for your organization before communication compliance events will be recorded. To enable auditing, see Enable the audit log.

To view communication compliance policy update activities, select the **Export policy updates** control on the

main page for any policy. You must be assigned the *Global Admin* or *Communication Compliance Admin* roles to export update activities. This action generates an audit file in the .csv format that contains the following information:

FIELD	DETAILS
CreationDate	The date the update activity was performed in a policy.
UserIds	The user that performed the update activity in a policy.
Operations	The update operations performed on the policy.
AuditData	This field is the main data source for all policy update activities. All update activities are recorded and separated by comma delimiters.

To view communication compliance review activities for a policy, select the **Export review activities** control on the **Overview** page for a specific policy. You must be assigned the *Global Admin* or *Communication Compliance Admin* roles to export review activities. This action generates an audit file in the .csv format that contains the following information:

FIELD	DETAILS
CreationDate	The date the review activity was performed in a policy.
UserIds	The user that performed the review activity in a policy.
Operations	The review operations performed on the policy.
AuditData	This field is the main data source for all policy review activities. All review activities are recorded and separated by comma delimiters.

You can also view audit activities in the unified audit log or with the Search-UnifiedAuditLog PowerShell cmdlet.

For example, the following example returns the activities for all the supervisory review activities (policies and rules):

```
Search-UnifiedAuditLog -StartDate $startDate -EndDate $endDate -RecordType AeD -Operations
SupervisoryReviewTag
```

This example returns the update activities for your communication compliance policies:

```
Search-UnifiedAuditLog -StartDate $startDate -EndDate $endDate -RecordType Discovery -Operations
SupervisionPolicyCreated,SupervisionPolicyUpdated,SupervisionPolicyDeleted
```

Ready to get started?

To configure communication compliance for your Microsoft 365 organization, see Configure communication compliance for your Microsoft 365 organization.

Case study - Contoso quickly configures an offensive language policy for Microsoft Teams, Exchange, and Yammer communications

2/18/2021 • 11 minutes to read • Edit Online

Communication compliance in Microsoft 365 helps minimize communication risks by helping you detect, capture, and act on inappropriate messages in your organization. Pre-defined and custom policies allow you to scan internal and external communications for policy matches so they can be examined by designated reviewers. Reviewers can investigate scanned email, Microsoft Teams, Yammer, or third-party communications in your organization and take appropriate remediation actions to make sure they're compliant with your organization's message standards.

The Contoso Corporation is a fictional organization that needs to quickly configure a policy to monitor for offensive language. They have been using Microsoft 365 primarily for email, Microsoft Teams, and Yammer support for their users but have new requirements to enforce company policy around workplace harassment. Contoso IT administrators and compliance specialists have a basic understanding of the fundamentals of working with Microsoft 365 and are looking for end-to-end guidance for how to quickly get started with communication compliance.

This case study will cover the basics for quickly configuring a communication compliance policy to monitor communications for offensive language. This guidance includes:

- Step 1 - Planning for communication compliance
- Step 2 - Accessing communication compliance in Microsoft 365
- Step 3 - Configuring prerequisites and creating a communication compliance policy
- Step 4 - Investigation and remediation of alerts

Step 1: Planning for communication compliance

Contoso IT administrators and compliance specialists attended online webinars about compliance solutions in Microsoft 365 and decided that communication compliance policies will help them meet the updated corporate policy requirements for reducing workplace harassment. Working together, they've developed a plan to create and enable a communication compliance policy that will monitor for offensive language for chats sent in Microsoft Teams, private messages and community conversations in Yammer, and in email messages sent in Exchange Online. Their plan includes identifying:

- The IT administrators that need access to communication compliance features.
- The compliance specialists that need to create and manage communication policies.
- The compliance specialists and other colleague in other departments (Human Resources, Legal, etc.) that need to investigate and remediate communication compliance alerts.
- The users that will be in-scope for the communication compliance offensive language policy.

Licensing

The first step is to confirm that Contoso's Microsoft 365 licensing includes support for the communication compliance solution. To access and use communication compliance, Contoso IT administrators need to verify that Contoso has one of the following:

- Microsoft 365 E5 subscription (paid or trial version)
- Microsoft 365 E3 subscription + the Microsoft 365 E5 Compliance add-on

- Microsoft 365 E3 subscription + the Microsoft 365 E5 Insider Risk Management add-on
- Microsoft 365 A5 subscription (paid or trial version)
- Microsoft 365 A3 subscription + the Microsoft 365 A5 Compliance add-on
- Microsoft 365 A3 subscription + the Microsoft 365 A5 Insider Risk Management add-on
- Microsoft 365 G5 subscription (paid or trial version)
- Microsoft 365 G5 subscription + the Microsoft 365 G5 Compliance add-on
- Microsoft 365 G5 subscription + the Microsoft 365 G5 Insider Risk Management add-on
- Office 365 Enterprise E5 subscription (paid or trial version)
- Office 365 Enterprise E3 subscription + the Office 365 Advanced Compliance add-on (no longer available for new subscriptions, see note)

They must also confirm that users included in communication compliance policies must be assigned one of the licenses above.

> **IMPORTANT**
> Office 365 Advanced Compliance is no longer sold as a standalone subscription. When current subscriptions expire, customers should transition to one of the subscriptions above, which contain the same or additional compliance features.

Contoso IT administrators take the following steps to verify the licensing support for Contoso:

1. IT administrators sign in to the **Microsoft 365 admin center** (https://admin.microsoft.com) and navigate to **Microsoft 365 admin center** > **Billing** > **Licenses**.
2. Here they confirm that they have one of the license options that includes support for communication compliance.

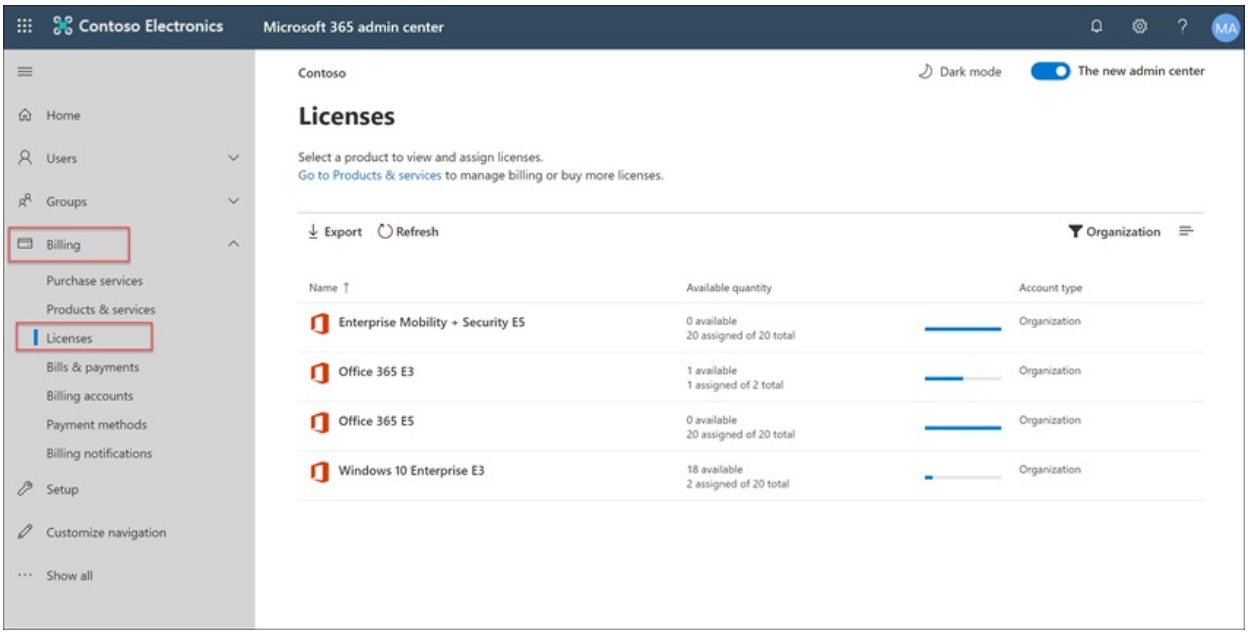

Permissions for communication compliance

There are five role groups used to configure permissions to manage communication compliance features. To make **Communication compliance** available as a menu option in Microsoft 365 compliance center and to continue with these configuration steps, Contoso administrators are assigned the *Communication Compliance Admin* role.

Contoso decides to use the *Communication Compliance* role group assign all the communication compliance administrators, analysts, investigators, and viewers to the group. This makes it easier for Contoso to get started quickly and best fits their compliance management requirements.

ROLE	ROLE PERMISSIONS
Communication Compliance	Use this role group to manage communication compliance for your organization in a single group. By adding all user accounts for designated administrators, analysts, investigators, and viewers, you can configure communication compliance permissions in a single group. This role group contains all the communication compliance permission roles. This configuration is the easiest way to quickly get started with communication compliance and is a good fit for organizations that do not need separate permissions defined for separate groups of users.
Communication Compliance Admin	Use this role group to initially configure communication compliance and later to segregate communication compliance administrators into a defined group. Users assigned to this role group can create, read, update, and delete communication compliance policies, global settings, and role group assignments. Users assigned to this role group cannot view message alerts.
Communication Compliance Analyst	Use this group to assign permissions to users that will act as communication compliance analysts. Users assigned to this role group can view policies where they are assigned as Reviewers, view message metadata (not message content), escalate to additional reviewers, or send notifications to users. Analysts cannot resolve pending alerts.
Communication Compliance Investigator	Use this group to assign permissions to users that will act as communication compliance investigators. Users assigned to this role group can view message metadata and content, escalate to additional reviewers, escalate to an Advanced eDiscovery case, send notifications to users, and resolve the alert.
Communication Compliance Viewer	Use this group to assign permissions to users that will manage communication reports. Users assigned to this role group can access all reporting widgets on the communication compliance home page and can view all communication compliance reports.

1. Contoso IT administrators sign into the **Office 365 Security & Compliance center** permissions page (https://protection.office.com/permissions) using credentials for a global administrator account and select the link to view and manage roles in Microsoft 365.
2. In the **Security & Compliance Center**, they go to **Permissions** and select the link to view and manage roles in Office 365.
3. The administrators select the *Communication Compliance* role group, then select **Edit role group**.
4. The administrators select **Choose members** from the left navigation pane, then select **Edit**.
5. They select **Add** and then select the checkbox for all Contoso users that will manage communication compliance, investigate, and review alerts.
6. The administrators select **Add**, then select **Done**.
7. They select **Save** to add Contoso users to the role group. They select **Close** to complete the steps.

Step 2: Accessing communication compliance in Microsoft 365

After configuring the permissions for communication compliance, Contoso IT administrators and compliance specialists assigned to the Communication Compliance role group can access the communication compliance

solution in Microsoft 365. Contoso IT administrators and compliance specialists have several ways to access communication compliance and get started creating a new policy:

- Starting directly from the communication compliance solution
- Starting from the Microsoft 365 compliance center
- Starting from the Microsoft 365 solution catalog
- Starting from the Microsoft 365 admin center

Starting directly from the communication compliance solution

The quickest way to access the solution is to sign in directly to the **Communication compliance** (https://compliance.microsoft.com/supervisoryreview) solution. Using this link, Contoso IT administrators and compliance specialists will be directed to the communication compliance Overview dashboard where you can quickly review the status of alerts and create new policies from the pre-defined templates.

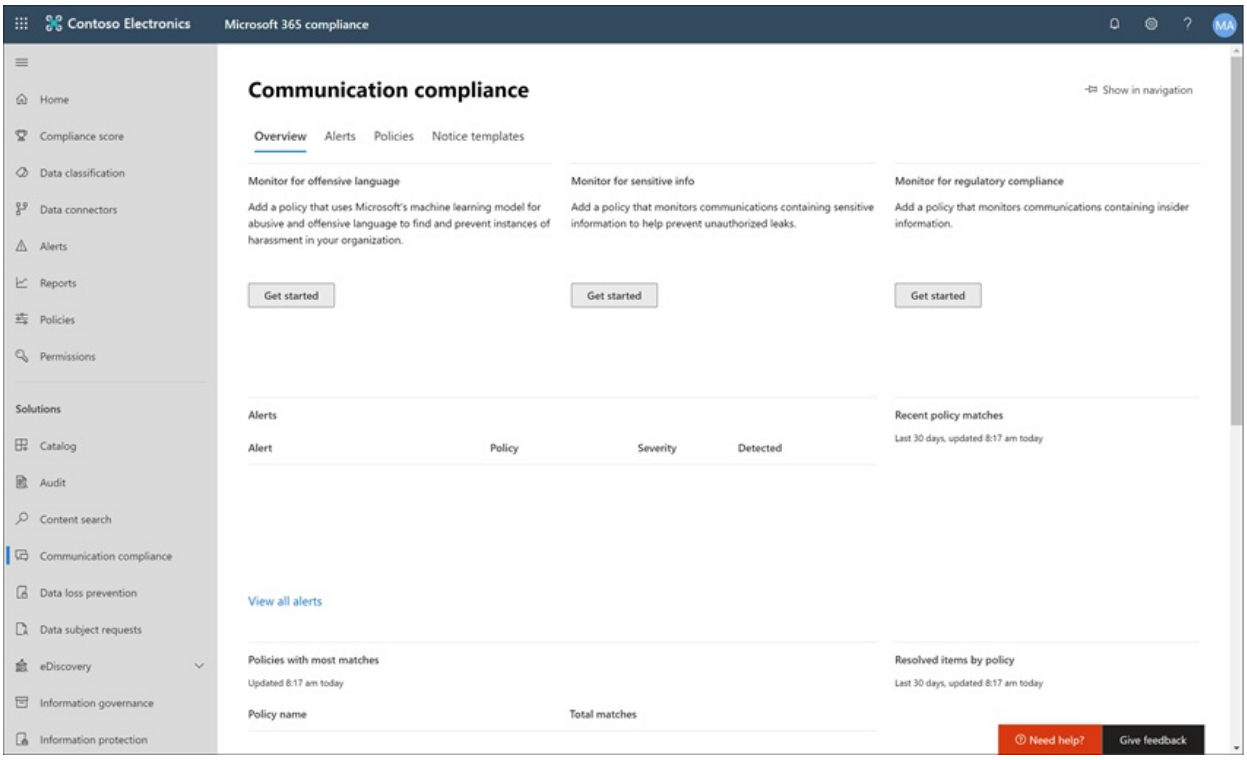

Starting from the Microsoft 365 compliance center

Another easy way for Contoso IT administrators and compliance specialists to access the communication compliance solution is to sign in directly to the **Microsoft 365 compliance center** (https://compliance.microsoft.com). After signing in, users simply need to select the **Show all** control to display all the compliance solutions and then select the **Communication compliance** solution to get started.

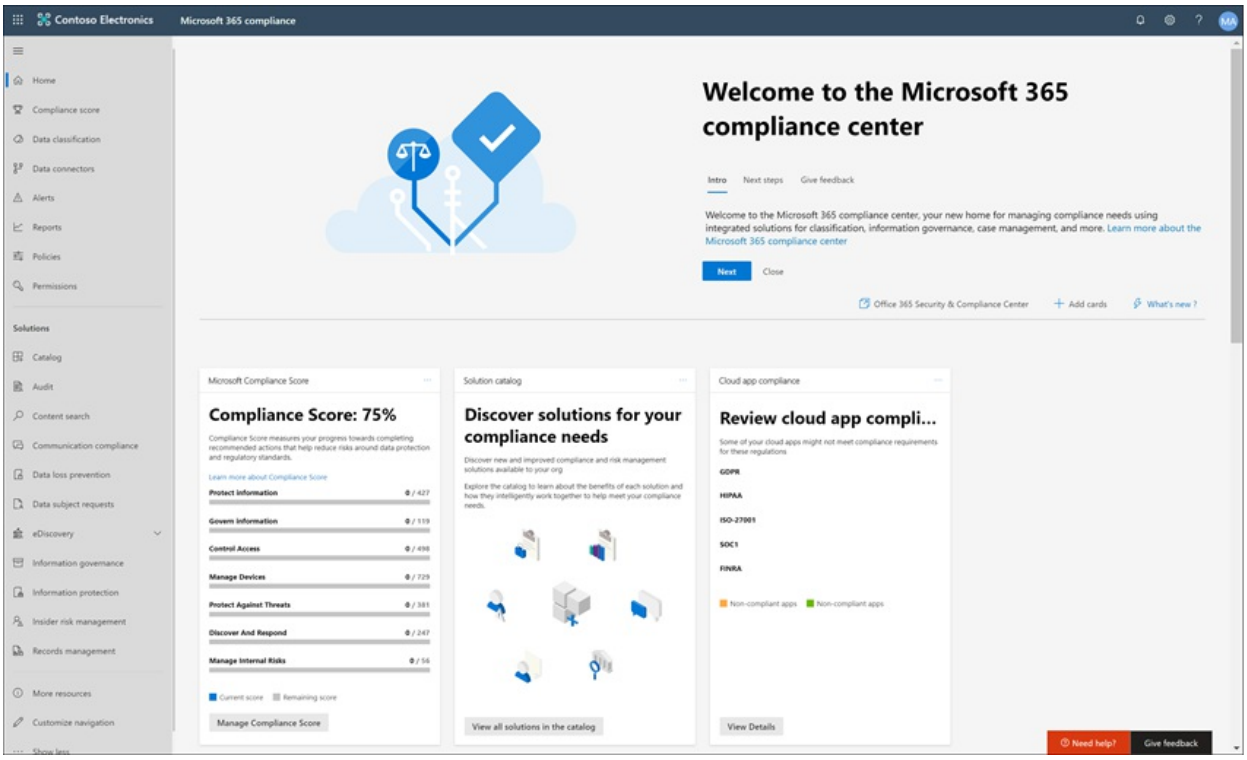

Starting from the Microsoft 365 solution catalog

Contoso IT administrators and compliance specialists could also choose to access the communication compliance solution by selecting the Microsoft 365 solution catalog. By selecting **Catalog** in **Solutions** section of the left navigation while in the **Microsoft 365 compliance center**, they can open the solution catalog listing all Microsoft 365 compliance solutions. Scrolling down to the **Insider risk management** section, Contoso IT administrators can select Communication compliance to get started. Contoso IT administrators also decide to use the Show in navigation control to pin the communication compliance solution to the left-navigation pane for quicker access when they sign in going forward.

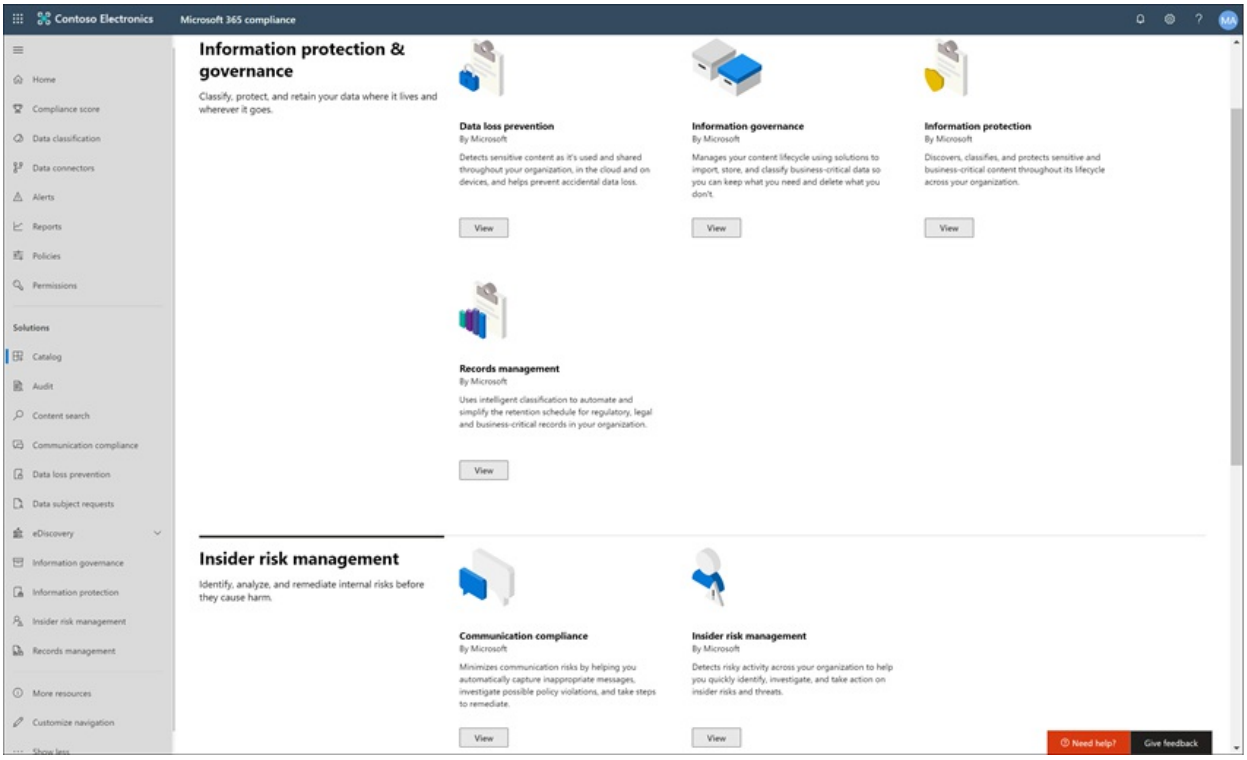

Starting from the Microsoft 365 admin center

To access communication compliance when starting from the Microsoft 365 admin center, Contoso IT administrators and compliance specialists sign in to the Microsoft 365 admin center (https://admin.microsoft.com) and navigate to **Microsoft 365 admin center** > **Compliance**.

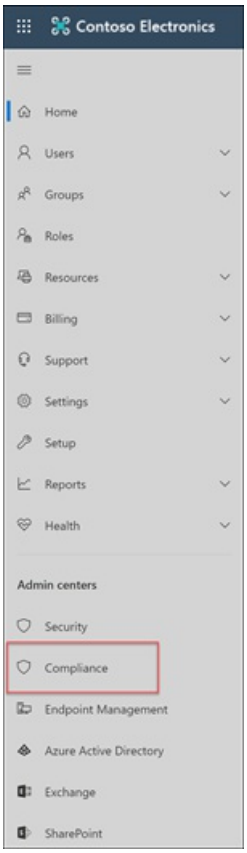

This action opens the **Office 365 Security and Compliance center**, and they must select the link to the **Microsoft 365 compliance center** provided in the banner at the top of the page.

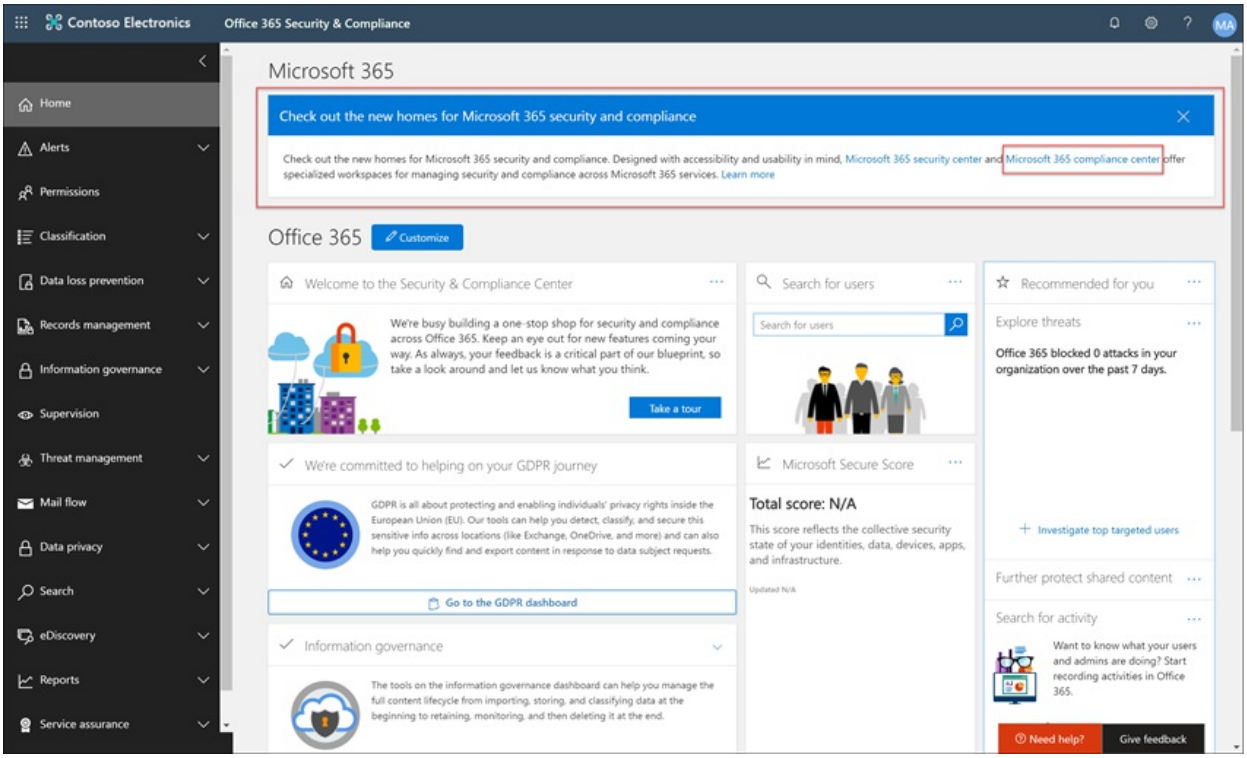

Once in the **Microsoft 365 compliance center**, Contoso IT administrators select **Show all** to display the full list of compliance solutions.

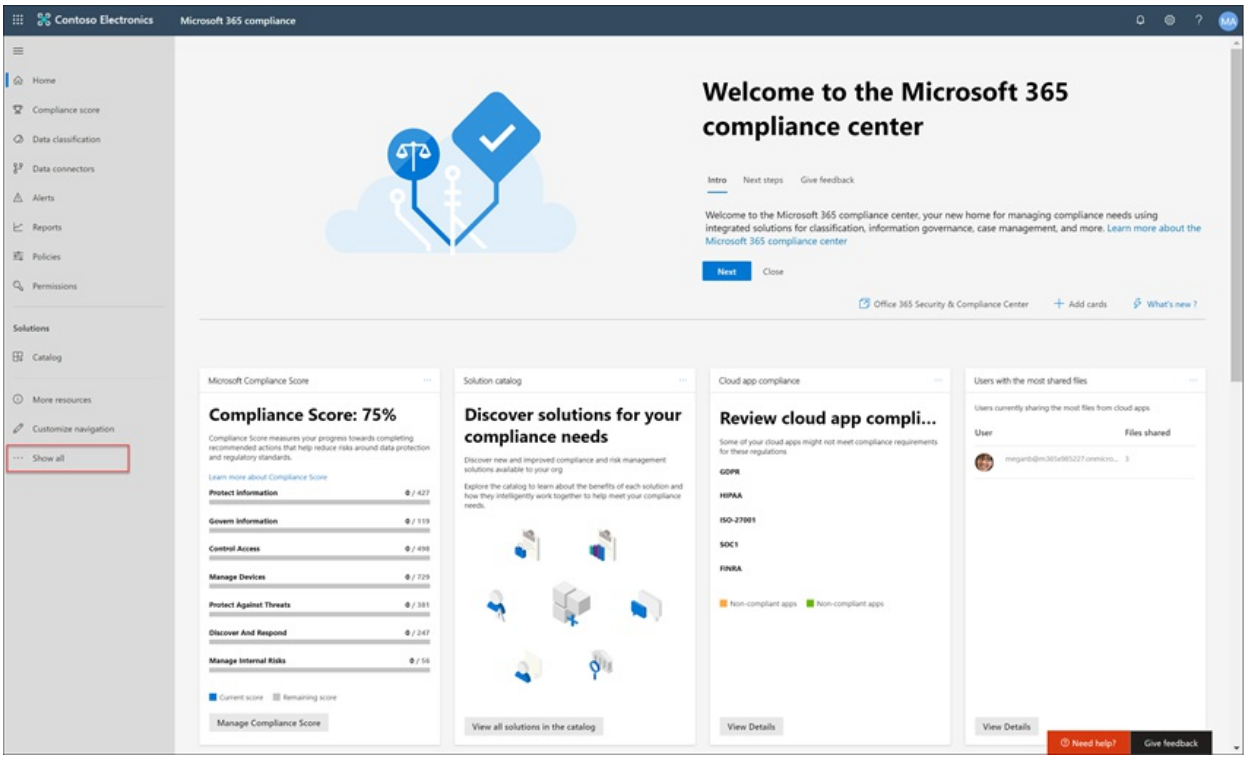

After selecting **Show all**, the Contoso IT administrators can access the communication compliance solution.

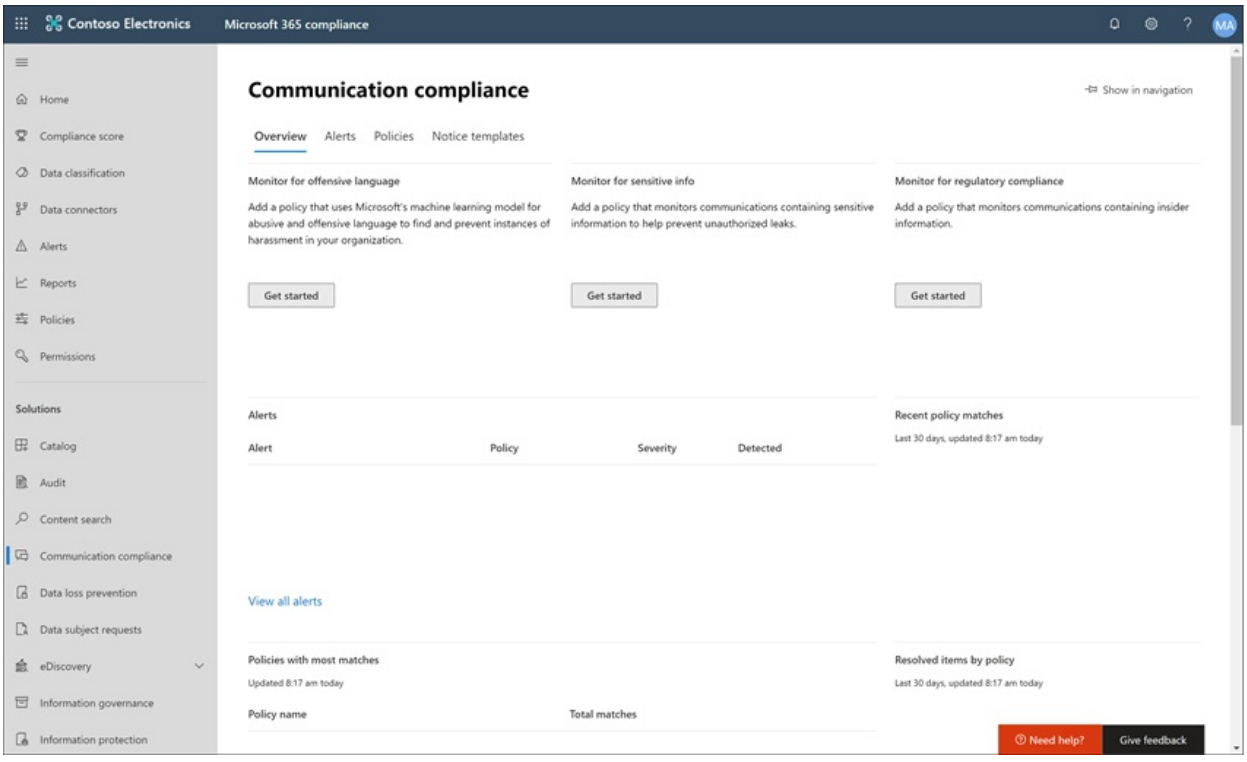

Step 3: Configuring prerequisites and creating a communication compliance policy

To get started with a communication compliance policy, there are several prerequisites that Contoso IT administrators need to configure before setting up the new policy to monitor for offensive language. After these prerequisites have been completed, Contoso IT administrators and compliance specialists can configure the new policy and compliance specialists can start investigation and remediating any generated alerts.

Enabling auditing in Microsoft 365

Communication compliance requires audit logs to show alerts and track remediation actions taken by reviewers. The audit logs are a summary of all activities associated with a defined organizational policy or anytime there is

a change to a communication compliance policy.

Contoso IT administrators review and complete the step-by-step instructions to turn on auditing. After they turn on auditing, a message is displayed that says the audit log is being prepared and that they can run a search in a couple of hours after the preparation is complete. The Contoso IT administrators only have to do this action once.

Configuring Yammer tenant for Native Mode

Communication compliance requires that the Yammer tenant for an organization is in Native Mode to monitor for offensive language in private messages and public community conversations.

Contoso IT administrators make sure they review the information in the Overview of Yammer Native Mode in Microsoft 365 article and follow the steps for running the migration tool in the Configure your Yammer network for Native Mode for Microsoft 365 article.

Setting up a group for in-scope users

Contoso compliance specialists want to add all users to the communication policy that will monitor for offensive language. They could decide to add each user account to the policy separately, but they've decided it is much easier and saves time to use an **All Users** distribution group for the users for this policy.

They need to create a new group to include all Contoso users, so they take the following steps:

1. Contoso IT administrators IT sign in to the **Microsoft 365 admin center** (https://admin.microsoft.com) and navigate to **Microsoft 365 admin center** > **Groups** > **Groups**.
2. They select **Add a group** and complete the wizard to create a new *Microsoft 365 group* or *Distribution group*.

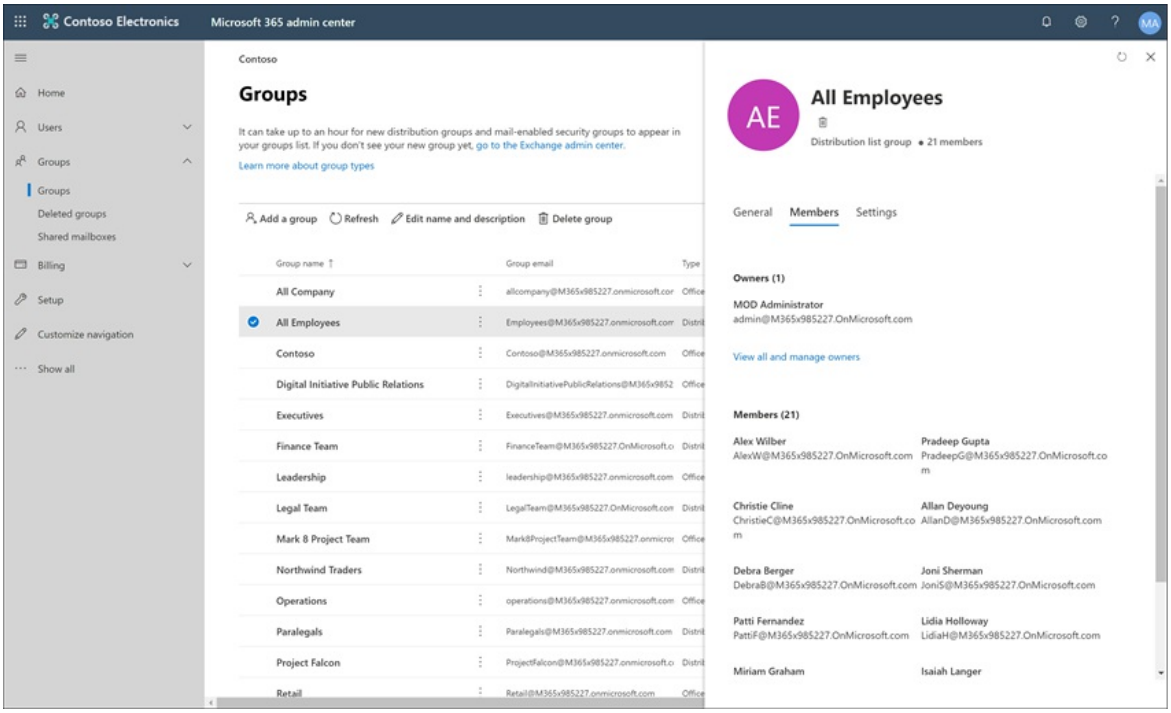

3. After the new group is created, they need to add all Contoso users to the new group. They open the **Exchange admin center** (https://outlook.office365.com/ecp) and navigate to **Exchange admin center** > **recipients** > **groups**. The Contoso IT administrators select the Membership area and the new *All Employees* group they created and select the **Edit** control to add all Contoso users to the new group in the wizard.

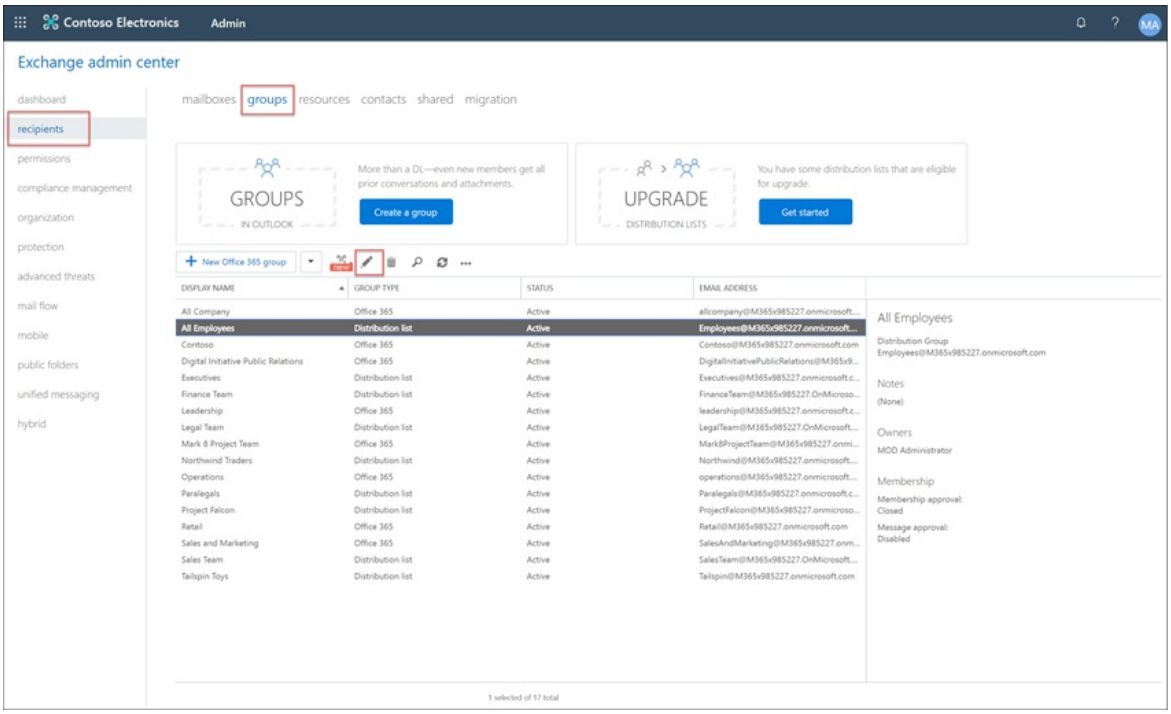

Creating the policy to monitor for offensive language

With all the prerequisites completed, the IT administrators and the compliance specialists for Contoso are ready to configure the communication compliance policy to monitor for offensive language. Using the new offensive language policy template, configuring this policy is simple and quick.

1. The Contoso IT administrators and compliance specialists sign into the **Microsoft 365 compliance center** and select **Communication compliance** from the left navigation pane. This action opens the **Overview** dashboard that has quick links for communication compliance policy templates. They choose the **Monitor for offensive language** template by selecting **Get started** for the template.

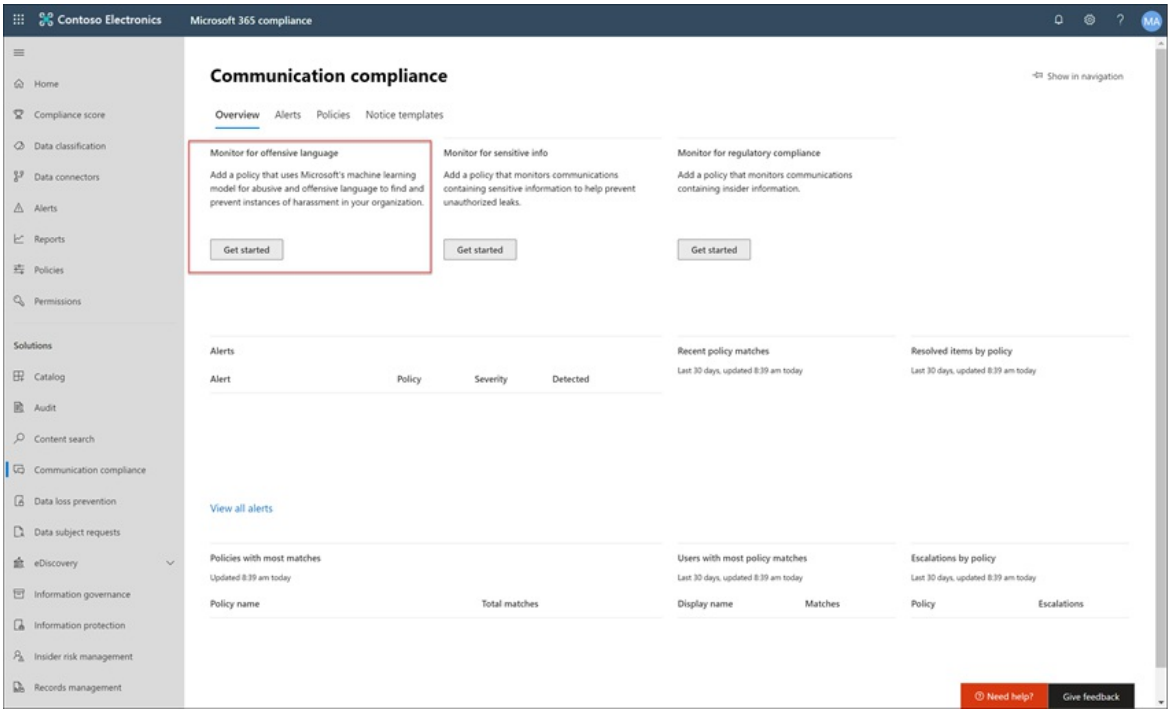

2. On the policy template wizard, the Contoso IT administrators and compliance specialists work together to complete the three required fields: **Policy name**, **Users or groups to supervise**, and **Reviewers**.

3. Since the policy wizard has already suggested a name for the policy, the IT administrators and compliance specialists decide to keep the suggested name and focus on the remaining fields. They select the *All users* group for the **Users or groups to supervise** field and select the compliance specialists that should

investigate and remediate policy alerts for the **Reviewers** field. The last step to configure the policy and start gathering alert information is to select **Create policy**.

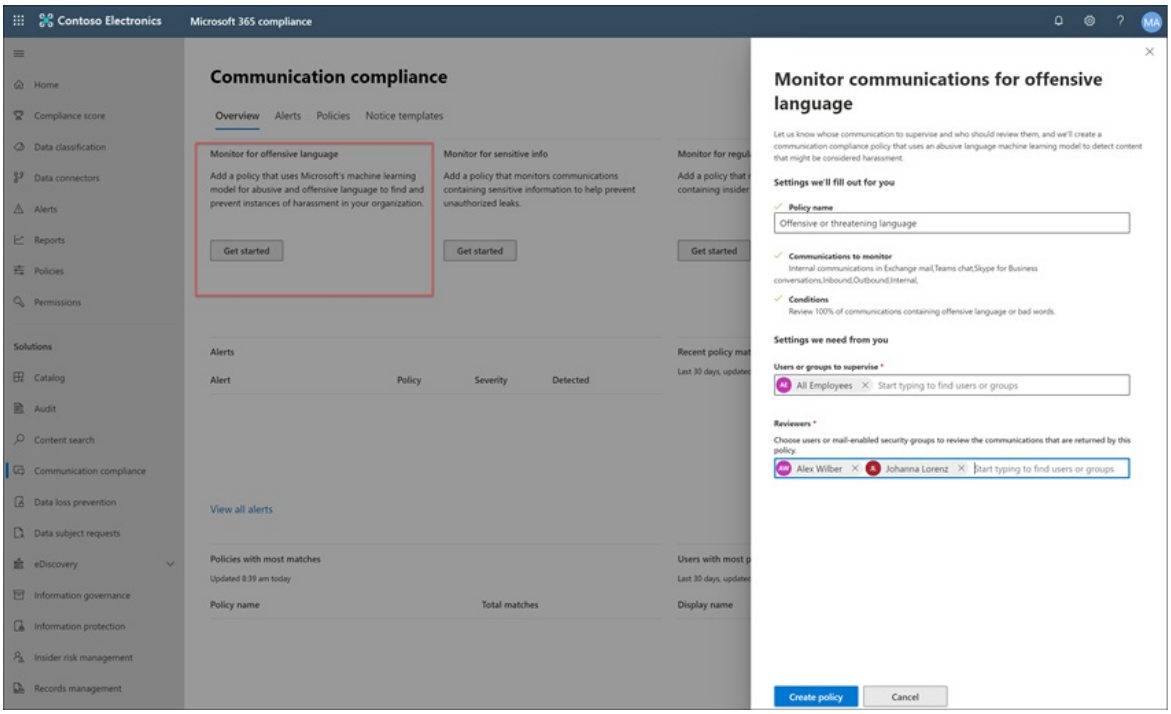

Step 4: Investigate and remediate alerts

Now that the communication compliance policy to monitor for offensive language is configured, the next step for the Contoso compliance specialists will be to investigate and remediate any alerts generated by the policy. It will take up to 24 hours for the policy to fully process communications in all the communication source channels and for alerts to show up in the **Alert dashboard**.

After alerts are generated, Contoso compliance specialists will follow the workflow instructions to investigate and remediate offensive language issues.

Supervision policies

11/2/2020 • 20 minutes to read • Edit Online

IMPORTANT

Following the release of communication compliance in Microsoft 365 Compliance in February 2020, Supervision in Office 365 is being retired. Supervision policies will no longer be available for creation, and policies will eventually be removed, after an extended period of read only access.

If you use Supervision, be aware that:

- Beginning June 15th, 2020, tenants will not have the ability to create new Supervision policies.
- Beginning August 31st, 2020, existing policies will stop capturing new messages.
- Beginning October 26th, 2020, existing policies will be deleted.

We actively encourage customers who are currently exploring or using Supervision in Office 365 to use the new communication compliance solution to address your communications monitoring or regulatory requirements with a much richer set of intelligent capabilities.

Supervision policies in Microsoft 365 allow you to capture employee communications for examination by designated reviewers. You can define specific policies that capture internal and external email, Microsoft Teams, or 3rd-party communications in your organization. Reviewers can then examine the messages to make sure that they are compliant with your organization's message standards and resolve them with classification type.

These policies can also help you overcome many modern compliance challenges, including:

- Monitoring increasing types of communication channels
- The increasing volume of message data
- Regulatory enforcement & the risk of fines

In some organizations, there may be a separation of duties between IT support and the compliance management group. Microsoft 365 supports the separation between supervision policy feature configuration and the configuration of policies for captured communications. For example, the IT group for an organization may be responsible for setting up role permissions and groups to support supervision policies that are configured and managed by the organization's compliance team.

For a quick overview of Supervision policies, see the Supervision policy video on the Microsoft Mechanics channel.

Transitioning from Supervision

Organizations using supervision policies and planning to transition to communication compliance policies in Microsoft 365 need to understand these important points:

- The supervision solution in Microsoft 365 will be fully replaced by the communication compliance solution in Microsoft 365. For organizations transitioning to communication compliance from supervision policies, we recommend creating new policies in communication compliance that have the same *conditions* as existing supervision policies to enable new investigation and remediation improvements. When transitioning to communication compliance in Microsoft 365, you should plan to export reporting data from supervision if you have internal compliance retention policy requirements.
- In the interim, organizations can use both solutions side by side until fully migrated, but policies used in each solution must have *unique policy names*. Groups and custom keyword dictionaries can be shared between

solutions during the transition period.

- Messages saved in supervision in Microsoft 365 policy matches cannot be moved or shared into communication compliance in Microsoft 365.

For retirement information for supervision in Office 365, see the Microsoft 365 Roadmap for details.

Scenarios for supervision policies

Supervision policies can assist monitoring communications in your organization in several areas:

- **Corporate policies**

  Employees must comply with acceptable use, ethical standards, and other corporate policies in all their business-related communications. Supervision policies can detect policy violations and help you take corrective actions to help mitigate these types of incidents. For example, you could monitor for potential human resources violations such as harassment or the use of inappropriate or offensive language in employee communications.

- **Risk management**

  Organizations are responsible to all communications distributed throughout their infrastructure and corporate network systems. Using supervision policies to help identify and manage potential legal exposure and risk can help minimize risks before they can damage corporate operations. For example, you could monitor your organization for unauthorized communications for confidential projects such as upcoming acquisitions, mergers, earnings disclosures, reorganizations, or leadership team changes.

- **Regulatory compliance**

  Most organizations must comply with some type of regulatory compliance standards as part of their normal operating procedures. These regulations often require organizations to implement some type of supervisory or oversight process for messaging that is appropriate for their industry. The Financial Industry Regulatory Authority (FINRA) Rule 3110 is a good example of a requirement for organizations to have supervisory procedures in place to monitor the activities of its employees and the types of businesses in which it engages. Another example may be a need to monitor broker-dealers in your organization to safeguard against potential money-laundering, insider trading, collusion, or bribery activities. Supervision policies can help your organization meet these requirements by providing a process to both monitor and report on corporate communications.

Components

Supervision policy

You create supervision policies in the Compliance center. These policies define which communications and users are subject to review in your organization, define custom conditions that the communications must meet, and specifies who should perform reviews. Users included in the Supervisory Review role group can set up policies and anyone who has this role assigned can access the Supervision page in the Compliance center.

Supervised users

Before you start using supervision, you must determine who needs their communications reviewed. In the policy, user email addresses identify individuals or groups of people to supervise. Some examples of these groups are Microsoft 365 Groups, Exchange-based distribution lists, and Microsoft Teams channels. You also can exclude specific users or groups from supervision with a supervised group or a list of groups.

IMPORTANT

Users monitored by supervision policies must have a Microsoft 365 E5 Compliance license, an Office 365 Enterprise E3 license with the Advanced Compliance add-on, or be included in an Office 365 Enterprise E5 subscription, or be included in a Microsoft 365 E5 subscription.If you don't have an existing Enterprise E5 plan and want to try supervision, you can sign up for a trial of Office 365 Enterprise E5.

Reviewers

When you create a supervision policy, you must determine who will perform the reviews of the messages of the supervised users. In the policy, user email addresses identify individuals or groups of people to review supervised communications. All reviewers must have mailboxes hosted on Exchange Online.

Groups for supervised users and reviewers

To simplify your setup, create groups for people who need their communications reviewed and groups for people who review those communications. If you're using groups, you might need several. For example, if you want to monitor communications between two distinct groups of people, or if you want to specify a group that isn't supervised.

When you select a Microsoft 365 group for supervised users, the policy monitors the content of the shared mailbox and the Microsoft Teams channels associated with the group. When you select a distribution list, the policy monitors individual user mailboxes.

Supported communication types

With supervision policies, you can choose to monitor messages in one or more of the following communication platforms:

- **Exchange email:** Mailboxes hosted on Exchange Online as part of your Microsoft 365 subscription are all eligible for message supervision. Emails and attachments matching supervision policy conditions are instantly available for monitoring and in supervision reports. Supported attachment types for supervision are the same as the file types supported for Exchange mail flow rule content inspections.
- **Microsoft Teams:** Chat communications and associated attachments in both public and private Microsoft Teams channels and individual chats can be supervised. Teams chats matching supervision policy conditions are processed once every 24 hours and then are available for monitoring and in supervision reports. Use the following group management configurations to supervise individual user chats and channel communications in Teams:
  - **For Teams chat supervision:** Assign individual users or assign a distribution group to the supervision policy. This configuration is for 1-to-1 or 1-to-many user/chat relationships.
  - **For Teams Channel communications:** Assign every Microsoft Team channel or Microsoft 365 group you want to monitor that contains a specific user to the supervision policy. If you add the same user to other Microsoft Teams channels or Microsoft 365 groups, be sure to add these new channels and groups to the supervision policy.
- **Skype for Business Online:** Chat communications and associated attachments in Skype for Business Online can be supervised. Skype for Business Online chats matching supervision policy conditions are processed once every 24 hours and then are available for monitoring and in supervision reports. Supervised chat conversations are sourced from previous conversations saved in Skype for Business Online. Use the following group management configuration to supervise user chat communications in Skype for Business Online:
  - **For Skype for Business Online chat supervision:** Assign individual users or assign a distribution group to the supervision policy. This configuration is for 1-to-1 or 1-to-many user/chat relationships.
- **Third-party sources:** You can supervise communications from third-party sources (like from Facebook or DropBox) for data imported into mailboxes in your organization. Learn how to import 3rd-party data.

Communications captured across these platforms are retained for seven years for each policy by default, even if users leave your organization and their mailbox is deleted.

Policy settings

Direction

By default, the **Direction is** condition is displayed and can't be removed. Communication direction settings in a policy are chosen individually or together:

- **Inbound**: You can choose **Inbound** to review communications sent **to** the people you chose to supervise **from** people not included in the policy.
- **Outbound**: You can choose **Outbound** if you want to review communications sent **from** the people you chose to supervise **to** people not included in the policy.
- **Internal**: You can choose **Internal** to review communications sent **between** the people you identified in the policy.

Sensitive information types

You have the option of including sensitive information types as part of your supervision policy. Sensitive information types are either pre-defined or custom data types that can help identify and protect credit card numbers, bank account numbers, passport numbers, and more. As a part of data loss prevention (DLP), the sensitive information configuration can use patterns, character proximity, confidence levels, and even custom data types to help identify and flag content that may be sensitive. The default sensitive information types are:

- Financial
- Medical and health
- Privacy
- Custom information type

To learn more about sensitive information details and the patterns included in the default types, see Sensitive information type entity definitions.

Custom keyword dictionaries

Configure custom keyword dictionaries (or lexicons) to provide simple management of keywords specific to your organization or industry. Keyword dictionaries support up to 100KB of terms (post compression) in the dictionary and support any language. The tenant limit is also 100KB after compression. If needed, you can apply multiple custom keyword dictionaries to a single policy or have a single keyword dictionary per policy. These dictionaries are assigned in a supervision policy and can be sourced from a file (such as a .csv or .txt list), or from a list you can import in the Compliance center.

Offensive language

Monitor sent or received email messages in your organization for offensive language. The model uses a combination of machine learning, artificial intelligence, and keywords to identify language in email messages likely to violate anti-harassment and bullying policies. The offensive language model currently supports English keywords and monitors the body of email messages.

> **NOTE**
>
> Create a data loss prevention policy with a custom keyword dictionary of blocked terms if you need to:
>
> - monitor Microsoft Teams communications in your organization for offensive language
> - prevent or block offensive language in communications in your organization

The model does not provide an exhaustive list of offensive language. Further, language and cultural standards continually change, and in light of these realities, Microsoft reserves the right to update the model in its discretion. While the model may assist your organization in monitoring offensive language, the model is not intended to provide your organization's sole means of monitoring or addressing such language. Your

organization, not Microsoft, remains responsible for all decisions related to monitoring and blocking offensive language.

The offensive language model monitors email for sentiment associated with the following types of language:

TYPE	DESCRIPTION
Profanities	Expressions that embarrass most people.
Slurs	Expressions that express prejudice against particular groups (for example, race, ethnicity, sexual orientation, disability).
Taunts	Expressions that taunt, condemn, ridicule, or could potentially cause anger or violence.
Disguised expressions	Expressions for which the meaning or pronunciation is the same as another more offensive term.

Conditional settings

The conditions you choose for the policy apply to communications from both email and 3rd-party sources in your organization (like from Facebook or DropBox).

The following table explains more about each condition.

CONDITION	HOW TO USE THIS CONDITION
Message is received from any of these domains **Message is not received from any of these domains**	Apply the policy to include or exclude specific domains or email addresses in received messages. Enter each domain or email address and separate multiple domains or email addresses with a comma. Each domain or email address entered is applied separately, only one domain or email address must apply for the policy to apply to the message. If you want to monitor all email from a specific domain but want to exclude messages that do not need review (newsletters, announcements, etc.), you must configure the condition a **Message is not received from any of these domains** condition that excludes the email address (example "newsletter@contoso.com").
Message is sent to any of these domains **Message is not sent to any of these domains**	Apply the policy to include or exclude specific domains or email addresses in sent messages. Enter each domain or email address and separate multiple domains or email addresses with a comma. Each domain or email address is applied separately, only one domain or email address must apply for the policy to apply to the message. If you want to monitor all email sent to a specific domain but want to exclude sent messages that do not need review, you must configure two conditions: - A **Message is sent to any of these domains** condition that defines the domain ("contoso.com"), AND - A **Message is not sent to any of these domains** condition that excludes the email address ("subscriptions@contoso.com").

CONDITION	HOW TO USE THIS CONDITION
Message is classified with any of these labels **Message is not classified with any of these labels**	To apply the policy when certain retention labels are included or excluded in a message. Retention labels must be configured separately and configured labels are chosen as part of this condition. Each label you choose is applied separately (only one of these labels must apply for the policy to apply to the message). For more information about retention labels, see Learn about retention policies and retention labels.
Message contains any of these words **Message contains none of these words**	To apply the policy when certain words or phrases are included or excluded in a message, enter each word or phrase and separate with a comma. Each word you enter is applied separately (only one word must apply for the policy to apply to the message). For more information about entering words or phrases, see the next section Matching words and phrases to emails or attachments.
Attachment contains any of these words **Attachment contains none of these words**	To apply the policy when certain words or phrases are included or excluded in a message attachment (such as a Word document), enter each word or phrase and separate with a comma. Each word you enter is applied separately (only one word must apply for the policy to apply to the attachment). For more information about entering words or phrases, see the next section Matching words and phrases to emails or attachments.
Attachment is any of these file types **Attachment is none of these file types**	To supervise communications that include or exclude specific types of attachments, enter the file extensions (such as .exe or .pdf). If you want to include or exclude multiple file extensions, enter these on separate lines. Only one attachment extension must match for the policy to apply.
Message size is larger than **Message size is not larger than**	To review messages based on a certain size, use these conditions to specify the maximum or minimum size a message can be before it is subject to review. For example, if you specify **Message size is larger than** > **1.0 MB**, all messages that are 1.01 MB and larger are subject to review. You can choose bytes, kilobytes, megabytes, or gigabytes for this condition.
Attachment is larger than **Attachment is not larger than**	To review messages based on the size of their attachments, specify the maximum or minimum size an attachment can be before the message and its attachments are subject to review. For example, if you specify **Attachment is larger than** > **2.0 MB**, all messages with attachments 2.01 MB and over are subject to review. You can choose bytes, kilobytes, megabytes, or gigabytes for this condition.

Matching words and phrases to emails or attachments

Each word you enter and separate with a comma is applied separately (only one word must apply for the policy condition to apply to the email or attachment). For example, let's use the condition, **Message contains any of these words**, with the keywords "banker" and "insider trading" separated by a comma (banker, insider trading). The policy applies to any messages that includes the word "banker" or the phrase "insider trading". Only one of these words or phrases must occur for this policy condition to apply. Words in the message or attachment must exactly match what you enter.

To scan both email messages and attachments for the same keywords, create a data loss prevention policy with a custom keyword dictionary for the terms you wish to monitor. This policy configuration identifies defined

keywords that appear in either the email message **OR** in the email attachment. Using the standard conditional policy settings (*Message contains any of these words* and *Attachment contains any of these words*) to identify terms in messages and in attachments requires the terms are present in **BOTH** the message and the attachment.

Enter multiple conditions

If you enter multiple conditions, Microsoft 365 uses all the conditions together to determine when to apply the policy to communication items. When you set up multiple conditions, all conditions must be met for the policy to apply, unless you enter an exception. For example, you need a policy that applies if a message contains the word "trade", and is larger than 2 MB. However, if the message also contains the words "Approved by Contoso financial", the policy should not apply. Thus, in this case, the three conditions would be as follows:

- **Message contains any of these words**, with the keyword "trade"
- **Message size is larger than**, with the value 2 MB
- **Message contains none of these words**, with the keywords "Approved by Contoso financial team"

Review percentage

If you want to reduce the amount of content to review, you can specify a percentage of all the communications governed by a supervision policy. A real-time, random sample of content is selected from the total percentage of content that matches chosen policy conditions. If you want reviewers to review all items, you can enter **100%** in a supervision policy.

Monitor & manage

It is easy to monitor the results of your supervision policies and apply a resolution tag. You can quickly see:

- The status of reviewed items
- Users and groups under supervision
- Users and groups designated as reviewers

Supervision policy dashboard

Use the supervision policy dashboard to manage supervision policy results and to resolve outstanding items. This dashboard allows reviewers to view items that need to be reviewed, act on an item, and review the results of previously reviewed and resolved items for each supervision policy. You can access the supervision policy dashboard in the Compliance center at **Supervision** > *Your Custom Policy* > **Open**.

Dashboard Home

The dashboard **Home** page has several sections to help you quickly act on your supervision policies. Here you can:

- Quickly review the pending and resolved highlights for the week
- See a list of the supervised users and supervised groups for the selected policy
- See a list of the reviewers and review teams for the selected policy
- See which communication platforms have content under supervision for the policy

Review tab

The **Review** tab is where reviewers classify and resolve items identified by the selected policy. Here you can:

- Filter by pending, compliant, non-compliant, and questionable items.
- Tag a single item as compliant, non-compliant, or questionable. You can also record a comment with the item to help clarify the tagging action taken.
- Bulk tag multiple items as compliant, non-compliant, or questionable. You can also record a comment with multiple items to help clarify the tagging action taken.
- View the history of the tagging for a single item. Includes who resolved the item, the date and time of the action, the resolution tag, and any included comments.

- Reclassify previously reviewed items as compliant, non-compliant, or questionable. You can also record a comment with single or multiple items to help clarify the reclassification action taken.

Resolved Items tab

The **Resolved Items** tab is where reviewers can view all previously resolved items for the selected policy. Here you can:

- Quickly view and sort the subject, sender, and date of resolved items
- View the classification and comment history of any selected item

Reports

Use the supervision reports to see the review activity at the policy and reviewer level. For each policy, you can also view live statistics on the current state of review activity. You can use the supervision reports to:

- Verify that your policies are working as you intended.
- Find out how many communications need review.
- Find out how many communications aren't compliant and which ones are passing review. This information can help you decide whether to fine-tune your policies or change the number of reviewers.

View the Supervision report

1. Sign into the Compliance center with credentials for an admin account with permissions to view supervision reports.
2. Go to either **Reports** > **Dashboard** or **Supervision** to view the supervision reporting widget for a summary of current supervision policy activity.
3. Select the **Supervision** widget to open the detailed report page.

> **NOTE**
> If you aren't able to access the **Reports** page, check that you're a member of the Supervisory Review role group, as described in Make supervision available in your organization. Inclusion in this role group lets you create and manage supervision polices and run the report.

How to use the report

This report lists each policy and the number of communications at each stage in the review process. Use the report to:

- View data for all or specific policies.
- View data grouped by tag type, reviewer, or message type.
- Export data to a CSV file based on activity date, policy, and by reviewer activity.
- Filter data based on activity date, tag type, reviewer, and message type.

Here's a breakdown of the values displayed the **Tag type** column.

TAG TYPE	WHAT IT MEANS
Not Reviewed	The number of emails not reviewed yet. These emails are awaiting review in the Microsoft 365 supervision dashboard.
Compliant	The number of emails reviewed and marked as compliant. These messages still need resolution.

TAG TYPE	WHAT IT MEANS
Questionable	The number of emails reviewed and marked questionable. Serves as a flag for other reviewers to help check whether an email needs investigation for compliance. These messages still need resolution.
Non-Compliant (Active)	The number of non-compliant emails that reviewers are currently investigating.
Non-Compliant (Resolved)	The number of non-compliant emails that reviewers investigated and resolved.
Hit Policy	The total number (daily) of messages from Exchange, Teams, and third-party data sources that matched one or more conditions defined in a supervision policy
In Purview	The total number (daily) of messages from Exchange, Teams, and third-party data sources scanned by a supervision policy
Resolved	The total number of messages from Exchange, Teams, and third-party data sources classified as **Resolved**

> **NOTE**
>
> Supervision policies must be provisioned before they appear in reports. If policies are deleted, historical data is still shown. However, they're indicated as a "Non-existent policy" and the **Export** function isn't available.

Audit

In some instances, you must provide information to regulatory or compliance auditors to prove supervision of employee activities and communications. This information may be a summary of all supervisory activities associated with a defined policy or anytime a supervision policy changes. Supervision policies have built-in audit trails for complete readiness for internal or external audits. Detailed audit histories of every action monitored by your supervision policies provide proof of supervisory procedures.

View audit activities in the unified audit log or with the Search-UnifiedAuditLog PowerShell cmdlet.

For example, the following example returns the activities for the all the supervisory review activities (policies and rules) and lists detailed information for each:

```
Search-UnifiedAuditLog -StartDate 3/1/2019 -EndDate ([System.DateTime]::Now) -RecordType DataGovernance -
ResultSize 5000 | Where-Object {$_.Operations -like "*SupervisoryReview*"}  | fl
CreationDate,Operations,UserIds,AuditData
```

This example returns the update activities for your communication compliance policies:

```
Search-UnifiedAuditLog -StartDate $startDate -EndDate $endDate -Operations
SupervisionPolicyCreated,SupervisionPolicyUpdated,SupervisionPolicyDeletedAuditData
```

In addition to information provided in the supervision reports and logs, you can also use the Get-SupervisoryReviewActivity PowerShell cmdlet to return a complete detailed listing of all supervision policy activities.

Ready to get started?

To configure supervision policies for your organization, see Configure supervision policies.

Configure supervision policies in Office 365

11/2/2020 • 8 minutes to read • Edit Online

IMPORTANT

Following the release of communication compliance in Microsoft 365 Compliance in February 2020, Supervision in Office 365 is being retired. Supervision policies will no longer be available for creation, and policies will eventually be removed, after an extended period of read only access.

If you use Supervision, be aware that:

- Beginning June 15th, 2020, tenants will not have the ability to create new Supervision policies.
- Beginning August 31st, 2020, existing policies will stop capturing new messages.
- Beginning October 26th, 2020, existing policies will be deleted.

We actively encourage customers who are currently exploring or using Supervision in Office 365 to use the new communication compliance solution to address your communications monitoring or regulatory requirements with a much richer set of intelligent capabilities.

Use supervision policies to capture employee communications for examination by internal or external reviewers. For more information about how supervision policies can help you monitor communications in your organization, see Supervision policies in Office 365.

NOTE

Users monitored by supervision policies must have a Microsoft 365 E5 Compliance license, an Office 365 Enterprise E3 license with the Advanced Compliance add-on, or be included in an Office 365 Enterprise E5 subscription, or be included in a Microsoft 365 E5 subscription. If you don't have an existing Enterprise E5 plan and want to try supervision, you can sign up for a trial of Office 365 Enterprise E5.

Follow these steps to set up and use supervision in your organization:

- **Step 1 (optional)**: Set up groups for supervision

  Before you start using supervision policies, determine who needs communications reviewed and who performs reviews. If you want to get started with just a few users to see how supervision works, you can skip setting up groups for now.

- **Step 2 (required)**: Make supervision available in your organization

  Add yourself to the Supervisory Review role group so you can set up policies. Anyone who has this role assigned can access the **Supervision** page in the Security & Compliance Center. If reviewable email is hosted on Exchange Online, each reviewer must have remote PowerShell access to Exchange Online.

- **Step 3 (optional)**: Create custom sensitive information types and custom keyword dictionaries

  If you need a custom sensitive info type or a custom keyword dictionary for your supervision policy, you need to create it before starting the supervision wizard.

- **Step 4 (required)**: Set up a supervision policy

  You create supervision policies in the Security & Compliance Center. These policies define which communications are subject to review in your organization and specifies who performs reviews.

Communications include email and Microsoft Teams communications, and 3rd-party platform communications (such as Facebook, Twitter, etc.). Supervision policies created in organizations are not supported in communication supervision in Microsoft 365 subscriptions.

- **Step 5 (optional)**: Test your communication supervision policy

  Test your supervision policy to make sure it functions as desired. It is important to ensure that your compliance strategy is meeting your standards.

Step 1: Set up groups for supervision (optional)

When you create a supervision policy, you define who has their communications scanned and who performs reviews. In the policy, you'll use email addresses to identify individuals or groups of people. To simplify your setup, you can create groups for people who have their communication scanned and groups for people who review those communications. If you're using groups, you may need several. For example, you want to monitor communications between two distinct groups of people or if you want to specify a group that isn't going to be supervised.

Use the following chart to help you configure groups in your organization for communication supervision policies:

POLICY MEMBER	SUPPORTED GROUPS	UNSUPPORTED GROUPS
Supervised users Non-supervised users	Distribution groups Microsoft 365 groups	Dynamic distribution groups
Reviewers	Mail-enabled security groups	Distribution groups Dynamic distribution groups

When you select a Microsoft 365 group for supervised users, the policy monitors the content of the shared mailbox and the Microsoft Teams channels associated with the group. When you select a distribution list, the policy monitors individual user mailboxes.

To manage supervised users in large enterprise organizations, you may need to monitor all users across large groups. You can use PowerShell to configure a distribution group for a global supervision policy for the assigned group. This enables you to monitor thousands of users with a single policy and keep the supervision policy updated as new employees join your organization.

1. Create a dedicated distribution group for your global supervision policy with the following properties: Make sure that this distribution group isn't used for other purposes or other Office 365 services.

   - **MemberDepartRestriction = Closed**. Ensures that users cannot remove themselves from the distribution group.
   - **MemberJoinRestriction = Closed**. Ensures that users cannot add themselves to the distribution group.
   - **ModerationEnabled = True**. Ensures that all messages sent to this group are subject to approval and that the group is not being used to communicate outside of the supervision policy configuration.

   ```
   New-DistributionGroup -Name <your group name> -Alias <your group alias> -MemberDepartRestriction
   'Closed' -MemberJoinRestriction 'Closed' -ModerationEnabled $true
   ```

2. Select an unused Exchange custom attribute to track users added to the supervision policy in your organization.

3. Run the following PowerShell script on a recurring schedule to add users to the supervision policy:

```
$Mbx = (Get-Mailbox -RecipientTypeDetails UserMailbox -ResultSize Unlimited -Filter {CustomAttribute9
-eq $Null})
$i = 0
ForEach ($M in $Mbx)
{
  Write-Host "Adding" $M.DisplayName
  Add-DistributionGroupMember -Identity <your group name> -Member $M.DistinguishedName -ErrorAction
SilentlyContinue
  Set-Mailbox -Identity $M.Alias -<your custom attribute name> SRAdded
  $i++
}
Write-Host $i "Mailboxes added to supervisory review distribution group."
```

For more information about setting up groups, see:

- Create and manage distribution groups
- Manage mail-enabled security groups
- Overview of Microsoft 365 Groups

Step 2: Make supervision available in your organization (required)

To make **Supervision** available as a menu option in Security & Compliance Center, you must be assigned the Supervisory Review Administrator role.

To do this, you can either add yourself as a member of the Supervisory Review role group, or you can create a role group.

Add members to the Supervisory Review role group

1. Sign into https://protection.office.com using credentials for an admin account in your organization.
2. In the Security & Compliance Center, go to **Permissions**.
3. Select the **Supervisory Review** role group and then click the Edit icon.
4. In the **Members** section, add the people who you want to manage communication supervision for your organization.

Create a new role group

1. Sign into https://protection.office.com/permissions using credentials for an admin account in your organization.
2. In the Security & Compliance Center, go to **Permissions** and then click Add (+).
3. In the **Roles** section, click Add (+) and scroll down to **Supervisory Review Administrator**. Add this role to the role group.
4. In the **Members** section, add the people who you want to manage communication supervision for your organization.

For more information about role groups and permissions, see Permissions in the Compliance Center.

Enable remote PowerShell access for reviewers (if email is hosted on Exchange Online)

1. Follow the guidance in Enable or disable access to Exchange Online PowerShell.

Step 3: Create custom sensitive information types and custom keyword dictionaries (optional)

In order to pick from existing custom sensitive information types or custom keyword dictionaries in the

supervision policy wizard, you first need to create these items if needed.

Create custom keyword dictionary/lexicon (optional)

Use a text editor (like Notepad), to create a file that includes the keyword terms you'd like to monitor in a supervision policy. Make sure that each term is on a separate line and save the file in the **Unicode/UTF-16 (Little Endian)** format.

Create custom sensitive information types

1. Create a new sensitive information type and add your custom dictionary in the Security & Compliance Center. Navigate to **Classifications** > **Sensitive info types** and follow the steps in the **New sensitive info type wizard**. Here you will:

   - Define a name and description for the sensitive info type
   - Define the proximity, confidence level, and primary pattern elements
   - Import your custom dictionary as a requirement for the matching element
   - Review your selections and create the sensitive info type

   For more detailed information, see Create a custom sensitive information type and Create a keyword dictionary

   After the custom dictionary/lexicon is created, you can view the configured keywords with the Get-DlpKeywordDictionary cmdlet or add and remove terms using the Set-DlpKeywordDictionary cmdlet.

Step 4: Set up a supervision policy (required)

1. Sign into https://protection.office.com using credentials for an admin account in your organization.
2. In the Security & Compliance Center, select **Supervision**.
3. Select **Create** and follow the wizard to set up the policy configuration. Using the wizard, you will:

   - Give the policy a name and description.
   - Choose the users or groups to supervise, including choosing users or groups you'd like to exclude.
   - Define the supervision policy conditions. You can choose from message address, keyword, file types, and size match conditions.
   - Choose if you'd like to include sensitive information types. This is where you can select default and custom sensitive info types.
   - Choose if you'd like to enable the offensive language model. This detects inappropriate language sent or received in the body of email messages.
   - Define the percentage of communications to review.
   - Choose the reviewers for the policy. Reviewers can be individual users or mail-enabled security groups. All reviewers must have mailboxes hosted on Exchange Online.
   - Review your policy selections and create the policy.

Step 5: Test your supervision policy (optional)

After you create a communication supervision policy, it's a good idea to test to make sure that the conditions you defined are being properly enforced by the policy. You may also want to test your data loss prevention (DLP) policies if your supervision policies include sensitive information types. Follow these steps to test your supervision policy:

1. Open an email client or Microsoft Teams logged in as a supervised user defined in the policy you want to test.
2. Send an email or Microsoft Teams chat that meets the criteria you've defined in the supervision policy. This can be a keyword, attachment size, domain, etc. Make sure that you determine if your configured

conditional settings in the policy are too restrictive or too lenient.

> **NOTE**
>
> Emails subject to defined policies are processed in near real-time and can be tested immediately after the policy is configured. Chats in Microsoft Teams can take up to 24 hours to fully process in a policy.

3. Log into Microsoft 365 as a reviewer designated in the communication supervision policy. Navigate to **Supervision** > *Your Custom Policy* > **Open** to view the report for the policy.

Insider risk management in Microsoft 365

2/18/2021 • 2 minutes to read • Edit Online

Increasingly, employees have more access to create, manage, and share data across a broad spectrum of platforms and services. In most cases, organizations have limited resources and tools to identify and mitigate organization-wide risks while also meeting compliance requirements and employee privacy standards. These risks may include data theft by departing employees and data leaks of information outside your organization by accidental oversharing or malicious intent.

Insider risk management in Microsoft 365 uses the full breadth of service and 3rd-party indicators to help you quickly identify, triage, and act on risky user activity. By using logs from Microsoft 365 and Microsoft Graph, insider risk management allows you to define specific policies to identify risk indicators and to take action to mitigate these risks.

Configure insider risk management for Microsoft 365

Use the following steps to configure insider risk management for your organization:

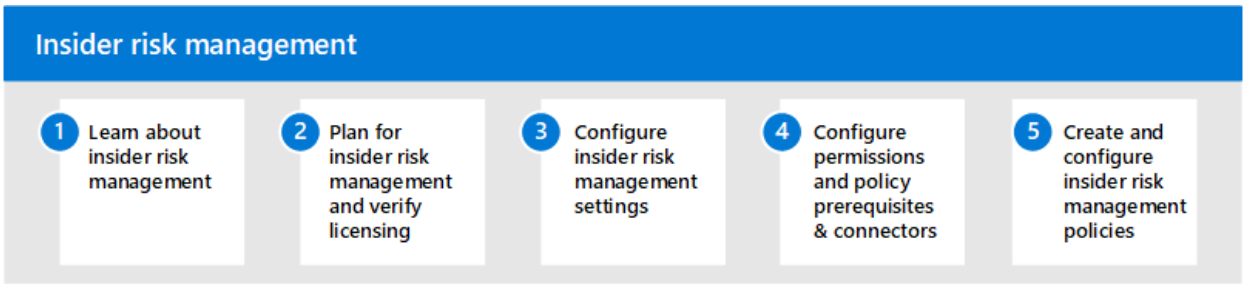

1. Learn about insider risk management in Microsoft 365
2. Plan for insider risk management and verify licensing
3. Configure insider risk management settings
4. Configure permissions and policy prerequisites & connectors
5. Create and configure insider risk management policies

More information about insider risk management

- Manage insider risk policies
- Investigate insider risk alerts
- Act on insider risk cases

Learn about insider risk management in Microsoft 365

2/18/2021 • 10 minutes to read • Edit Online

Insider risk management is a compliance solution in Microsoft 365 that helps minimize internal risks by enabling you to detect, investigate, and act on malicious and inadvertent activities in your organization. Insider risk policies allow you to define the types of risks to identify and detect in your organization, including acting on cases and escalating cases to Microsoft Advanced eDiscovery if needed. Risk analysts in your organization can quickly take appropriate actions to make sure users are compliant with your organization's compliance standards.

Watch the video below to learn how insider risk management can help your organization prevent, detect, and contain risks while prioritizing your organization values, culture, and user experience:

Modern risk pain points

Managing and minimizing risk in your organization starts with understanding the types of risks found in the modern workplace. Some risks are driven by external events and factors that are outside of direct control. Other risks are driven by internal events and user activities that can be minimized and avoided. Some examples are risks from illegal, inappropriate, unauthorized, or unethical behavior and actions by users in your organization. These behaviors include a broad range of internal risks from users:

- Leaks of sensitive data and data spillage
- Confidentiality violations
- Intellectual property (IP) theft
- Fraud
- Insider trading
- Regulatory compliance violations

Users in the modern workplace have access to create, manage, and share data across a broad spectrum of platforms and services. In most cases, organizations have limited resources and tools to identify and mitigate organization-wide risks while also meeting user privacy standards.

Insider risk management uses the full breadth of service and 3rd-party indicators to help you quickly identify, triage, and act on risk activity. By using logs from Microsoft 365 and Microsoft Graph, insider risk management allows you to define specific policies to identify risk indicators. These policies allow you to identify risky activities and to act to mitigate these risks.

Insider risk management is centered around the following principles:

- **Transparency**: Balance user privacy versus organization risk with privacy-by-design architecture.
- **Configurable**: Configurable policies based on industry, geographical, and business groups.
- **Integrated**: Integrated workflow across Microsoft 365 compliance solutions.
- **Actionable**: Provides insights to enable user notifications, data investigations, and user investigations.

Workflow

The insider risk management workflow helps you identify, investigate, and take action to address internal risks in your organization. With focused policy templates, comprehensive activity signaling across the Microsoft 365 service, and alert and case management tools, you can use actionable insights to quickly identify and act on risky behavior.

Identifying and resolving internal risk activities and compliance issues with insider risk management in Microsoft 365 uses the following workflow:

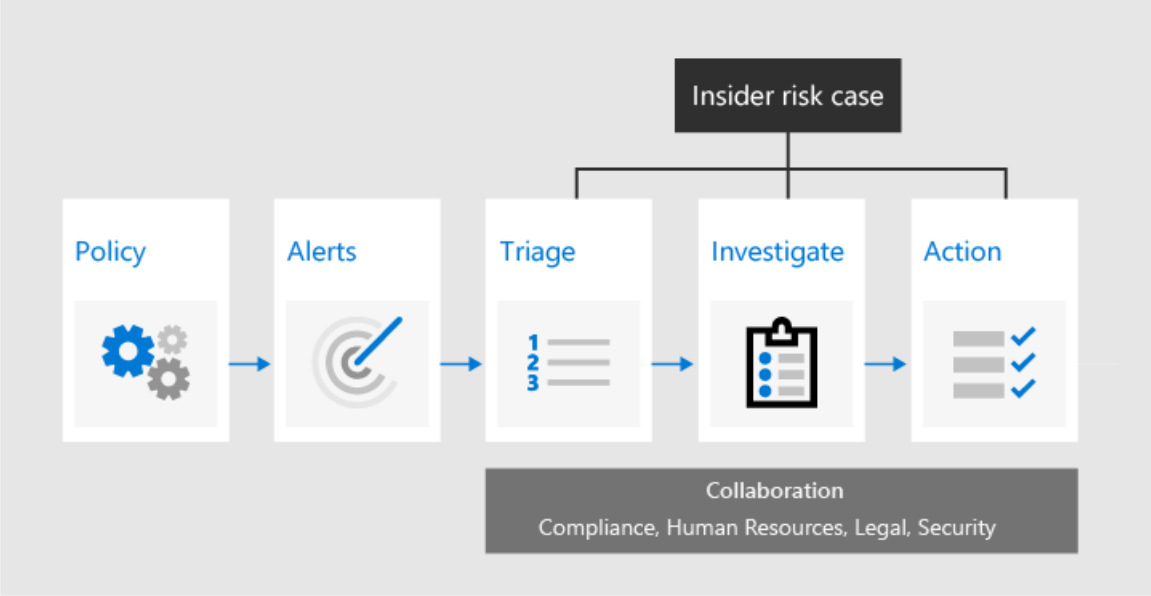

Policies

Insider risk management policies are created using pre-defined templates and policy conditions that define what triggering events and risk indicators are examined in your organization. These conditions include how risk indicators are used for alerts, what users are included in the policy, which services are prioritized, and the monitoring time period.

You can select from the following[policy templates to quickly get started with insider risk management:

- Data theft by departing users
- General data leaks
- Data leaks by priority users (preview)
- Data leaks by disgruntled users (preview)
- General security policy violations (preview)
- Security policy violations by departing users (preview)
- Security policy violations by priority users (preview)
- Security policy violations by disgruntled users (preview)

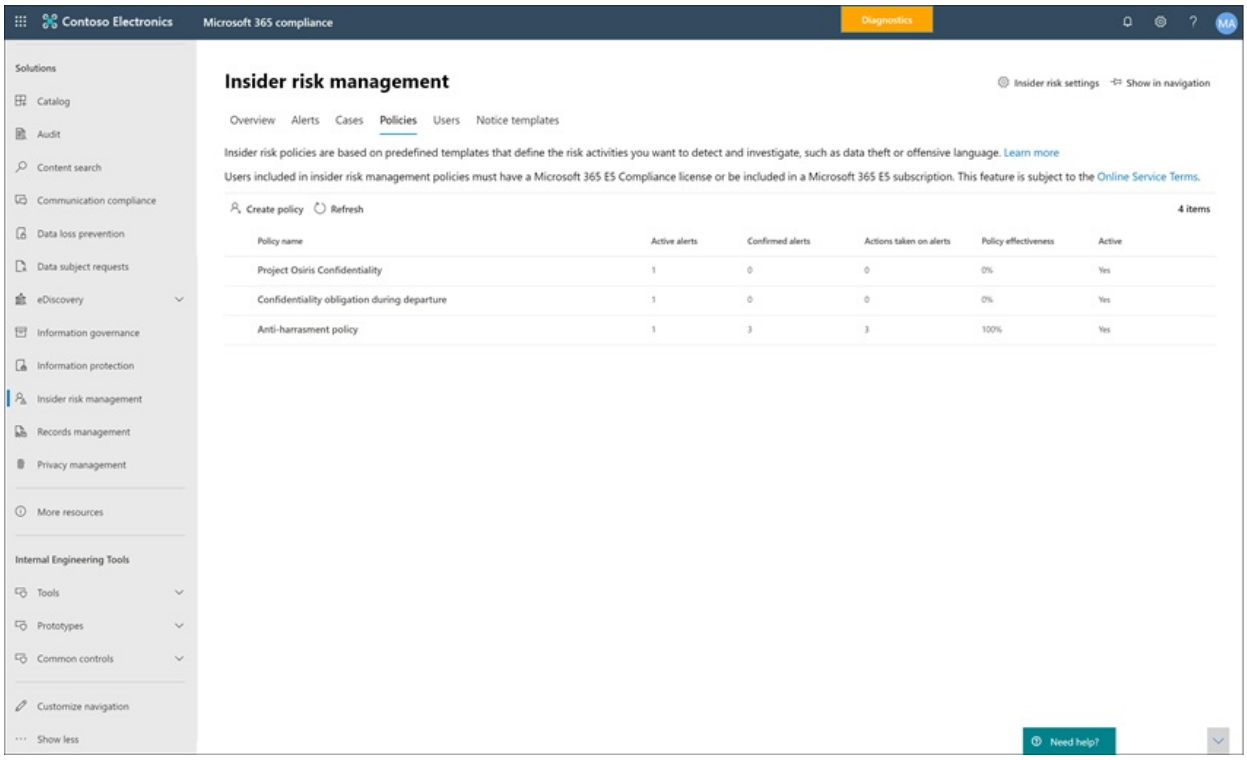

Alerts

Alerts are automatically generated by risk indicators that match policy conditions and are displayed in the Alerts dashboard. This dashboard enables a quick view of all alerts needing review, open alerts over time, and alert statistics for your organization. All policy alerts are displayed with the following information to help you quickly identify the status of existing alerts and new alerts that need action:

- Status
- Severity
- Time detected
- Case
- Case status

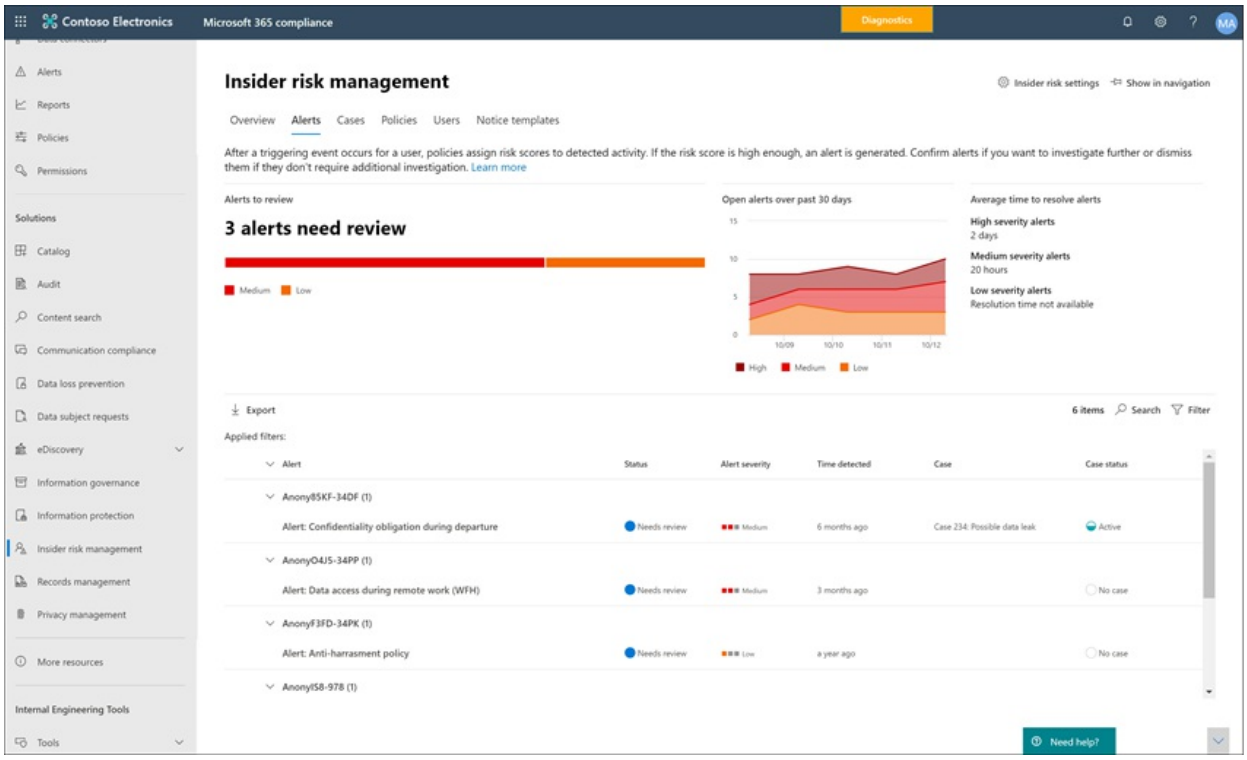

Triage

New user activities that need investigation automatically generate alerts that are assigned a *Needs review* status. Reviewers can quickly identify and review, evaluate, and triage these alerts.

Alerts are resolved by opening a new case, assigning the alert to an existing case, or dismissing the alert. Using alert filters, it's easy to quickly identify alerts by status, severity, or time detected. As part of the triage process, reviewers can view alert details for the activities identified by the policy, view user activity associated with the policy match, see the severity of the alert, and review user profile information.

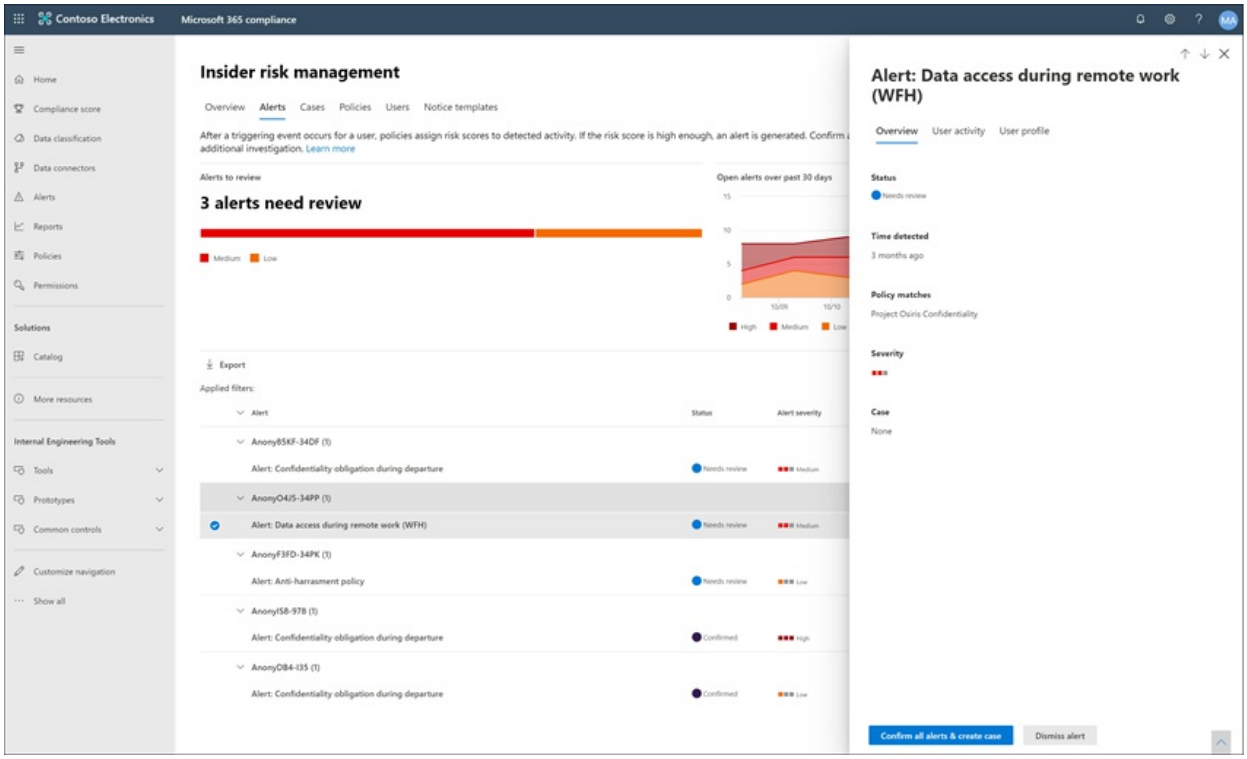

Investigate

Cases are created for alerts that require deeper review and investigation of the activity details and circumstances around the policy match. The **Case dashboard** provides an all-up view of all active cases, open cases over time, and case statistics for your organization. Reviewers can quickly filter cases by status, the date the case was opened, and the date the case was last updated.

Selecting a case on the case dashboard opens the case for investigation and review. This step is the heart of the insider risk management workflow. This area is where risk activities, policy conditions, alerts details, and user details are synthesized into an integrated view for reviewers. The primary investigation tools in this area are:

- **User activity**: User activity is automatically displayed in an interactive chart that plots activities over time and by risk level for current or past risk activities. Reviewers can quickly filter and view the entire risk history for the user and drill into specific activities for more details.
- **Content Explorer**: All data files and email messages associated with alert activities are automatically captured and displayed in the Content Explorer. Reviewers can filter and view files and messages by data source, file type, tags, conversation, and many more attributes.
- **Case notes**: Reviewers can provide notes for a case in the Case Notes section. This list consolidates all notes in a central view and include reviewer and date submitted information.

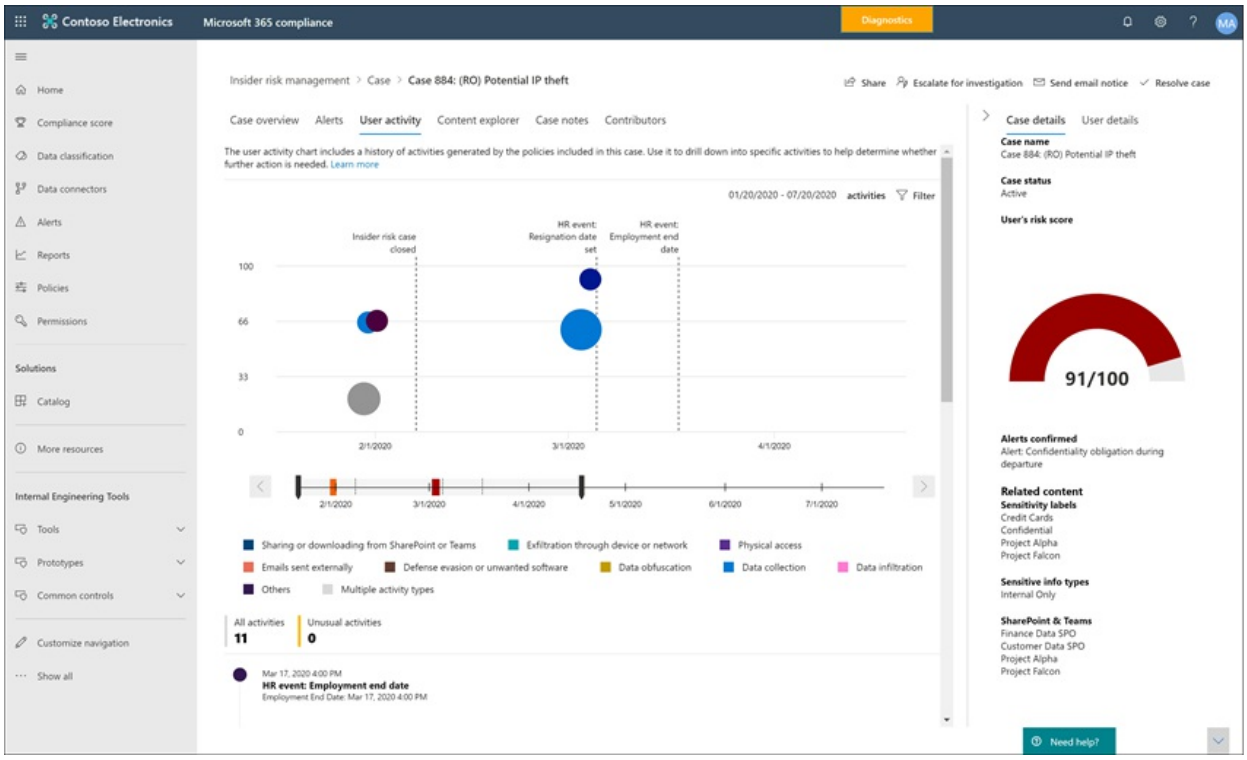

Action

After cases are investigated, reviewers can quickly act to resolve the case or collaborate with other risk stakeholders in your organization. If users accidentally or inadvertently violate policy conditions, a simple reminder notice can be sent to the user from notice templates you can customize for your organization. These notices may serve as simple reminders or may direct the user to refresher training or guidance to help prevent future risky behavior. For more information, see Insider risk management notice templates.

In the more serious situations, you may need to share the insider risk management case information with other reviewers or services in your organization. Insider risk management is tightly integrated with other Microsoft 365 compliance solutions to help you with end-to-end risk resolution.

- **Advanced eDiscovery**: Escalating a case for investigation allows you to transfer data and management of the case to Advanced eDiscovery in Microsoft 365. Advanced eDiscovery provides an end-to-end workflow to preserve, collect, review, analyze, and export content that's responsive to your organization's internal and external investigations. It allows legal teams to manage the entire legal hold notification workflow. To learn more about Advanced eDiscovery cases, see Overview of Advanced eDiscovery in Microsoft 365.
- **ServiceNow (preview)**: ServiceNow is a popular cloud computing platform that helps organizations manage digital workflows for enterprise operations. Insider risk management supports sharing case alerts with your ServiceNow service and allows you to create incidents and change requests related to individual insider risk cases. To learn more about sharing alert information with ServiceNow, see Share a case with ServiceNow.
- **Office 365 Management APIs integration (preview)**: Insider risk management supports exporting alert information to security information and event management (SIEM) services via the Office 365 Management APIs. Having access to alert information in the platform the best fits your organization's risk processes gives you more flexibility in how to act on risk activities. To learn more about exporting alert information with Office 365 Management APIs, see Export alerts.

> **NOTE**
>
> Thank you for your feedback and support during the preview of the ServiceNow connector. We've decided to end the preview of ServiceNow connector and discontinue support in insider risk management on November 30, 2020. We are actively evaluating alternative methods to provide customers with ServiceNow integration in insider risk management.

Scenarios

Insider risk management can help you detect, investigate, and take action to mitigate internal risks in your organization in several common scenarios:

Data theft by departing users

When users leave an organization, either voluntarily or as the result of termination, there is often legitimate concerns that company, customer, and user data are at risk. Users may innocently assume that project data isn't proprietary, or they may be tempted to take company data for personal gain and in violation of company policy and legal standards. Insider risk management policies that use the Data theft by departing users policy template automatically detect activities typically associated with this type of theft. With this policy, you'll automatically receive alerts for suspicious activities associated with data theft by departing users so you can take appropriate investigative actions. Configuring a Microsoft 365 HR connector for your organization is required for this policy template.

Intentional or unintentional leak of sensitive or confidential information

In most cases, users try their best to properly handle sensitive or confidential information. But occasionally users may make mistakes and information is accidentally shared outside your organization or in violation of your information protection policies. In other circumstances, users may intentionally leak or share sensitive and confidential information with malicious intent and for potential personal gain. Insider risk management policies created using the following Data leaks policy templates automatically detect activities typically associated with sharing sensitive or confidential information:

- General data leaks
- Data leaks by priority users (preview)
- Data leaks by disgruntled users (preview)

Intentional or unintentional security policy violations (preview)

Users typically have a large degree of control when managing their devices in the modern workplace. This may include permissions to install or uninstall applications needed in the performance of their duties or the ability to temporarily disable device security features. Whether this activity is inadvertent, accidental, or malicious, this conduct can pose risk to your organization and is important to identify and act to minimize. To help identity these risky security activities, the following insider risk management security policy violation templates scores security risk indicators and uses Microsoft Defender for Endpoint alerts to provide insights for security-related activities:

- General security policy violations (preview)
- Security policy violations by departing users (preview)
- Security policy violations by priority users (preview)
- Security policy violations by disgruntled users (preview)

Policies for users based on position, access level, or risk history (preview)

Users in your organization may have different levels of risk depending on their position, level of access to sensitive information, or risk history. This may include members of your organization's executive leadership team, IT administrators that have extensive data and network access privileges, or users with a past history of risky activities. In these circumstances, closer inspection and more aggressive risk scoring are important to help surface alerts for investigation and quick action. To help identify risky activities for these types of users, you can create priority user groups and create policies from the following policy templates:

- Security policy violations by priority users (preview)

- Data leaks by priority users (preview)

Actions and behaviors by disgruntled users (preview)

Employment stresses events can impact user behavior in several ways that relate to insider risks. These stressors may be a poor performance review, a position demotion, or the user being placement on a performance review plan. Though most users do not respond maliciously to these events, the stress of these actions may result in some users to take actions they may not normally consider during normal circumstances. To help identity these types risky activities, the following insider risk management policy templates use the Microsoft 365 HR connector and starts scoring risk indicators relating to behaviors that may occur near employment stressor events:

- Data leaks by disgruntled users (preview)
- Security policy violations by disgruntled users (preview)

Ready to get started?

- See Plan for insider risk management for how to prepare to enable insider risk management policies in your organization.
- See Get started with insider risk management settings to configure global settings for insider risk policies.
- See Get started with insider risk management to configure prerequisites, create policies, and start receiving alerts.

Plan for insider risk management

2/18/2021 • 6 minutes to read • Edit Online

Before getting started with insider risk management in your organization, there are important planning activities and considerations that should be reviewed by your information technology and compliance management teams. Thoroughly understanding and planning for deployment in the following areas will help ensure that your implementation and use of insider risk management features goes smoothly and is aligned with the best practices for the solution.

Work with stakeholders in your organization

Identify the appropriate stakeholders in your organization to collaborate for taking actions on insider risk management alerts and cases. Some recommended stakeholders to consider including in initial planning and the end-to-end insider risk management workflow are people from the following areas of your organization:

- Information technology
- Compliance
- Privacy
- Security
- Human resources
- Legal

Determine any regional compliance requirements

Different geographic and organizational areas may have compliance and privacy requirements that are different from other areas of your organization. Work with the stakeholders in these areas to ensure they understand the compliance and privacy controls in insider risk management and how they should be used across different areas of your organization. In some scenarios, compliance and privacy requirements might require policies that designate or restrict some stakeholders from investigations and cases based on the case for a user or regulatory or policy requirements for the area.

If you have requirements for specific stakeholders to be involved in case investigations that involve users in certain regions, roles, or divisions, you may want to implement separate (even if identical) insider risk management policies targeting the different regions and populations. This configuration will make it easier for the right stakeholders to triage and manage cases that are relevant to their roles and regions. Additionally, you may want to consider creating processes and policies for regions where investigators and reviewers speak the same language as the users to help streamline the escalation process for insider risk management alerts and cases.

Plan for the review and investigation workflow

Select dedicated stakeholders to monitor and review the alerts and cases on a regular cadence in the Microsoft 365 compliance center. Make sure understand how you will assign different stakeholders to the different role groups available in insider risk management.

Depending on the structure of your compliance management team, you have options to assign users to specific role groups to manage different sets of insider risk management features. To view the **Permissions** tab in the Office 365 Security & Compliance Center and manage role groups, you need to be assigned to the *Organization Management* role group or need to be assigned the *Role Management* role. Choose from these role group options when configuring insider risk management:

ROLE GROUP	ROLE PERMISSIONS
Insider Risk Management	Use this role group to manage insider risk management for your organization in a single group. By adding all user accounts for designated administrators, analysts, and investigators, you can configure insider risk management permissions in a single group. This role group contains all the insider risk management permission roles. This configuration is the easiest way to quickly get started with insider risk management and is a good fit for organizations that do not need separate permissions defined for separate groups of users.
Insider Risk Management Admin	Use this role group to initially configure insider risk management and later to segregate insider risk administrators into a defined group. Users in this role group can create, read, update, and delete insider risk management policies, and global settings.
Insider Risk Management Analysts	Use this group to assign permissions to users that will act as insider risk case analysts. Users in this role group can access to all insider risk management alerts, cases, and notices templates. They cannot access the insider risk Content Explorer.
Insider Risk Management Investigators	Use this group to assign permissions to users that will act as insider risk data investigators. Users in this role group can access to all insider risk management alerts, cases, notices templates, and the Content Explorer for all cases.

Understand requirements and dependencies

Depending on how you plan to implement insider risk management policies, you need to have the proper Microsoft 365 licensing subscriptions and understand and plan for some solution prerequisites.

Licensing: Insider risk management is available as part of wide selection of Microsoft 365 licensing subscriptions. For details, see the Getting started with insider risk management article.

If you don't have an existing Microsoft 365 Enterprise E5 plan and want to try insider risk management, you can add Microsoft 365 to your existing subscription or sign up for a trial of Microsoft 365 Enterprise E5.

Policy template requirements: Depending on the policy template you choose, there are requirements that you need to understand and plan for prior to configuring insider risk management in your organization:

- When using the **Data theft by departing users** template, you must configure a Microsoft 365 HR connector to periodically import resignation and termination date information for users in your organization. See the Import data with the HR connector article for step-by-step guidance to configure the Microsoft 365 HR connector for your organization.
- When using **Data leaks** templates, you must configure at least one Data Loss Prevention (DLP) policy to define sensitive information in your organization and to receive insider risk alerts for High Severity DLP policy alerts. See the Create, test, and tune a DLP policy article for step-by-step guidance to configure DLP policies for your organization.
- When using **Security policy violation** templates, you must enable Microsoft Defender for Endpoint for insider risk management integration in the Defender Security Center to import security violation alerts. See the Configure advanced features in Microsoft Defender article for step-by-step guidance to enable Defender for Endpoint integration with insider risk management.
- When using **Disgruntled user** templates, you must configure a Microsoft 365 HR connector to periodically

import performance or demotion status information for users in your organization. See the Import data with the HR connector article for step-by-step guidance to configure the Microsoft 365 HR connector for your organization.

Test with a small group of users in a production environment

Before enabling the solution broadly in your production environment, you may consider testing the policies with a small set of production users while conducting for the necessary compliance, privacy, and legal reviews in your organization. Evaluating insider risk management in a test environment would require that you generate simulated user actions and other signals to create alerts for triage and cases for processing. This approach isn't practical for most organizations, so testing insider risk management with a small group of users in a production environment is preferred.

Keep the anonymization feature in policy settings enabled to anonymize user display names in the insider risk management console during this testing to maintain privacy within the tool. This setting helps protect the privacy of users that have policy matches and can help promote objectivity in data investigation and analysis reviews for insider risk alerts.

If you don't see any alerts immediately after configuring an insider risk management policy, it may mean the minimum risk threshold has not been met yet. A good way to check if the policy is triggered and working as expected is to see if the user is in-scope for the policy on the **Users** page.

Resources for stakeholders

Share insider risk management documentation with the stakeholders in your organization that are included in your management and remediation workflow:

- Create and manage insider risk policies
- Investigate insider risk alerts
- Take action on insider risk cases
- Review case data with the insider risk Content Explorer
- Create insider risk notice templates

Ready to get started?

Ready to configure insider risk management for your organization? Review the following articles:

- Get started with insider risk management settings to configure global policy settings.
- Get started with insider risk management to configure prerequisites, create policies, and start receiving alerts.

Get started with insider risk management settings

2/18/2021 • 39 minutes to read • Edit Online

Insider risk management settings apply to all insider risk management policies, regardless of the template you chose when creating a policy. Settings are configured using the **Insider risk settings** control located at the top of all insider risk management tabs. These settings control policy components for the following areas:

- Privacy
- Indicators
- Policy timelines
- Intelligent detections
- Export alerts (preview)
- Priority user groups (preview)
- Priority physical assets (preview)
- Power Automate flows (preview)
- Microsoft Teams (preview)

Before you get started and create insider risk management policies, it's important to understand these settings and choose setting levels best for the compliance needs for your organization.

Privacy

Protecting the privacy of users that have policy matches is important and can help promote objectivity in data investigation and analysis reviews for insider risk alerts. For users with an insider risk policy match, you can choose one of the following settings:

- **Show anonymized versions of usernames**: Names of users are anonymized to prevent admins, data investigators, and reviewers from seeing who is associated with policy alerts. For example, a user 'Grace Taylor' would appear with a randomized pseudonym such as 'AnonIS8-988' in all areas of the insider risk management experience. Choosing this setting anonymizes all users with current and past policy matches and applies to all policies. User profile information in the insider risk alert and case details will not be available when this option is chosen. However, usernames are displayed when adding new users to existing policies or when assigning users to new policies. If you choose to turn off this setting, usernames will be displayed for all users that have current or past policy matches.
- **Do not show anonymized versions of usernames**: Usernames are displayed for all current and past policy matches for alerts and cases. User profile information (the name, title, alias, and organization or department) is displayed for the user for all insider risk management alerts and cases.

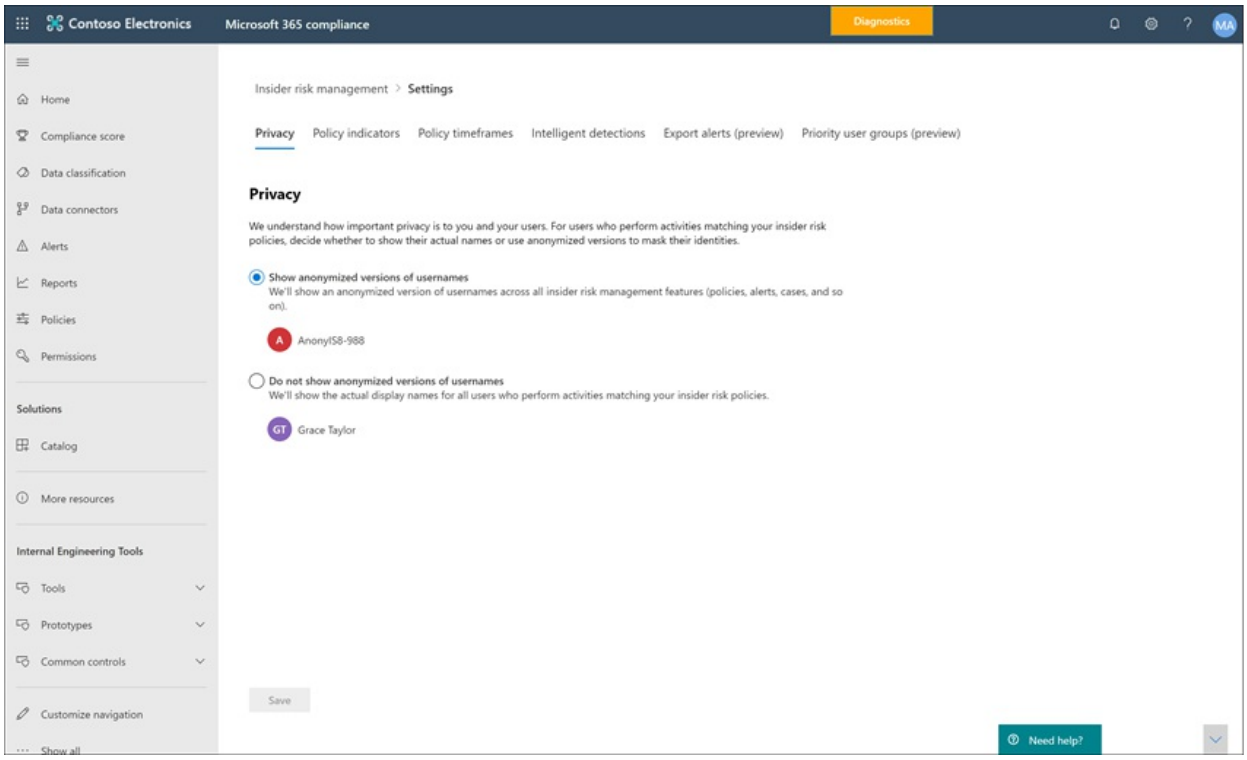

Indicators

Insider risk policy templates define the type of risk activities that you want to detect and investigate. Each policy template is based on specific indicators that correspond to specific triggers and risk activities. All indicators are disabled by default, and you must select one or more policy indicators before configuring an insider risk management policy.

Alerts are triggered by policies when users perform activities related to policy indicators that meet a required threshold. Insider risk management uses two types of indicators:

- **Triggering events**: Events that determine if a user is active for an insider risk management policy. If a user is added to an insider risk management policy does not have a triggering event, the user activity is not evaluated by the policy. For example, User A is added to a policy created from the *Data theft by departing users* policy template and the policy and Microsoft 365 HR connector are properly configured. Until User A has a termination date reported by the HR connector, User A activities aren't evaluated by this insider risk management policy for risk. Another example of a triggering event is if a user has a *High* severity DLP policy alert when using *Data leaks* policies.
- **Policy indicators**: Indicators included in insider risk management policies used to determine a risk score for an in-scope user. These policy indicators are only activated after a triggering event occurs for a user. Some examples of policy indicators are when a user copies data to personal cloud storage services or portable storage devices, or if a user shares internal files and folders with unauthorized external parties.

Policy indicators are segmented into the following areas. You can choose the indicators to activate and customize indicator event limits for each indicator level when creating an insider risk policy:

- **Office indicators**: These include policy indicators for SharePoint sites, Teams, and email messaging.
- **Device indicators**: These include policy indicators for activity such as sharing files over the network or with devices. Indicators include activity involving Microsoft Office files, .CSV files, and .PDF files. If you select **Device indicators**, activity is processed only for devices with Windows 10 Build 1809 or higher. For more information on configuring devices for integration with insider risk, see the following Enable device indicators and onboard devices section.
- **Security policy violation indicator**: These include indicators from Microsoft Defender for Endpoint related to unapproved or malicious software installation or bypassing security controls. To receive alerts in

insider risk management, you must have an active Defender for Endpoint license and insider risk integration enabled. For more information on configuring Defender for Endpoint for insider risk management integration, see Configure advanced features in Microsoft Defender for Endpoint.

- **Risk score boosters**: These include raising the risk score for unusual activities or past policy violations. Enabling risk score boosters increase risk scores and the likelihood of alerts for these types of activities. Risk score boosters can only be selected if one or more indicators are selected.

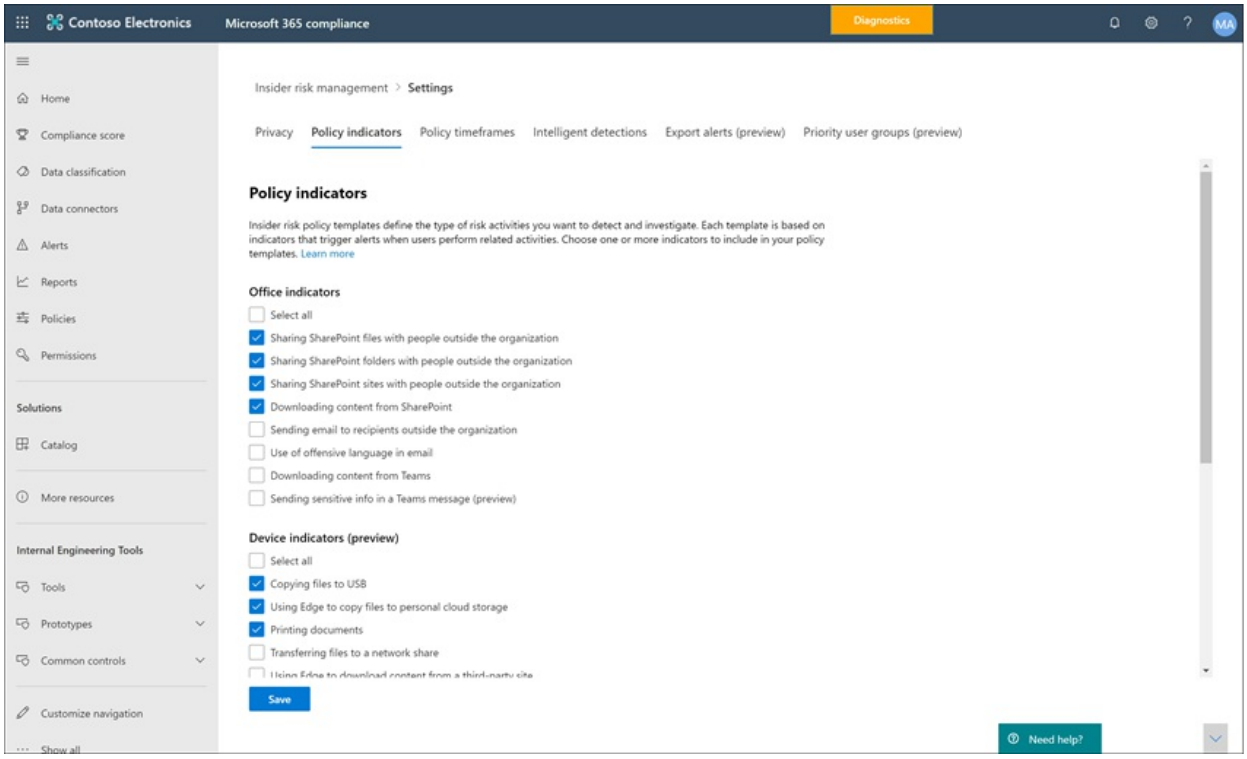

In some cases, you may want to limit the insider risk policy indicators that are applied to insider risk policies in your organization. You can turn off the policy indicators for specific areas by disabling them from all insider risk policies. Triggering events cannot be modified for insider risk policy templates.

To define the insider risk policy indicators that are enabled in all insider risk policies, navigate to **Insider risk settings** > **Indicators** and select one or more policy indicators. The indicators selected on the Indicators settings page cannot be individually configured when creating or editing an insider risk policy in the policy wizard.

> **NOTE**
>
> It may take several hours for new manually-added users to appear in the **Users dashboard**. Activities for the previous 90 days for these users may take up to 24 hours to display. To view activities for manually added users, select the user on the **Users dashboard** and open the **User activity** tab on the details pane.

Enable device indicators and onboard devices

To enable the monitoring of risk activities on devices and include policy indicators for these activities, your devices must meet the following requirements and you must complete the following onboarding steps.

Step 1: Prepare your endpoints

Make sure that the Windows 10 devices that you plan on reporting in insider risk management meet these requirements.

1. Must be running Windows 10 x64 build 1809 or later and must have installed the Windows 10 update (OS Build 17763.1075) from February 20, 2020.
2. All devices must be Azure Active Directory (AAD) joined, or Hybrid Azure AD joined.
3. Install Microsoft Chromium Edge browser on the endpoint device to monitor actions for the cloud upload

activity. See, Download the new Microsoft Edge based on Chromium.

Step 2: Onboarding devices

You must enable device monitoring and onboard your endpoints before you can monitor for insider risk management activities on a device. Both of these actions are done in the Microsoft 365 Compliance portal.

When you want to onboard devices that haven't been onboarded yet, you'll download the appropriate script and deploy as outlined in the following steps.

If you already have devices onboarded into Microsoft Defender for Endpoint, they will already appear in the managed devices list. Follow Step 3: If you have devices onboarded into Microsoft Defender for Endpoint in the next section.

In this deployment scenario, you'll onboard devices that have not been onboarded yet, and you just want to monitor insider risk activities on Windows 10 devices.

1. Open the Microsoft compliance center.
2. Open the Compliance Center settings page and choose **Onboard devices**.

   > **NOTE**
   > While it usually takes about 60 seconds for device onboarding to be enabled, please allow up to 30 minutes before engaging with Microsoft support.

3. Choose **Device management** to open the **Devices** list. The list will be empty until you onboard devices.
4. Choose **Onboarding** to begin the onboarding process.
5. Choose the way you want to deploy to these more devices from the **Deployment method** list and then **download package**.
6. Follow the appropriate procedures in Onboarding tools and methods for Windows 10 machines. This link takes you to a landing page where you can access Microsoft Defender for Endpoint procedures that match the deployment package you selected in step 5:

   - Onboard Windows 10 machines using Group Policy
   - Onboard Windows machines using Microsoft Endpoint Configuration Manager
   - Onboard Windows 10 machines using Mobile Device Management tools
   - Onboard Windows 10 machines using a local script
   - Onboard non-persistent virtual desktop infrastructure (VDI) machines.

Once done and endpoint is onboarded, it should be visible in the devices list and the endpoint will start reporting audit activity logs to insider risk management.

> **NOTE**
> This experience is under license enforcement. Without the required license, data will not be visible or accessible.

Step 3: If you have devices onboarded into Microsoft Defender for Endpoint

If Microsoft Defender for Endpoint is already deployed and there are endpoints reporting in, all these endpoints will appear in the managed devices list. You can continue to onboard new devices into insider risk management to expand coverage by using the Step 2: Onboarding devices section.

1. Open the Microsoft compliance center.
2. Open the Compliance Center settings page and choose **Enable device monitoring**.

3. Choose **Device management** to open the **Devices** list. You should see the list of devices that are already reporting in to Microsoft Defender for Endpoint.
4. Choose **Onboarding** if you need to onboard more devices.
5. Choose the way you want to deploy to these more devices from the **Deployment method** list and then **Download package**.
6. Follow the appropriate procedures in Onboarding tools and methods for Windows 10 machines. This link takes you to a landing page where you can access Microsoft Defender for Endpoint procedures that match the deployment package you selected in step 5:
   - Onboard Windows 10 machines using Group Policy
   - Onboard Windows machines using Microsoft Endpoint Configuration Manager
   - Onboard Windows 10 machines using Mobile Device Management tools
   - Onboard Windows 10 machines using a local script
   - Onboard non-persistent virtual desktop infrastructure (VDI) machines.

Once done and endpoint is onboarded, it should be visible under the **Devices** table and the endpoint will start reporting audit activity logs to insider risk management.

> **NOTE**
> This experience is under license enforcement. Without the required license, data will not be visible or accessible.

Indicator level settings (preview)

When creating a policy in the policy wizard, you can configure how the daily number of risk events should influence the risk score for insider risk alerts. These indicator settings help you control how the number of occurrences of risk events in your organization should affect the risk score, and consequently the associated alert severity, for these events. If you prefer, you can also choose to keep the default event threshold levels recommended by Microsoft for all enabled indicators.

For example, you decide to enable SharePoint indicators in the insider risk policy settings and to set custom thresholds for SharePoint events when configuring indicators for a new insider risk *Data leaks* policy. While in the insider risk policy wizard, you configure three different daily event levels for each SharePoint indicator to influence the risk score for alerts associated with these events.

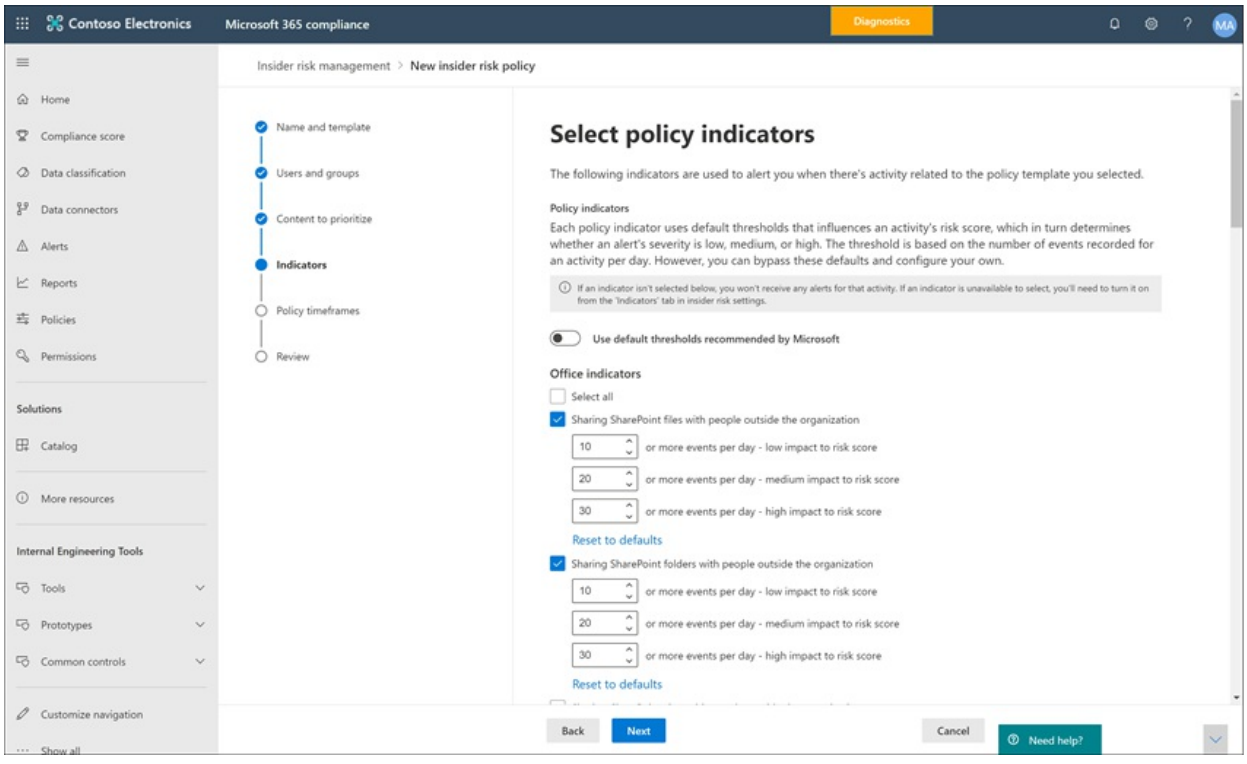

For the first daily event level, you set the threshold at *10 or more events per day* for a lower impact to the risk score for the events, *20 or more events per day* for a medium impact to the risk score for the events, and *30 or more events per day* a higher impact to the risk score for the events. These settings effectively mean:

- If there are 1-9 SharePoint events that take place after triggering event, risk scores are minimally impacted and would tend not to generate an alert.
- If there are 10-19 SharePoint events that take place after a triggering event, the risk score is inherently lower and alert severity levels would tend to be at a low level.
- If there are 20-29 SharePoint events that take place after a triggering, the risk score is inherently higher and alert severity levels would tend to be at a medium level.
- If there are 30 or more SharePoint events that take place after a triggering, the risk score is inherently higher and alert severity levels would tend to be at a high level.

Policy timeframes

Policy timeframes allow you to define past and future review periods that are triggered after policy matches based on events and activities for the insider risk management policy templates. Depending on the policy template you choose, the following policy timeframes are available:

- **Activation window**: Available for all policy templates, the *Activation window* is the defined number of days that the window activates **after** a triggering event. The window activates for 1 to 30 days after a triggering event occurs for any user assigned to the policy. For example, you've configured an insider risk management policy and set the *Activation window* to 30 days. Several months have passed since you configured the policy and a triggering event occurs for one of the users included in the policy. The triggering event activates the *Activation window* and the policy is active for that user for 30 days after the triggering event occurred.
- **Past activity detection**: Available for all policy templates, the *Past activity detection* is the defined number of days that the window activates **before** a triggering event. The window activates for 0 to 180 days before a triggering event occurs for any user assigned to the policy. For example, you've configured an insider risk management policy and set the *Past activity detection* to 90 days. Several months have passed since you configured the policy and a triggering event occurs for one of the users included in the policy. The triggering event activates the *Past activity detection* and the policy gathers historic activities for that user for 90 days prior to the triggering event.

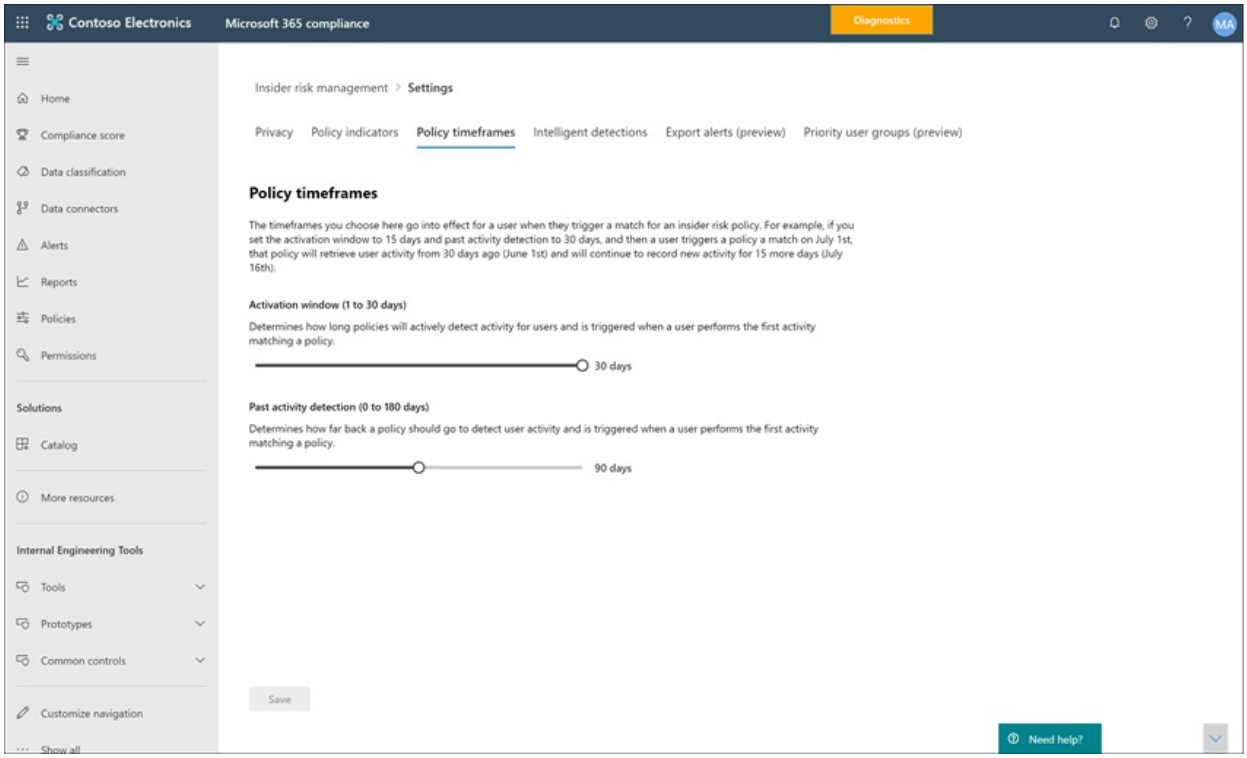

Intelligent detections

Intelligent detection settings help refine how the detections of risky activities are processed for alerts. In certain circumstances, you may need to define files types to ignore or you want to enforce a detection level for files to help define a minimum bar for alerts. Use these settings to control overall alert volume, file type exclusions, and file volume limits.

Anomaly detections

Anomalous detections include settings for file type exclusions and file volume limits.

- **File type exclusions**: To exclude specific file types from all insider risk management policy matching, enter file type extensions separated by commas. For example, to exclude certain types of music files from policy matches you may enter *aac,mp3,wav,wma* in the **File type exclusions** field. Files with these extensions would be ignored by all insider risk management policies.
- **File volume cut-off limit**: To define a minimum file level before activity alerts are reported in insider risk policies, enter the number of files. For example, you would enter '10' if you do not want to generate insider risk alerts when a user downloads 10 files or less, even if the policies consider this activity an anomaly.

Alert volume

User activities detected by insider risk policies are assigned a specific risk score, which in turn determines the alert severity (low, medium, high). By default, we'll generate a certain amount of low, medium, and high severity alerts, but you can increase or decrease the volume to suit your needs. To adjust the volume of alerts for all insider risk management policies, choose one of the following settings:

- **Fewer alerts**: You'll see all high severity alerts, fewer medium severity alerts, and no low severity ones. This setting level means you might miss some true positives.
- **Default volume**: You'll see all high severity alerts and a balanced amount of medium and low severity alerts.
- **More alerts**: You'll see all medium and high severity alerts and most low severity alerts. This setting level might result in more false positives.

Microsoft Defender for Endpoint (preview)

Microsoft Defender for Endpoint is an enterprise endpoint security platform designed to help enterprise networks prevent, detect, investigate, and respond to advanced threats. To have better visibility of security violation in your organization, you can import and filter Defender for Endpoint alerts for activities used in policies created from insider risk management security violation policy templates.

Depending on the types of signals you are interested in, you can choose to import alerts to insider risk management based on the Defender for Endpoint alert triage status. You can define one or more of the following alert triage statuses in the global settings to import:

- Unknown
- New
- In progress
- Resolved

Alerts from Defender for Endpoint are imported daily. Depending on the triage status you choose, you may see multiple user activities for the same alert as the triage status changes in Defender for Endpoint.

For example, if you select *New*, *In progress*, and *Resolved* for this setting, when a Microsoft Defender for Endpoint alert is generated and the status is *New*, an initial alert activity is imported for the user in insider risk. When the Defender for Endpoint triage status changes to *In progress*, a second activity for this alert is imported for the user in insider risk. When the final Defender for Endpoint triage status of *Resolved* is set, a third activity for this alert is imported for the user in insider risk. This functionality allows investigators to follow the progression of the Defender for Endpoint alerts and choose the level of visibility that their investigation

requires.

> **IMPORTANT**
> You'll need to have Microsoft Defender for Endpoint configured in your organization and enable Defender for Endpoint for insider risk management integration in the Defender Security Center to import security violation alerts. For more information on configuring Defender for Endpoint for insider risk management integration, see Configure advanced features in Defender for Endpoint.

Domains (preview)

Domain settings help you define risk levels for communications to specific domains. These communications include sharing files, email messages, or downloading content. By specifying domains in these settings, you can increase or decrease the risk scoring for activity that takes place with these domains. For example, to specify contoso.com and sales.wingtiptoys.com as allowed domains, you will enter 'contoso.com sales.wingtiptoys.com' in the **Allowed domains** field.

For each of the following domain settings, you can enter up to 500 domains:

- **Unallowed domains:** By specifying unallowed domains, activity that takes place with these domains will have *higher* risk scores.
- **Allowed domains:** By specifying allowed domains in settings, activity that takes place with these domains will have *lower* risk scores and is treated similarly to how internal organization activity is treated. For example, email activities to these domains are analyzed similarly to how internal email activity is analyzed.
- **Third party domains:** Third party domains are domains used for business purposes at your organization and sensitive content may be stored across these locations. By specifying a third party domain, you can receive alerts for any risky activity on these domains.

Export alerts (preview)

Insider risk management alert information is exportable to security information and event management (SIEM) services via the Office 365 Management Activity API schema. You can use the Office 365 Management Activity APIs to export alert information to other applications your organization may use to manage or aggregate insider risk information.

To use the APIs to review insider risk alert information:

1. Enable Office 365 Management Activity API support in **Insider risk management** > **Settings** > **Export**. By default, this setting is disabled for your Microsoft 365 organization.
2. Filter the common Office 365 audit activities by *SecurityComplianceAlerts*.
3. Filter *SecurityComplianceAlerts* by the *InsiderRiskManagement* category.

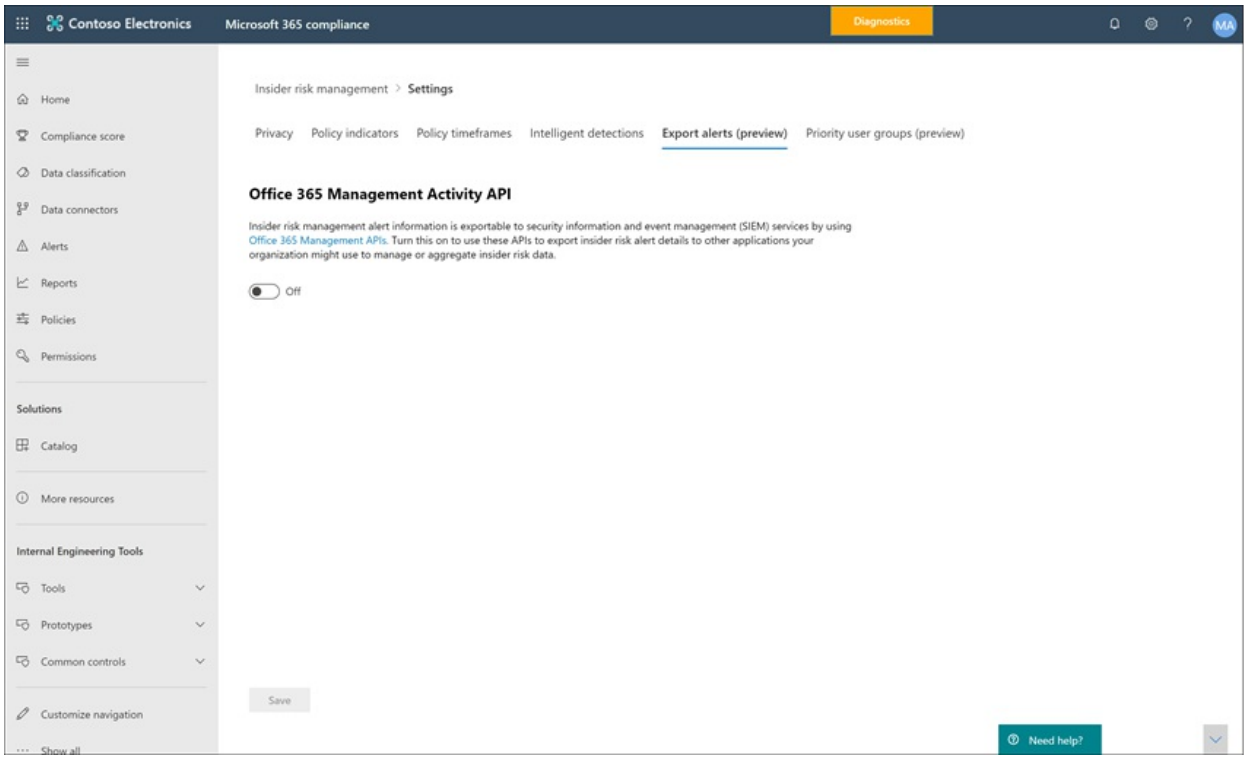

Alert information contains information from the security and compliance alert schema and the Office 365 Management Activity API common schema.

The following fields and values are exported for insider risk management alerts for the Security & Compliance alert schema:

ALERT PARAMETER	DESCRIPTION
AlertType	Type of the alert is *Custom*.
AlertId	The GUID of the alert. Insider risk management alerts are mutable. As alert status changes, a new log with the same AlertID is generated. This AlertID can be used to correlate updates for an alert.
Category	The category of the alert is *InsiderRiskManagement*. This category can be used to distinguish from these alerts from other Security & Compliance alerts.
Comments	Default comments for the alert. Values are *New Alert* (logged when an alert is created) and *Alert Updated* (logged when there is an update to an alert). Use the AlertID to correlate updates for an alert.
Data	The data for the alert, includes the unique user ID, user principal name, and date and time (UTC) when user was triggered into a policy.
Name	Policy name for insider risk management policy that generated the alert.
PolicyId	The GUID of the insider risk management policy that triggered the alert.
Severity	The severity of the alert. Values are *High*, *Medium*, or *Low*.

ALERT PARAMETER	DESCRIPTION
Source	The source of the alert. The value is *Office 365 Security & Compliance*.
Status	The status of the alert. Values are *Active* (*Needs Review* in insider risk), *Investigating* (*Confirmed* in insider risk), *Resolved* (*Resolved* in insider risk), *Dismissed* (*Dismissed* in insider risk).
Version	The version of the security and compliance alert schema.

The following fields and values are exported for insider risk management alerts for the Office 365 Management Activity API common schema.

- UserId
- Id
- RecordType
- CreationTime
- Operation
- OrganizationId
- UserType
- UserKey

Priority user groups (preview)

Users in your organization may have different levels of risk depending on their position, level of access to sensitive information, or risk history. Prioritizing the examination and scoring of the activities of these users can help alert you to potential risks that may have higher consequences for your organization. Priority user groups in insider risk management help define the users in your organization that need closer inspection and more sensitive risk scoring. Coupled with the *Security policy violations by priority users* and *Data leaks by priority users* policy templates, users added to a priority user group have an increased likelihood of insider risk alerts and alerts with higher severity levels.

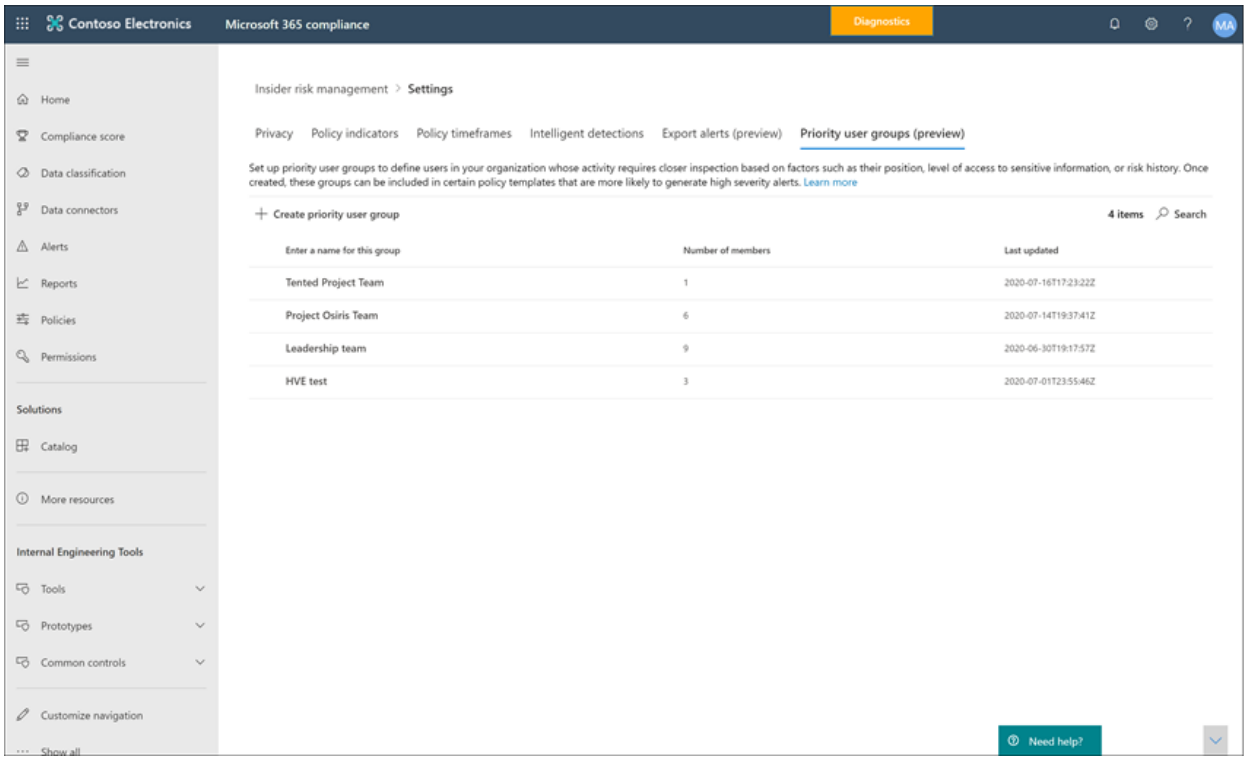

For example, you need to protect against data leaks for a highly confidential project where users have access to sensitive information. You choose to create *Confidential Project Users* priority user group for users in your organization that work on this project. Using the policy wizard and the *Data leaks by priority users* policy template, you create a new policy and assign the *Confidential Project Users* priority users group to the policy. Activities examined by the policy for members of the *Confidential Project Users* priority user group are more sensitive to risk and activities by these users will be more likely to generate an alert and have alerts with higher severity levels.

Create a priority user group

To create a new priority user group, you'll use setting controls in the **Insider risk management** solution in the Microsoft 365 compliance center. To create a priority user group, you must be a member of the *Insider Risk Management* or *Insider Risk Management Admin* role group.

Complete the following steps to create a priority user group:

1. In the Microsoft 365 compliance center, go to **Insider risk management** and select **Insider risk settings**.
2. Select the **Priority user groups** tab
3. On the **Priority user groups** tab, select **Create priority user group** to start the group creation wizard.
4. On the **Define group** page, complete the following fields:
   - **Name (required)**: Enter a friendly name for the priority user group. You can't change the name of the priority user group after you complete the wizard.
   - **Description (optional)**: Enter a description for the priority user group.
5. Select **Next** to continue.
6. On the **Choose members** page, select **Choose members** to search and select which mail-enabled user accounts are included in the group or select the **Select all** checkbox to add all users in your organization to the group. Select **Add** to continue or **Cancel** to close without adding any users to the group.
7. Select **Next** to continue.
8. On the **Review** page, review the settings you've chosen for the priority user group. Select **Edit** to change any of the group values or select **Submit** to create and activate the priority user group.
9. On the confirmation page, select **Done** to exit the wizard.

Update a priority user group

To update an existing priority user group, you'll use setting controls in the **Insider risk management** solution in the Microsoft 365 compliance center. To update a priority user group, you must be a member of the *Insider Risk Management* or *Insider Risk Management Admin* role group.

Complete the following steps to edit a priority user group:

1. In the Microsoft 365 compliance center, go to **Insider risk management** and select **Insider risk settings**.
2. Select the **Priority user groups** tab
3. Select the priority user group you want to edit and select **Edit group**.
4. On the **Define group** page, update the Description field if needed. You can't update the name of the priority user group. Select **Next** to continue.
5. On the **Choose members** page, add new members to the group using the **Choose members** control. To remove a user from the group, select the 'X' next to the user you wish to remove. Select **Next** to continue.
6. On the **Review** page, review the update settings you've chosen for the priority user group. Select **Edit** to change any of the group values or select **Submit** to update the priority user group.
7. On the confirmation page, select **Done** to exit the wizard.

Delete a priority user group

To delete an existing priority user group, you'll use setting controls in the **Insider risk management** solution in the Microsoft 365 compliance center. To delete a priority user group, you must be a member of the *Insider*

Risk Management or *Insider Risk Management Admin* role group.

> **IMPORTANT**
> Deleting a priority user group will remove it from any active policy to which it is assigned. If you delete a priority user group that is assigned to an active policy, the policy will not contain any in-scope users and will effectively be idle and will not create alerts.

Complete the following steps to delete a priority user group:

1. In the Microsoft 365 compliance center, go to **Insider risk management** and select **Insider risk settings**.
2. Select the **Priority user groups** tab
3. Select the priority user group you want to edit and select **Delete** from the dashboard menu.
4. On the **Delete** dialog, select **Yes** to delete the priority user group or select **Cancel** to return to the dashboard.

Priority physical assets (preview)

Identifying access to priority physical assets and correlating access activity to user events is an important component of your compliance infrastructure. These physical assets represent priority locations in your organization, such as company buildings, data centers, or server rooms. Insider risk activities may be associated with users working unusual hours, attempting to access these unauthorized sensitive or secure areas, and requests for access to high-level areas without legitimate needs.

With priority physical assets enabled and the Physical badging data connector configured, insider risk management integrates signals from your physical control and access systems with other user risk activities. By examining patterns of behavior across physical access systems and correlating these activities with other insider risk events, insider risk management can help compliance investigators and analysts make more informed response decisions for alerts. Access to priority physical assets are scored and identified in insights differently from access to non-priority assets.

For example, your organization has a badging system for users that monitors and approves physical access to normal working and sensitive project areas. You have several users working on a sensitive project and these users will return to other areas of your organization when the project is completed. As the sensitive project nears completion, you want to make sure that the project work remains confidential and that access to the project areas is tightly controlled.

You choose to enable the Physical badging data connector in Microsoft 365 to import access information from your physical badging system and specify priority physical assets in insider risk management. By importing information from your badging system and correlating physical access information with other risk activities identified in insider risk management, you notice that one of the users on the project is accessing the project offices after normal working hours and is also exporting large amounts of data to a personal cloud storage service from their normal work area. This physical access activity associated with the online activity may point to possible data theft and compliance investigators and analysts can take appropriate actions as dictated by the circumstances for this user.

Configure priority physical assets

To configure priority physical assets, you'll configure the Physical badging connector and use setting controls in the **Insider risk management** solution in the Microsoft 365 compliance center. To configure priority physical assets, you must be a member of the *Insider Risk Management* or *Insider Risk Management Admin role group*.

Complete the following steps to configure priority physical assets:

1. Follow the configuration steps for insider risk management in the Getting started with insider risk management article. In Step 3, make sure you configure the Physical badging connector.

> **IMPORTANT**
> For insider risk management policies to use and correlate signal data related to departing and terminated users with event data from your physical control and access platforms, you must also configure the Microsoft 365 HR connector. If you enable the Physical badging connector without enabling the Microsoft 365 HR connector, insider risk management policies will only process events for physical access activities for users in your organization.

2. In the Microsoft 365 compliance center, go to **Insider risk management** and select **Insider risk settings** > **Priority physical assets**.
3. On the **Priority physical assets** page, you can either manually add the physical asset IDs you want to monitor for the asset events imported by the Physical badging connector or import a .CSV file of all physical assets IDs imported by the Physical badging connector: a) To manually add physical assets IDs, choose **Add priority physical assets**, enter a physical asset ID, then select **Add**. Enter other physical asset IDs and then select **Add priority physical assets** to save all the assets entered. b) To add a list of physical asset IDs from a .CSV file, choose **Import priority physical assets**. From the file explorer dialog, select the .CSV file you wish to import, then select **Open**. The physical asset IDs from the .CSV files are added to the list.
4. Navigate to the **Policy indicators** tab in Settings.
5. On the **Policy indicators** page, navigate to the **Physical access indicators** section and select the checkbox for **Physical access after termination or failed access to sensitive asset**.
6. Select **Save** to configure and exit.

Delete a priority physical asset

To delete an existing priority physical asset, you'll use setting controls in the Insider risk management solution in the Microsoft 365 compliance center. To delete a priority physical asset, you must be a member of the Insider Risk Management or Insider Risk Management Admin role group.

> **IMPORTANT**
> Deleting a priority physical asset removes it from examination by any active policy to which it was previously included. Alerts generated by activities associated with the priority physical asset aren't deleted.

Complete the following steps to delete a priority physical asset:

1. In the Microsoft 365 compliance center, go to **Insider risk management** and select **Insider risk settings** > **Priority physical assets**.
2. On the **Priority physical assets** page, select the asset you want to delete.
3. Select **Delete** on the action menu to delete the asset.

Power Automate flows (preview)

Microsoft Power Automate is a workflow service that automates actions across applications and services. By using flows from templates or created manually, you can automate common tasks associated with these applications and services. When you enable Power Automate flows for insider risk management, you can automate important tasks for cases and users. You can configure Power Automate flows to retrieve user, alert, and case information and share this information with stakeholders and other applications, as well as automate actions in insider risk management, such as posting to case notes. Power Automate flows are applicable for cases and any user in scope for a policy.

Customers with Microsoft 365 subscriptions that include insider risk management do not need additional Power

Automate licenses to use the recommended insider risk management Power Automate templates. These templates can be customized to support your organization and cover core insider risk management scenarios. If you choose to use premium Power Automate features in these templates, create a custom template using the Microsoft 365 compliance connector, or use Power Automate templates for other compliance areas in Microsoft 365, you may need more Power Automate licenses.

The following Power Automate templates are provided to customers to support process automation for insider risk management users and cases:

- **Notify users when they're added to an insider risk policy**: This template is for organizations that have internal policies, privacy, or regulatory requirements that users must be notified when they are subject to insider risk management policies. When this flow is configured and selected for a user in the users page, users and their managers are sent an email message when the user is added to an insider risk management policy. This template also supports updating a SharePoint list hosted on a SharePoint site to help track notification message details like date/time and the message recipient. If you've chosen to anonymize users in **Privacy settings**, flows created from this template will not function as intended so that user privacy is maintained. Power Automate flows using this template are available on the **Users dashboard**.

- **Request information from HR or business about a user in an insider risk case**: When acting on a case, insider risk analysts and investigators may need to consult with HR or other stakeholders to understand the context of the case activities. When this flow is configured and selected for a case, analysts and investigators send an email message to HR and business stakeholders configured for this flow. Each recipient is sent a message with pre-configured or customizable response options. When recipients select a response option, the response is recorded as a case note and includes recipient and date/time information. If you've chosen to anonymize users in **Privacy settings**, flows created from this template will not function as intended so that user privacy is maintained. Power Automate flows using this template are available on the **Cases dashboard**.

- **Notify manager when a user has an insider risk alert**: Some organizations may need to have immediate management notification when a user has an insider risk management alert. When this flow is configured and selected, the manager for the case user is sent an email message with the following information about all case alerts:

  - Applicable policy for the alert
  - Date/Time of the alert
  - Severity level of the alert

  The flow automatically updates the case notes that the message was sent and that the flow was activated. If you've chosen to anonymize users in **Privacy settings**, flows created from this template will not function as intended so that user privacy is maintained. Power Automate flows using this template are available on the **Cases dashboard**.

- **Add calendar reminder to follow up on an insider risk case**: This template allows risk investigators and analysts to add calendar reminders for cases to their Office 365 Outlook calendar. This flow eliminates the need for users to exit or switch out of the insider risk management workflow when processing cases and triaging alerts. When this flow is configured and selected, a reminder is added to Office 365 Outlook calendar for the user running the flow. Power Automate flows using this template are available on the **Cases dashboard**.

- **Create record for insider risk case in ServiceNow**: This template is for organizations that want to use their ServiceNow solution to track insider risk management cases. When in a case, insider risk analysts and investigators can create a record for the case in ServiceNow. You can customize this template to populate selected fields in ServiceNow based on your organization's requirements. Power Automate flows using this template are available on the **Cases dashboard**. For more information on available ServiceNow fields, see the ServiceNow Connector reference article.

Create a Power Automate flow from insider risk management template

To create a Power Automate flow from a recommended insider risk management template, you'll use the settings controls in the **Insider risk management** solution in the Microsoft 365 compliance center or the **Manage Power Automate flows** option from the **Automate** control when working directly in the **Cases** or **Users dashboards**.

To create a Power Automate flow in the settings area, you must be a member of the *Insider Risk Management* or *Insider Risk Management Admin* role group. To create a Power Automate flow with the **Manage Power Automate flows** option, you must be a member of at least one insider risk management role group.

Complete the following steps to create a Power Automate flow from a recommended insider risk management template:

1. In the Microsoft 365 compliance center, go to **Insider risk management** and select **Insider risk settings** > **Power Automate flows**. You can also access from the **Cases** or **Users dashboards** pages by choosing **Automate** > **Manage Power Automate flows**.
2. On the **Power Automate flows** page, select a recommended template from the **Insider risk management templates you may like** section on the page.
3. The flow lists the embedded connections needed for the flow and will note if the connection statuses are available. If needed, update any connections that aren't displayed as available. Select **Continue**.
4. By default, the recommended flows are pre-configured with the recommended insider risk management and Microsoft 365 service data fields required to complete the assigned task for the flow. If needed, customize the flow components by using the **Show advanced options** control and configuring the available properties for the flow component.
5. If needed, add any other steps to the flow by selecting the **New step** button. In most cases, this should not be needed for the recommended default templates.
6. Select **Save draft** to save the flow for further configuration or select **Save** to complete the configuration for the flow.
7. Select **Close** to return to the **Power Automate flow** page. The new template will be listed as a flow on the **My flows** tabs and is automatically available from the **Automate** dropdown control when working with insider risk management cases for the user creating the flow.

> **IMPORTANT**
> If other users in your organization need access to the flow, the flow must be shared.

Create a custom Power Automate flow for insider risk management

Some processes and workflows for your organization may be outside of the recommended insider risk management flow templates and you may have the need to create custom Power Automate flows for insider risk management areas. Power Automate flows are flexible and support extensive customization, but there are steps that need to be taken to integrate with insider risk management features.

Complete the following steps to create a custom Power Automate template for insider risk management:

1. **Check your Power Automate flow license**: To create customized Power Automate flows that use insider risk management triggers, you'll need a Power Automate license. The recommended insider risk management flow templates do not require extra licensing and are included as part of your insider risk management license.
2. **Create an automated flow**: Create a flow that performs one or more tasks after it's triggered by an insider risk management event. For details on how to create an automated flow, see Create a flow in Power Automate.
3. **Select the Microsoft 365 compliance connector**: Search for and select the Microsoft 365 compliance connector. This connector enables insider risk management triggers and actions. For more information on

connectors, see the Connector reference overview article.

4. **Choose insider risk management triggers for your flow**: Insider risk management has two triggers available for custom Power Automate flows:
   - **For a selected insider risk management case**: Flows with this trigger can be selected from the insider risk management Cases dashboard page.
   - **For a selected insider risk management user**: Flows with this trigger can be selected from the insider risk management Users dashboard page.
5. Choose insider risk management actions for your flow: You can choose from several actions for insider risk management to include in your custom flow:
   - Get insider risk management alert
   - Get insider risk management case
   - Get insider risk management user
   - Get insider risk management alerts for a case
   - Add insider risk management case note

Share a Power Automate flow

By default, Power Automate flows created by a user are only available to that user. For other insider risk management users to have access and use a flow, the flow must be shared by the flow creator. To share a flow, you'll use the settings controls in the **Insider risk management solution** in the Microsoft 365 compliance center or the **Manage Power Automate flows** option from the Automate control when working directly in the **Cases** or **Users dashboard** pages. Once you have shared a flow, everyone who it has been shared with can access the flow in the **Automate** control dropdown in the **Case** and **User dashboards**.

To share a Power Automate flow in the settings area, you must be a member of the *Insider Risk Management* or *Insider Risk Management Admin* role group. To share a Power Automate flow with the **Manage Power Automate flows** option, you must be a member of at least one insider risk management role group.

Complete the following steps to share a Power Automate flow:

1. In the Microsoft 365 compliance center, go to **Insider risk management** and select **Insider risk settings** > **Power Automate flows**. You can also access from the **Cases** or **Users dashboards** pages by choosing **Automate** > **Manage Power Automate flows**.
2. On the **Power Automate flows** page, select the **My flows** or **Team flows** tab.
3. Select the flow to share, then select **Share** from the flow options menu.
4. On the flow sharing page, enter the name of the user or group you want to add as an owner for the flow.
5. On the **Connection Used** dialog, select **OK** to acknowledge that the added user or group will have full access to the flow.

Edit a Power Automate flow

To edit a flow, you'll use the settings controls in the **Insider risk management** solution in the Microsoft 365 compliance center or the **Manage Power Automate flows** option from the **Automate** control when working directly in the **Cases** or **Users dashboards**.

To edit a Power Automate flow in the settings area, you must be a member of the *Insider Risk Management* or *Insider Risk Management Admin* role group. To edit a Power Automate flow with the **Manage Power Automate flows** option, you must be a member of at least one insider risk management role group.

Complete the following steps to edit a Power Automate flow:

1. In the Microsoft 365 compliance center, go to **Insider risk management** and select **Insider risk settings** > **Power Automate flows**. You can also access from the **Cases** or **Users dashboards** pages by choosing **Automate** > **Manage Power Automate flows**.
2. On the **Power Automate flows** page, select a flow to edit and select **Edit** from the flow control menu.

3. Select the **ellipsis** > **Settings** to change a flow component setting or **ellipsis** > **Delete** to delete a flow component.
4. Select **Save** and then **Close** to complete editing the flow.

Delete a Power Automate flow

To delete a flow, you'll use the settings controls in the **Insider risk management** solution in the Microsoft 365 compliance center or the **Manage Power Automate flows** option from the **Automate** control when working directly in the **Cases** or **Users dashboards**. When a flow is deleted, it is removed as an option for all users.

To delete a Power Automate flow in the settings area, you must be a member of the *Insider Risk Management* or *Insider Risk Management Admin* role group. To delete a Power Automate flow with the **Manage Power Automate flows** option, you must be a member of at least one insider risk management role group.

Complete the following steps to delete a Power Automate flow:

1. In the Microsoft 365 compliance center, go to **Insider risk management** and select **Insider risk settings** > **Power Automate flows**. You can also access from the **Cases** or **Users dashboards** pages by choosing **Automate** > **Manage Power Automate flows**.
2. On the **Power Automate flows** page, select a flow to delete and select **Delete** from the flow control menu.
3. On the deletion confirmation dialog, select **Delete** to remove the flow or select **Cancel** to exit the deletion action.

Microsoft Teams (preview)

Compliance analysts and investigators can easily use Microsoft Teams for collaboration on insider risk management cases. They can coordinate and communicate with other stakeholders in Microsoft Teams to:

- Coordinate and review response activities for cases in private Teams channels
- Securely share and store files and evidence related to individual cases
- Track and review response activities by analysts and investigators

After Microsoft Teams is enabled for insider risk management, a dedicated Microsoft Teams team is created every time an alert is confirmed and a case is created. By default, the team automatically includes all members of the *Insider Risk Management*, *Insider Risk Management Analysts*, and *Insider Risk Management Investigators* role groups (up to 100 initial users). Additional organization contributors may be added to the team after it is created and as appropriate. For existing cases created before enabling Microsoft Teams, analysts and investigators can choose to create a new Microsoft Teams team when working in a case if needed. Once you resolve the associated case in insider risk management, the team is automatically archived (moved to hidden and read-only).

For more information on how to use teams and channels in Microsoft Teams, see Overview of teams and channels in Microsoft Teams.

Enabling Microsoft Teams support for cases is quick and easy to configure. To enable Microsoft Teams for insider risk management, complete the following steps:

1. In the Microsoft 365 compliance center, go to **Insider risk management** > **Insider risk settings**.
2. Select the **Microsoft Teams** tab.
3. Enable Microsoft Teams integration for insider risk management.
4. Select **Save** to configure and exit.

Create a Microsoft Teams team for existing cases

If you enable Microsoft Teams support for insider risk management after you have existing cases, you'll need to manually create a team for each case as needed. After enabling Microsoft Teams support in insider risk management settings, new cases will automatically create a new Microsoft Teams team.

Users need permission to create Microsoft 365 groups in your organization to create a Microsoft Teams team from a case. For more information about managing permissions for Microsoft 365 Groups, see Manage who can create Microsoft 365 Groups.

To create a team for a case, you'll use the Create Microsoft Team control when working directly in an existing case. Complete the following steps to create a new team:

1. In the Microsoft 365 compliance center, go to **Insider risk management** > **Cases** and select an existing case.
2. On the case action menu, select **Create Microsoft Team**.
3. In the **Team name** field, enter a name for the new Microsoft Teams team.
4. Select **Create Microsoft team** and then select **Close**.

Depending on the number of users assigned to insider risk management role groups, it may take 15 minutes for all investigators and analysts to be added to the Microsoft Teams team for a case.

Get started with insider risk management

2/18/2021 • 14 minutes to read • Edit Online

Use insider risk management policies to identify risky activities and management tools to act on risk alerts in your organization. Complete the following steps to set up prerequisites and configure an insider risk management policy.

> **IMPORTANT**
>
> The Microsoft 365 insider risk management solution provides a tenant level option to help customers facilitate internal governance at the user level. Tenant level administrators can set up permissions to provide access to this solution for members of your organization and set up data connectors in the Microsoft 365 compliance center to import relevant data to support user level identification of potentially risky activity. Customers acknowledge insights related to the individual user's behavior, character, or performance materially related to employment can be calculated by the administrator and made available to others in the organization. In addition, customers acknowledge that they must conduct their own full investigation related to the individual user's behavior, character, or performance materially related to employment, and not just rely on insights from the insider risk management service. Customers are solely responsible for using the Microsoft 365 insider risk management service, and any associated feature or service in compliance with all applicable laws, including laws relating to individual user identification and any remediation actions.

For more information about how insider risk policies can help you manage risk in your organization, see Insider risk management in Microsoft 365.

Subscriptions and licensing

Before you get started with insider risk management, you should confirm your Microsoft 365 subscription and any add-ons. To access and use insider risk management, your organization must have one of the following subscriptions or add-ons:

- Microsoft 365 E5 subscription (paid or trial version)
- Microsoft 365 E3 subscription + the Microsoft 365 E5 Compliance add-on
- Microsoft 365 E3 subscription + the Microsoft 365 E5 Insider Risk Management add-on
- Microsoft 365 A5 subscription (paid or trial version)
- Microsoft 365 A3 subscription + the Microsoft 365 A5 Compliance add-on
- Microsoft 365 A3 subscription + the Microsoft 365 A5 Insider Risk Management add-on
- Microsoft 365 G5 subscription (paid or trial version)
- Microsoft 365 G3 subscription + the Microsoft 365 G5 Compliance add-on
- Microsoft 365 G3 subscription + the Microsoft 365 G5 Insider Risk Management add-on
- Office 365 E3 subscription + Enterprise Mobility and Security E3 + the Microsoft 365 E5 Compliance add-on

Users included in insider risk management policies must be assigned one of the licenses above.

If you don't have an existing Microsoft 365 Enterprise E5 plan and want to try insider risk management, you can add Microsoft 365 to your existing subscription or sign up for a trial of Microsoft 365 Enterprise E5.

Step 1: Enable permissions for insider risk management

IMPORTANT

After configuring your role groups, it may take up to 30 minutes for the role group permissions to apply to assigned users across your organization.

There are four roles groups used to configure permissions to manage insider risk management features. To continue with these configuration steps, your tenant administrators must first assign you to the **Insider Risk Management** or **Insider Risk Management Admin** role group. To access and manage insider risk management features after initial configuration, users must be a member of at least one insider risk management role group.

Depending on the structure of your compliance management team, you have options to assign users to specific role groups to manage different sets of insider risk management features. To view the **Permissions** tab in the Office 365 Security & Compliance Center and manage role groups, you need to be assigned to the *Organization Management* role group or need to be assigned the *Role Management* role. Choose from these role group options when configuring insider risk management:

ROLE GROUP	ROLE PERMISSIONS
Insider Risk Management	Use this role group to manage insider risk management for your organization in a single group. By adding all user accounts for designated administrators, analysts, and investigators, you can configure insider risk management permissions in a single group. This role group contains all the insider risk management permission roles. This configuration is the easiest way to quickly get started with insider risk management and is a good fit for organizations that do not need separate permissions defined for separate groups of users.
Insider Risk Management Admin	Use this role group to initially configure insider risk management and later to segregate insider risk administrators into a defined group. Users in this role group can create, read, update, and delete insider risk management policies, and global settings.
Insider Risk Management Analysts	Use this group to assign permissions to users that will act as insider risk case analysts. Users in this role group can access all insider risk management alerts, cases, and notices templates. They cannot access the insider risk Content Explorer.
Insider Risk Management Investigators	Use this group to assign permissions to users that will act as insider risk data investigators. Users in this role group can access all insider risk management alerts, cases, notices templates, and the Content Explorer.

NOTE

These role groups are currently not supported on Privileged Identity Management (PIM). To learn more about PIM, see Assign Azure AD roles in Privileged Identity Management.

Add users to an insider risk management role group

Complete the following steps to add users to an insider risk management role group:

1. Sign into https://protection.office.com/permissions using credentials for an admin account in your

Microsoft 365 organization.

2. In the Security & Compliance Center, go to **Permissions**. Select the link to view and manage roles in Office 365.
3. Select the insider risk management role group you want to add users to, then select **Edit role group**.
4. Select **Choose members** from the left navigation pane, then select **Edit**.
5. Select **Add** and then select the checkbox for all users you want to add to the role group.
6. Select **Add**, then select **Done**.
7. Select **Save** to add the users to the role group. Select **Close** to complete the steps.

Step 2: Enable the audit log

Insider risk management uses audit logs for user insights and activities configured in policies. The audit logs are a summary of all activities associated with an insider risk management policy or anytime a policy is changed.

For step-by-step instructions to turn on auditing, see Turn audit log search on or off. After you turn on auditing, a message is displayed that says the audit log is being prepared and that you can run a search in a couple of hours after the preparation is complete. You only have to do this action once. For more information about the using the audit log, see Search the audit log.

Step 3: Configure prerequisites for templates

Most insider risk management templates have prerequisites that must be configured for policy indicators to generate relevant activity alerts. Configure the appropriate prerequisites depending on the policies you plan to configure for your organization.

Configure Microsoft 365 HR connector

Insider risk management supports importing user and log data imported from 3rd-party risk management and human resources platforms. The Microsoft 365 Human Resources (HR) data connector allows you to pull in human resources data from CSV files, including user termination dates, last employment dates, performance improvement plan notifications, performance review actions, and job level change status. This data helps drive alert indicators in insider risk management policies and is an important part of configuring full risk management coverage in your organization. If you configure more than one HR connector for your organization, insider risk management will automatically pull indicators from all HR connectors.

The Microsoft 365 HR connector is required when using the following policy templates:

- Departing user data theft
- Security policy violations by departing users
- Security policy violations by disgruntled users
- Data leaks by disgruntled users

See the Set up a connector to import HR data article for step-by-step guidance to configure the Microsoft 365 HR connector for your organization. After you've configured the HR connector, return to these configuration steps.

Configure Data Loss Prevention (DLP) policies

Insider risk management supports using DLP policies to help identify the intentional or accidental exposure of sensitive information to unwanted parties for High severity level DLP alerts. When configuring an insider risk management policy with any of the **Data leaks** templates, you must assign a specific DLP policy to the policy.

DLP policies help identify users to activate risk scoring in insider risk management for high severity DLP alerts

for sensitive information and are an important part of configuring full risk management coverage in your organization. For more information about insider risk management and DLP policy integration and planning considerations, see Insider risk management policies.

> **IMPORTANT**
>
> Make sure you've completed the following:
>
> - You understand and properly configure the in-scope users in both the DLP and insider risk management policies to produce the policy coverage you expect.
> - Make sure the **Incident reports** setting in the DLP policy for insider risk management used with these templates are configured for *High* severity level alerts. Insider risk management alerts won't be generated from DLP policies with the **Incident reports** field set at *Low* or *Medium*.

A DLP policy is required when using the following policy templates:

- General data leaks
- Data leaks by priority users

See the Create, test, and tune a DLP policy article for step-by-step guidance to configure DLP policies for your organization. After you've configured a DLP policy, return to these configuration steps.

Configure priority user groups

Insider risk management includes support for assigning priority user groups to policies to help identity unique risk activities for user with critical positions, high levels of data and network access, or a past history of risk behavior. Creating a priority user group and assigning users to the group help scope policies to the unique circumstances presented by these users.

A priority user group is required when using the following policy templates:

- Security policy violations by priority users
- Data leaks by priority users

See the Getting started with insider risk management settings article for step-by-step guidance to create a priority user group. After you've configured a priority user group, return to these configuration steps.

Configure Physical badging connector (optional)

Insider risk management supports importing user and log data imported from physical control and access platforms. The Physical badging connector allows you to pull in access data from JSON files, including user IDs, access point IDs, access time and dates, and access status. This data helps drive alert indicators in insider risk management policies and is an important part of configuring full risk management coverage in your organization. If you configure more than one Physical badging connector for your organization, insider risk management automatically pulls indicators from all Physical badging connectors. Information from the Physical badging connector supplements other insider risk signals when using all insider risk policy templates.

> **IMPORTANT**
>
> For insider risk management policies to use and correlate signal data related to departing and terminated users with event data from your physical control and access platforms, you must also configure the Microsoft 365 HR connector. If you enable the Physical badging connector without enabling the Microsoft 365 HR connector, insider risk management policies will only process events for unauthorized physical access for users in your organization.

See the Set up a connector to import physical badging data article for step-by-step guidance to configure the Physical badging connector for your organization. After you've configured the connector, return to these configuration steps.

Step 4: Configure insider risk settings

Insider risk settings apply to all insider risk management policies, regardless of the template you chose when creating a policy. Settings are configured using the **Insider risk settings** control located at the top of all insider risk management tabs. These settings control privacy, indicators, monitoring windows, and intelligent detections.

Before configuring a policy, define the following insider risk settings:

1. In the Microsoft 365 compliance center, go to **Insider risk management** and select **Insider risk settings** from the top-right corner of any page.
2. On the **Privacy** page, select a privacy setting for displaying usernames for policy alerts.
3. On the **Indicators** page, select the alert indicators you want to apply to all insider risk policies.

   > **IMPORTANT**
   > In order to receive alerts for risky activity defined in your policies, you must select one or more indicators.

4. On the **Policy timeframes** page, select the policy timeframes to go into effect for a user when they trigger a match for an insider risk policy.
5. On the **Intelligent detections** page, configure the following settings for insider risk policies:
   - Anomaly detections
   - Alert volume level
   - Microsoft Defender for Endpoint alert status
   - Domain settings
6. On the **Export alerts** page, enable export of insider risk alert information using the Office 365 Management APIs if needed.
7. On the **Priority user groups** page, create a priority user group and add users if not created in **Step 3**.
8. On the **Power Automate flows** page, configure a flow from insider risk flow templates or create a new flow. See the Getting started with insider risk management settings article for step-by-step guidance.
9. On the **Priority assets page**, configure priority assets to use data from your physical control and access platform imported by the Physical badging connector. See the Getting started with insider risk management settings article for step-by-step guidance.
10. On the **Microsoft Teams** page, enable Microsoft Teams integration with insider risk management to automatically create a team for case or user collaboration. See the Getting started with insider risk management settings article for step-by-step guidance.
11. Select **Save** to enable these settings for your insider risk policies.

Step 5: Create an insider risk management policy

Insider risk management policies include assigned users and define which types of risk indicators are configured for alerts. Before activities can trigger alerts, a policy must be configured.

1. In the Microsoft 365 compliance center, go to **Insider risk management** and select the **Policies** tab.
2. Select **Create policy** to open the policy wizard.
3. On the **New insider risk policy** page, complete the following fields:
   - **Name (required)**: Enter a friendly name for the policy.

- **Description (optional)**: Enter a description for the policy.
- **Choose policy template (required)**: Select one of the policy templates to define the types of risk indicators are monitored by the policy.

> **IMPORTANT**
> Most policy templates have prerequisites that must be configured for the policy to generate relevant alerts. If you haven't configured the applicable policy prerequisites, see **Step 3** above.

4. Select **Next** to continue.
5. On the **Users** page, select **Add user or group** or **Choose Priority user groups** to define which users or priority user groups are included in the policy, depending on the policy template you've selected. Select **All users and mail-enabled groups** checkbox if applicable (if you haven't selected a priority user-based template). Select **Next** to continue.
6. On the **Specify what content to prioritize (optional)** page, you can assign the sources to prioritize for increased risk scores. However, some activities won't generate an alert at all unless the related content contains built-in or custom sensitive info types or was specified as a priority on this page:
   - **SharePoint sites**: Select **Add SharePoint site** and select the SharePoint organizations you want to prioritize. For example, *"group1@contoso.sharepoint.com/sites/group1"*.
   - **Sensitive info type**: Select **Add sensitive info type** and select the sensitivity types you want to prioritize. For example, *"U.S. Bank Account Number"* and *"Credit Card Number"*.
   - **Sensitivity labels**: Select **Add sensitivity label** and select the labels you want to prioritize. For example, *"Confidential"* and *"Secret"*.
7. Select **Next** to continue.
8. On the **Select policy indicators** page, you'll see the indicators that you've defined as available on the **Insider risk settings** > **Indicators** page. If you selected a *Data leaks* template at the beginning of the wizard, you must select a DLP policy from the **DLP policy** dropdown list to enable triggering indicators for the policy. Select the indicators you want to apply to the policy. If you prefer not to use the default policy threshold settings for these indicators, disable the **Use default thresholds recommended by Microsoft** and enter the threshold values for each selected indicator. If you've selected at least one *Office* or *Device* indicator, select the **Risk score boosters** as appropriate. Risk score boosters are only applicable for selected indicators.

   > **IMPORTANT**
   > If indicators on this page can't be selected, you'll need to select the indicators you want to enable for all policies on the **Insider risk management** > **Settings** > **Policy indicators** page.

9. Select **Next** to continue.
10. On the **Policy timeframes** page, you'll see the activation window conditions for the policy that on the **Insider risk settings** > **Policy timeframes** page.
11. Select **Next** to continue.
12. On the **Review** page, review the settings you've chosen for the policy. Select **Edit** to change any of the policy values or select **Submit** to create and activate the policy.

Next steps

After you've completed these steps to create your first insider risk management policy, you'll start to receive

alerts from activity indicators after about 24 hours. Configure additional policies as needed using the guidance in Step 4 of this article or the steps in Create a new insider risk policy.

To learn more about investigating insider risk alerts and the **Alerts dashboard**, see Insider risk management alerts.

Insider risk management policies

2/18/2021 • 19 minutes to read • Edit Online

Insider risk management policies determine which users are in-scope and which types of risk indicators are configured for alerts. You can quickly create a policy that applies to all users in your organization or define individual users or groups for management in a policy. Policies support content priorities to focus policy conditions on multiple or specific Microsoft Teams, SharePoint sites, data sensitivity types, and data labels. Using templates, you can select specific risk indicators and customize event thresholds for policy indicators, effectively customizing risk scores and level and frequency of alerts. Additionally, risk score boosters and anomaly detections help identify user activity that is of higher importance or more unusual. Policies windows allow you to define the time frame to apply the policy to alert activities and are used to determine the duration of the policy once activated.

Policy dashboard

The **Policy dashboard** allows you to quickly see the policies in your organization and the current status of alerts associated with each policy.

- **Policy name**: The name assigned to the policy in the policy wizard.
- **Active alerts**: The number of active alerts for each policy.
- **Confirmed alerts**: The total number of alerts the resulted in cases from the policy in the last 365 days.
- **Actions taken on alerts**: The total number of alerts that were confirmed or dismissed for the last 365 days.
- **Policy effectiveness**: The percentage determined by total confirmed alerts divided by total actions taken on alerts (which is the sum of alerts that were confirmed or dismissed over the past year).
- **Active**: The status of the case, either *Yes* or *No*.

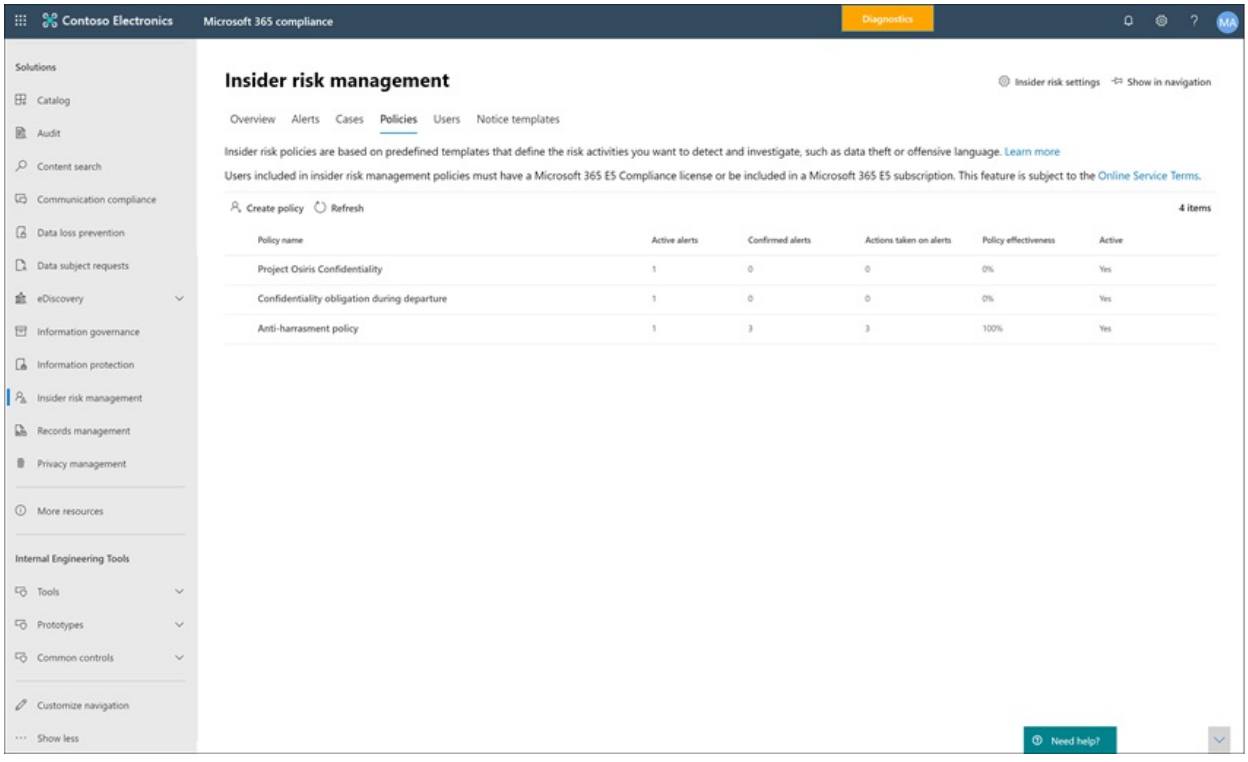

Policy templates

Insider risk management templates are pre-defined policy conditions that define the types of risk indicators and

risk scoring model used by the policy. Each policy must have a template assigned in the policy creation wizard before the policy is created. Insider risk management supports up to five policies for each policy template. When you create a new insider risk policy with the policy wizard, you'll choose from one of the following policy templates:

Data theft by departing users

When users leave your organization, there are specific risk indicators typically associated with data theft by departing users. This policy template uses indicators for risk scoring and focuses detection and alerts to this risk area. Data theft for departing users may include downloading files from SharePoint Online, printing files, and copying data to personal cloud messaging and storage services near their employment resignation and end dates. This template starts scoring for risk indicators relating to these activities and how they correlate with user employment status.

> **IMPORTANT**
> When using this template, you must configure a Microsoft 365 HR connector to periodically import resignation and termination date information for users in your organization. See the Import data with the HR connector article for step-by-step guidance to configure the Microsoft 365 HR connector for your organization.

General data leaks

Protecting data and preventing data leaks is a constant challenge for most organizations, particularly with the rapid grow of new data created by users, devices, and services. Users are empowered to create, store, and share information across services and devices that make managing data leaks increasingly more complex and difficult. Data leaks can include accidental oversharing of information outside your organization or data theft with malicious intent. In conjunction with an assigned Data Loss Prevention (DLP) policy, this template starts scoring real-time detections of suspicious SharePoint Online data downloads, file and folder sharing, printing files, and copying data to personal cloud messaging and storage services.

When using a **Data leaks** template, you must assign a DLP policy to trigger indicators in the insider risk policy for high severity alerts in your organization. Whenever a high severity alert is generated by a DLP policy rule is added to the Office 365 audit log, insider risk policies created with this template automatically examine the high severity DLP alert. If the alert contains an in-scope user defined in the insider risk policy, the alert is processed by the insider risk policy as a new alert and assigned an insider risk severity and risk score. This policy allows you to evaluate this alert in context with other activities included in the case.

Data leaks policy guidelines

When creating or modifying DLP policies for use with insider risk management policies, consider the following guidelines:

- Prioritize data exfiltration events and be selective when assigning **Incident reports** settings to *High* when configuring rules in your DLP policies. For example, emailing sensitive documents to a known competitor should be a *High* alert level exfiltration event. Over-assigning the *High* level in the **Incident reports** settings in other DLP policy rules can increase the noise in the insider risk management alert workflow and make it more difficult for your data investigators and analysts to properly evaluate these alerts. For example, assigning *High* alert levels to access denial activities in DLP policies makes it more challenging to evaluate truly risky user behavior and activities.
- Make sure you understand and properly configure the in-scope users in both the DLP and insider risk management policies. Only users defined as in-scope for insider risk management policies using the **Data leaks** template will have high severity DLP policy alerts processed. Additionally, only users defined as in-scope in a rule for a high severity DLP alert will be examined by the insider risk management policy for consideration. It is important that you don't unknowingly configure in-scope users in both your DLP and insider risk policies in a conflicting manner.

For example, if your DLP policy rules are scoped to only users on the Sales Team and the insider risk policy created from the **Data leaks** template has defined all users as in-scope, the insider risk policy will only actually process high severity DLP alerts for the users on the Sales Team. The insider risk policy won't receive any high priority DLP alerts for users to process that aren't defined in the DLP rules in this example. Conversely, if your insider risk management policy created from **Data leaks** templates is scoped to only users on the Sales Team and the assigned DLP policy is scoped to all users, the insider risk policy will only process high severity DLP alerts for members of the Sales Team. The insider risk management policy will ignore high severity DLP alerts for all users not on the Sales Team.

- Make sure the **Incident reports** rule setting in the DLP policy used for this insider risk management template is configured for *High* severity level alerts. The *High* severity level is the triggering events and insider risk management alerts won't be generated from rules in DLP policies with the **Incident reports** field set at *Low* or *Medium*.

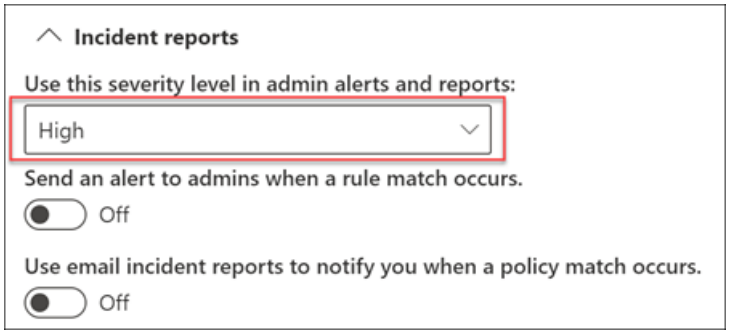

> **NOTE**
> When creating a new DLP policy using the built-in templates, you'll need to select the **Create or customize advanced DLP rules** option to configure the **Incident reports** setting for the *High* severity level.

Each insider risk management policy created from the **Data leaks** template can only have one DLP policy assigned. Consider creating a dedicated DLP policy that combines the different activities you want to detect and act as triggering events for insider risk policies that use the **Data leaks** template.

See the Create, test, and tune a DLP policy article for step-by-step guidance to configure DLP policies for your organization.

Data leaks by priority users (preview)

Protecting data and preventing data leaks for users in your organization may depend on their position, level of access to sensitive information, or risk history. Data leaks can include accidental oversharing of highly sensitive information outside your organization or data theft with malicious intent. In conjunction with an assigned Data Loss Prevention (DLP) policy, this template starts scoring real-time detections of suspicious activity and result in an increased likelihood of insider risk alerts and alerts with higher severity levels. Priority users are defined in priority user groups configured in the insider risk management settings area.

As with the **General data leaks template**, you must assign a DLP policy to trigger indicators in the insider risk policy for high severity alerts in your organization. Follow the Data leaks policy guidelines above when creating a policy using this template. Additionally, you will need to assign priority user groups created in **Insider risk management** > **Settings** > **Priority user groups** to the policy.

Data leaks by disgruntled users (preview)

When users experience employment stressors, they may become disgruntled which may increase the chances of insider risk activity. This template starts scoring user activity when an indicator associated with disgruntlement is identified. Examples include performance improvement notifications, poor performance reviews, or changes to job level status. Data leaks for disgruntled users may include downloading files from SharePoint Online and copying data to personal cloud messaging and storage services near employment stressor events.

When using this template, you must also configure a Microsoft 365 HR connector to periodically import performance improvement notifications, poor performance review status, or job level change information for users in your organization. See the Import data with the HR connector article for step-by-step guidance to configure the Microsoft 365 HR connector for your organization.

General security policy violations (preview)

In many organizations, users have permissions to install software on their devices or to modify device settings to help with their tasks. Either inadvertently or with malicious intent, users may install malware or disable important security features that help protect information on their device or on your network resources. This policy template uses security alerts from Microsoft Defender for Endpoint to start scoring these activities and focus detection and alerts to this risk area. Use this template to provide insights for security policy violations in scenarios when users may have a history of security policy violations that may be an indicator of insider risk.

You'll need to have Microsoft Defender for Endpoint configured in your organization and enable Defender for Endpoint for insider risk management integration in the Defender Security Center to import security violation alerts. For more information on configuring Defender for Endpoint for insider risk management integration, see Configure advanced features in Defender for Endpoint.

Security policy violations by departing users (preview)

Departing users, whether leaving on positive or negative terms, may be higher risks for security policy violations. To help protect against inadvertent or malicious security violations for departing users, this policy template uses Defender for Endpoint alerts to provide insights into security-related activities. These activities include the user installing malware or other potentially harmful applications and disabling security features on their devices. Policy indicators are activated after users have a resignation or termination date imported from the Microsoft 365 HR Connector as a triggering event.

When using this template, you must configure a Microsoft 365 HR connector to periodically import resignation and termination date information for users in your organization. See the Import data with the HR connector article for step-by-step guidance to configure the Microsoft 365 HR connector for your organization.

You'll need to have Microsoft Defender for Endpoint configured in your organization and enable Defender for Endpoint for insider risk management integration in the Defender Security Center to import security violation alerts. For more information on configuring Defender for Endpoint for insider risk management integration, see Configure advanced features in Defender for Endpoint.

Security policy violations by priority users (preview)

Protecting against security violations for users in your organization may depend on their position, level of access to sensitive information, or risk history. Because security violations by priority users may have an outsized impact on your organization's critical areas, this policy template starts scoring on these indicators and uses Microsoft Defender for Endpoint alerts to provide insights into security-related activities for these users. These may include the priority users installing malware or other potentially harmful applications and disabling security features on their devices. Priority users are defined in priority user groups configured in the insider risk management settings area.

You'll need to have Microsoft Defender for Endpoint configured in your organization and enable Defender for Endpoint for insider risk management integration in the Defender Security Center to import security violation alerts. For more information on configuring Defender for Endpoint for insider risk management integration, see Configure advanced features in Defender for Endpoint. Additionally, you will need to assign priority user groups created in **Insider risk management** > **Settings** > **Priority user groups** to the policy.

Security policy violations by disgruntled users (preview)

Users that experience employment stressors may be at a higher risk for inadvertent or malicious security policy violations. These stressors may include the user being placed on a performance improvement plan, poor performance review status, or being demoted from their current position. This policy template starts risk scoring based these indicators and activities associated with these events for these users.

When using this template, you must also configure a Microsoft 365 HR connector to periodically import performance improvement notifications, poor performance review status, or job level change information for users in your organization. See the Import data with the HR connector article for step-by-step guidance to configure the Microsoft 365 HR connector for your organization.

You'll also need to have Microsoft Defender for Endpoint configured in your organization and enable Defender for Endpoint for insider risk management integration in the Defender Security Center to import security violation alerts. For more information on configuring Defender for Endpoint for insider risk management integration, see Configure advanced features in Defender for Endpoint.

Policy template prerequisites and triggering events

Depending on the template you choose for an insider risk management policy, the triggering events and policy prerequisites vary. Triggering events are prerequisites that determine if a user is active for an insider risk management policy. If a user is added to an insider risk management policy but does not have a triggering event, the user activity is not evaluated by the policy unless they are manually added in the Users dashboard. Policy prerequisites are required items so that the policy receives the signals or activities necessary to evaluate risk.

The following table lists the triggering events and prerequisites for policies created from each insider risk management policy template:

POLICY TEMPLATE	TRIGGERING EVENTS FOR POLICIES	PREREQUISITES
Data theft by departing users	Resignation or termination date indicator from HR connector	Microsoft 365 HR connector configured for termination and resignation date indicators
General data leaks	Data leak policy activity that creates a High severity alert	DLP policy configured for High severity alerts
Data leaks by priority users	Data leak policy activity that creates a High severity alert	DLP policy configured for High severity alerts Priority user groups configured in insider risk settings
Data leaks by disgruntled users	Performance improvement, poor performance, or job level change indicators from HR connector	Microsoft 365 HR connector configured for disgruntlement indicators
General security policy violations	Defensive evasion of security controls or unwanted software detected by Microsoft Defender for Endpoint	Active Microsoft Defender for Endpoint subscription Microsoft Defender for Endpoint integration with Microsoft 365 compliance center configured
Security policy violations by departing users	Resignation or termination date indicators from HR connector	Microsoft 365 HR connector configured for termination and resignation date indicators Active Microsoft Defender for Endpoint subscription Microsoft Defender for Endpoint integration with Microsoft 365 compliance center configured

POLICY TEMPLATE	TRIGGERING EVENTS FOR POLICIES	PREREQUISITES
Security policy violations by priority users	Defensive evasion of security controls or unwanted software detected by Microsoft Defender for Endpoint	Active Microsoft Defender for Endpoint subscription Microsoft Defender for Endpoint integration with Microsoft 365 compliance center configured Priority user groups configured in insider risk settings
Security policy violations by disgruntled user	Performance improvement, poor performance, or job level change indicators from HR connector	Microsoft 365 HR connector configured for disgruntlement indicators Active Microsoft Defender for Endpoint subscription Microsoft Defender for Endpoint integration with Microsoft 365 compliance center configured

Prioritize content in policies

Insider risk management policies support specifying a higher priority for content depending where it is stored or how it is classified. Specifying content as a priority increases the risk score for any associated activity, which in turn increases the chance of generating a high severity alert. However, some activities won't generate an alert at all unless the related content contains built-in or custom sensitive info types or was specified as a priority in the policy.

For example, your organization has a dedicated SharePoint site for a highly confidential project. Data leaks for information in this SharePoint site could compromise the project and would have a significant impact on its success. By prioritizing this SharePoint site in a Data leaks policy, risk scores for qualifying activities are automatically increased. This prioritization increases the likelihood that these activities generate an insider risk alert and raises the severity level for the alert.

When you create an insider risk management policy in the policy wizard, you can choose from the following priorities:

- **SharePoint sites**: Any activity associated with all file types in defined SharePoint sites is assigned a higher risk score.
- **Sensitive information types**: Any activity associated with content that contains sensitive information types are assigned a higher risk score.
- **Sensitivity labels**: Any activity associated with content that has specific sensitivity labels applied are assigned a higher risk score.

Create a new policy

To create a new insider risk management policy, you'll use the policy wizard in **Insider risk management** solution in the Microsoft 365 compliance center.

Complete the following steps to create a new policy:

1. In the Microsoft 365 compliance center, go to **Insider risk management** and select the **Policies** tab.
2. Select **Create policy** to open the policy wizard

3. On the **New insider risk policy** page, complete the following fields:

   - **Name (required)**: Enter a friendly name for the policy.
   - **Description (optional)**: Enter a description for the policy.
   - **Choose policy template (required)**: Select one of the policy templates to define the types of risk indicators are monitored by the policy.

   > **IMPORTANT**
   > Most policy templates have prerequisites that must be configured for the policy to generate relevant alerts. If you haven't configured the applicable policy prerequisites, see Get started with insider risk management.

4. Select **Next** to continue.

5. On the **Users** page, select **Add user or group** or **Choose Priority user groups** to define which users or priority user groups are included in the policy, depending on the policy template you've selected. Select **All users and mail-enabled groups** checkbox if applicable (if you haven't selected a priority user-based template). Select **Next** to continue.

6. On the **Specify what content to prioritize (optional)** page, you can assign the sources to prioritize for increased risk scores. However, some activities won't generate an alert at all unless the related content contains built-in or custom sensitive info types or was specified as a priority on this page:

   - **SharePoint sites**: Select **Add SharePoint site** and select the SharePoint organizations you want to prioritize. For example, *"group1@contoso.sharepoint.com/sites/group1"*.
   - **Sensitive info type**: Select **Add sensitive info type** and select the sensitivity types you want to prioritize. For example, *"U.S. Bank Account Number"* and *"Credit Card Number"*.
   - **Sensitivity labels**: Select **Add sensitivity label** and select the labels you want to prioritize. For example, *"Confidential"* and *"Secret"*.

7. Select **Next** to continue.

8. On the **Select policy indicators** page, you'll see the indicators that you've defined as available on the **Insider risk settings** > **Indicators** page. If you selected a *Data leaks* template at the beginning of the wizard, you must select a DLP policy from the **DLP policy** dropdown list to enable triggering indicators for the policy. Select the indicators you want to apply to the policy. If you prefer not to use the default policy threshold settings for these indicators, disable the **Use default thresholds recommended by Microsoft** and enter the threshold values for each selected indicator. If you've selected at least one *Office* or *Device* indicator, select the **Risk score boosters** as appropriate. Risk score boosters are only applicable for selected indicators.

   > **IMPORTANT**
   > If indicators on this page can't be selected, you'll need to select the indicators you want to enable for all policies on the **Insider risk management** > **Settings** > **Policy indicators** page.

9. Select **Next** to continue.

10. On the **Policy timeframes** page, you'll see the activation window conditions for the policy that on the **Insider risk settings** > **Policy timeframes** page.

11. Select **Next** to continue.

12. On the **Review** page, review the settings you've chosen for the policy. Select **Edit** to change any of the policy values or select **Submit** to create and activate the policy.

Update a policy

To update an existing insider risk management policy, you'll use the policy wizard in **Insider risk management** solution in the Microsoft 365 compliance center.

Complete the following steps to manage an existing policy:

1. In the Microsoft 365 compliance center, go to **Insider risk management** and select the **Policies** tab.
2. On the policy dashboard, select the policy you want to manage.
3. On the policy details page, select **Edit policy**
4. In the policy wizard, you cannot edit the following fields:
   - **Name**: The friendly name for the policy
   - **Choose policy template**: The template used to define the types of risk indicators monitored by the policy.
5. Enter a new description for the policy in the **Description** field.
6. Select **Next** to continue.
7. On the **Users** page, select **Add user or group** or **Choose Priority user groups** to define which users or priority user groups are included in the policy, depending on the policy template you've selected. Select **All users and mail-enabled groups** checkbox if applicable (if you haven't selected a priority user-based template). Select **Next** to continue.
8. On the **Specify what content to prioritize (optional)** page, you can assign the sources to prioritize for increased risk scores. However, some activities won't generate an alert at all unless the related content contains built-in or custom sensitive info types or was specified as a priority on this page:
   - **SharePoint sites**: Select **Add SharePoint site** and select the SharePoint organizations you want to prioritize. For example, *"group1@contoso.sharepoint.com/sites/group1"*.
   - **Sensitive info type**: Select **Add sensitive info type** and select the sensitivity types you want to prioritize. For example, *"U.S. Bank Account Number"* and *"Credit Card Number"*.
   - **Sensitivity labels**: Select **Add sensitivity label** and select the labels you want to prioritize. For example, *"Confidential"* and *"Secret"*.
9. Select **Next** to continue.
10. On the **Select policy indicators** page, you'll see the indicators that you've defined as available on the **Insider risk settings** > **Indicators** page. If you selected a *Data leaks* template at the beginning of the wizard, you must select a DLP policy from the **DLP policy** dropdown list to enable triggering indicators for the policy. Select the indicators you want to apply to the policy. If you prefer not to use the default policy threshold settings for these indicators, disable the **Use default thresholds recommended by Microsoft** and enter the threshold values for each selected indicator. If you've selected at least one *Office* or *Device* indicator, select the **Risk score boosters** as appropriate. Risk score boosters are only applicable for selected indicators.

    > **IMPORTANT**
    > If indicators on this page can't be selected, you'll need to select the indicators you want to enable for all policies on the **Insider risk management** > **Settings** > **Policy indicators** page.

11. Select **Next** to continue.
12. On the **Policy timeframes** page, you'll see the activation window conditions for the policy that on the **Insider risk settings** > **Policy timeframes** page.

13. Select **Next** to continue.
14. On the **Review** page, review the settings you've updated for the policy. Select **Edit** to change any of the policy values or select **Submit** to update and activate the policy.

Delete a policy

> **NOTE**
> Deleting a policy does not delete active or archived alerts generated from the policy.

To delete an existing insider risk management policy, complete the following steps:

1. In the Microsoft 365 compliance center, go to **Insider risk management** and select the **Policies** tab.
2. On the policy dashboard, select the policy you want to delete.
3. Select **Delete** on the dashboard toolbar.
4. On the **Delete** dialog, Select **Yes** to delete the policy, or select **Cancel** to close the dialog.

Insider risk management alerts

11/2/2020 • 7 minutes to read • Edit Online

Insider risk management alerts are automatically generated by risk indicators defined in insider risk management policies. These alerts give compliance analysts and investigators an all-up view of the current risk status and allow your organization to triage and take actions for discovered risks. By default, policies generate a certain amount of low, medium, and high severity alerts, but you can increase or decrease the alert volume to suit your needs. Additionally, you can configure the alert threshold for policy indicators when creating a new policy with the policy wizard.

Alert dashboard

The insider risk **Alert dashboard** allows you to view and act on alerts generated by insider risk policies. Each report widget displays information for last 30 days.

- **Alerts to review**: The total number of alerts needing review and triage are listed, including a breakdown by alert severity.
- **Open alerts over past 30 days**: The total number of alerts created by policy matches over the last 30 days, sorted by high, medium, and low alert severity levels.
- **Average time to resolve alerts**: A summary of useful alert statistics:
  - Average time to resolve high severity alerts, listed in hours, days, or months.
  - Average time to resolve medium severity alerts, listed in hours, days, or months.
  - Average time to resolve low severity alerts, listed in hours, days, or months.

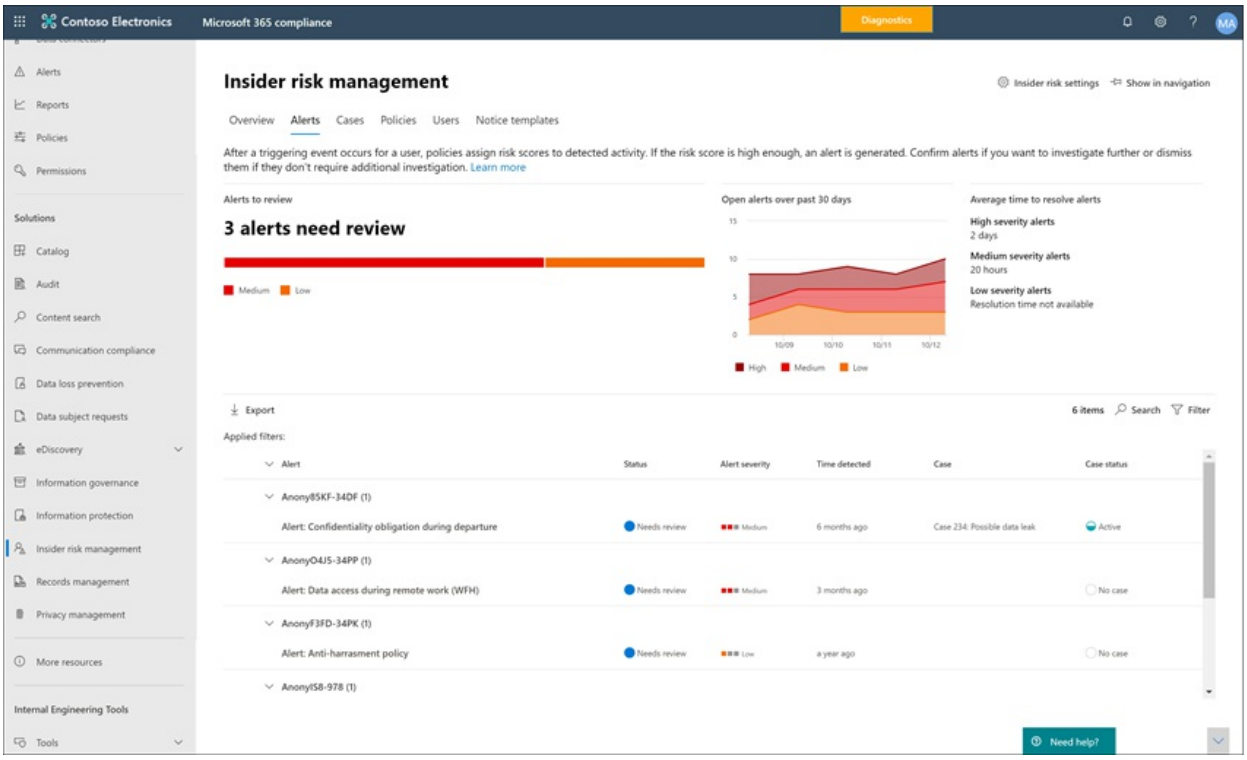

> **NOTE**
>
> Insider risk management uses built-in alert throttling to help protect and optimize your risk investigation and review experience. This throttling guards against issues that might result in an overload of policy alerts, such as misconfigured data connectors or DLP policies. As a result, there might be a delay in displaying new alerts for a user.

Alert status and severity

You can triage alerts into one of the following statuses:

- **Confirmed**: An alert confirmed and assigned to a new or existing case.
- **Dismissed**: An alert dismissed as benign in the triage process.
- **Needs review**: A new alert where triage actions have not yet been taken.
- **Resolved**: An alert that is part of a closed and resolved case.

Alert risk scores are automatically calculated from several risk activity indicators. These indicators include the type of risk activity, the number and frequency of the activity occurrence, the history of user risk activity, and the addition of activity risks that may boost the seriousness of the activity. The alert risk score drives the programmatic assignment of a risk severity level for each alert and cannot be customized. If alerts remain untriaged and risk activities continue to accrue to the alert, the risk severity level can increase. Risk analysts and investigators can use the alert risk severity to help triage alerts in accordance with your organization's risk policies and standards.

Alert risk severity levels are:

- **High severity**: The activities and indicators for the alert pose significant risk. The associated risk activities are serious, repetitive, and corelate strongly to other significant risk factors.
- **Medium severity**: The activities and indicators for the alert pose a moderate risk. The associated risk activities are moderate, frequent, and have some correlation to other risk factors.
- **Low severity**: The activities and indicators for the alert pose a minor risk. The associated risk activities are minor, more infrequent, and do not corelate to other significant risk factors.

Filter alerts on the Alert dashboard

Depending on the number and type of active insider risk management policies in your organization, reviewing a large queue of alerts can be challenging. Using alert filters can help analysts and investigators sort alerts by several attributes. To filter alerts on the **Alerts dashboard**, select the **Filter** control. You can filter alerts by one or more attributes:

- **Status**: Select one or more status values to filter the alert list. The options are *Confirmed*, *Dismissed*, *Needs review*, and *Resolved*.
- **Severity**: Select one or more alert risk severity levels to filter the alert list. The options are *High*, *Medium*, and *Low*.
- **Time detected**: Select the start and end dates for when the alert was created.
- **Policy**: Select one or more policies to filter the alerts generated by the selected policies.

Search alerts on the Alert dashboard

To search the alert name for a specific word, select the **Search** control and type the word to search. The search results display any policy alert containing the word defined in the search.

Triage alerts

To triage an insider risk alert, complete the following steps:

1. In the Microsoft 365 compliance center, go to **Insider risk management** and select the **Alerts** tab.
2. On the **Alerts dashboard**, select the alert you want to triage.
3. On the **Alerts detail pane**, you can review the following tabs and triage the alert:
   - **Summary**: This tab contains general information about the alert and allows you to confirm the alert and create a new case or allows you to dismiss the alert. It includes the current status for the alert and the alert risk severity level, listed as *High*, *Medium*, or *Low*. The severity level may increase or decrease over time if the alert is not triaged.
     - **What happened**: Displays the top three risk activities and policy matches during the activity evaluation period, including the type of violation associated with the activity.
     - **User details**: Displays general information about the user assigned to the alert. If anonymization is enabled, the username, email address, alias, and organization fields are anonymized.
     - **Alert details**: Includes the length of time since the alert was generated, the policies that generated the alert are listed, and the case generated from the alert is listed. For new alerts, the **Case** field displays None.
     - **Content detected**: Includes content associated with the risk activities for the alert and summarizes activity events by key areas. Selecting an activity link opens the Activity explorer and displays additional details about the activity.
   - **User activity**: This tab displays the activity history for the user associated with the alert. This history includes other alerts and activities related to risk indicators defined in the template assigned to the policy for this alert. This history allows risk analysts and investigators to factor in any past risky behavior for the employee as part of the triage process.
   - **Actions**: The following actions are available for each alert:
     - **Open expanded view**: Opens the **Activity explorer** dashboard.
     - **Confirm and create case**: Use this action to confirm and create a new case for all the alerts associated with a user. This action automatically changes the alert status to *Confirmed*.
     - **Dismiss alert**: Use this action to dismiss the alert. This action changes the alert status to *Resolved*.

Activity explorer (preview)

> **NOTE**
> Activity explorer is available in the alert management area for users with triggering events after this feature is available in your organization.

The Activity explorer provides risk investigators and analysts with a comprehensive analytic tool that provides detailed information about alerts. With the Activity explorer, reviewers can quickly review a timeline of detected risky activity and identify and filter all risk activities associated with alerts. To filter alerts on the Activity explorer, select the Filter control. You can filter alerts by one or more attributes listed in the details pane for the alert. Activity explorer also supports customizable columns to help investigators and analysts focus the dashboard on the information most important to them.

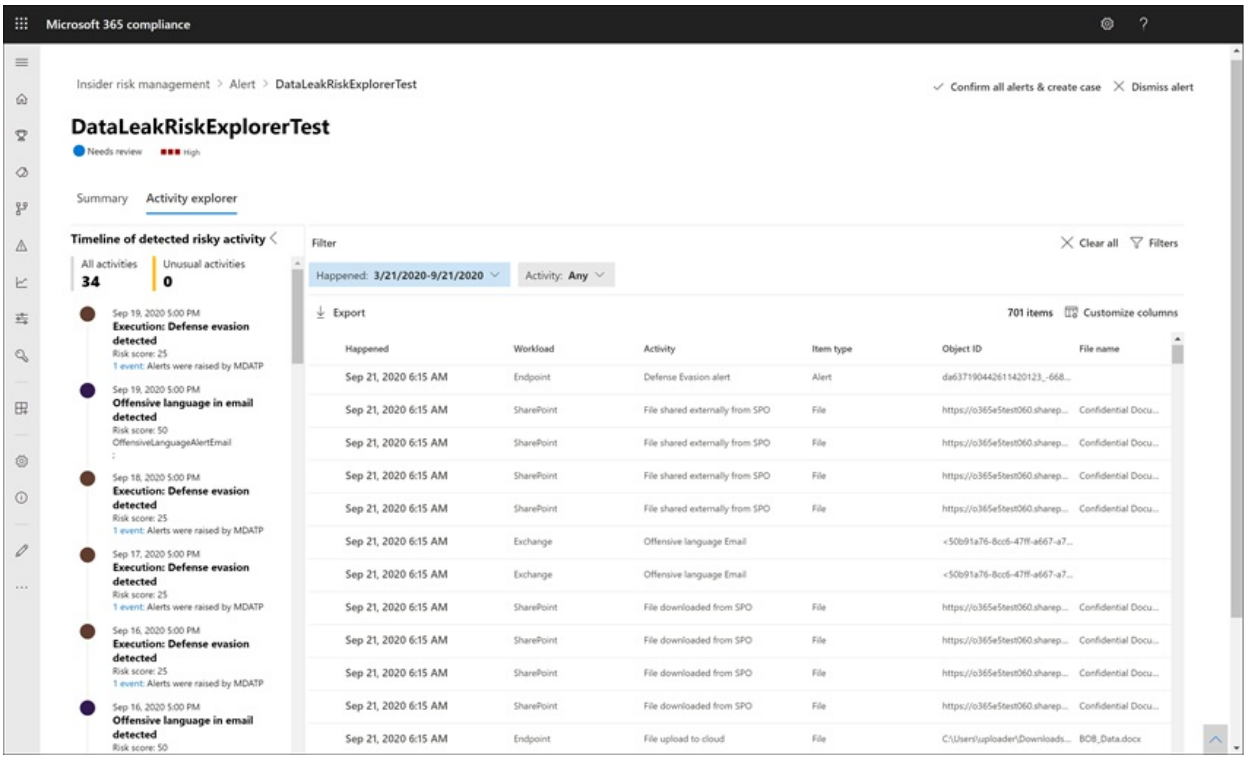

To use the **Activity explorer**, complete the following steps:

1. In the Microsoft 365 compliance center, go to **Insider risk management** and select the **Alerts** tab.
2. On the **Alerts dashboard**, select the alert you want to triage.
3. On the **Alerts detail pane**, select **Open expanded view**.
4. On the page for the selected alert, select the **Activity explorer** tab.

When reviewing activities in the Activity explorer, investigators and analysts can select a specific activity and open the activity details pane. The pane displays detailed information about the activity that investigators and analysts can use during the alert triage process. The detailed information may provide context for the alert and assist with identifying the full scope of the risk activity that triggered the alert.

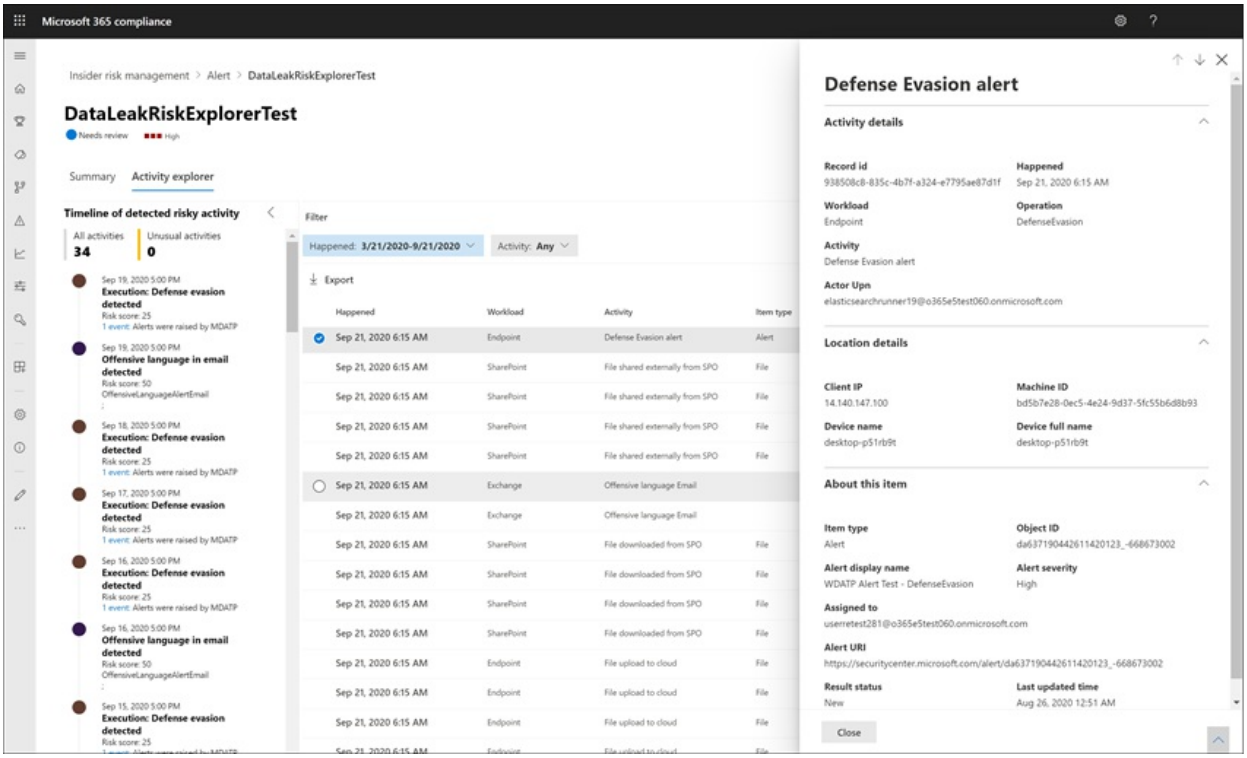

Create a case for an alert

As alert is reviewed and triaged, you can create a new case to further investigate the risk activity. To create a case for an alert, follow these steps:

1. In the Microsoft 365 compliance center, go to **Insider risk management** and select the **Alerts** tab.
2. On the **Alerts dashboard**, select the alert you want to confirm and create a new case for.
3. On the **Alerts details pane**, select **Actions** > **Confirm alerts & create case**.
4. On the **Confirm alert and create insider risk case** dialog, enter a name for the case, select users to add as contributors, and add comments as applicable. Comments are automatically added to the case as a case note.
5. Select **Create case** to create a new case or select **Cancel** to close the dialog without creating a case.

After the case is created, investigators and analysts can manage and act on the case. See the Insider risk management case article for more details.

Insider risk management cases

2/18/2021 • 16 minutes to read • Edit Online

Cases are the heart of insider risk management and allow you to deeply investigate and act on issues generated by risk indicators defined in your policies. Cases are manually created from alerts in situations where further action is needed to address a compliance-related issue for a user. Each case is scoped to a single user and multiple alerts for the user can be added to an existing case or to a new case.

After investigating the details of a case, you can take action by:

- sending the user a notice
- resolving the case as benign
- sharing the case with your ServiceNow instance or with an email recipient
- escalating the case for an Advanced eDiscovery investigation

Cases dashboard

The insider risk management **Cases dashboard** allows you to view and act on cases. Each report widget on the dashboard displays information for last 30 days.

- **Active cases**: The total number of active cases under investigation.
- **Cases over past 30 days**: The total number of cases created, sorted by *Active* and *Closed* status.
- **Statistics**: Average time of active cases, listed in hours, days, or months.

The case queue lists all active and closed cases for your organization, in addition to the current status of the following case attributes:

- **Case name**: The name of the case, defined when an alert is confirmed and the case is created.
- **Status**: The status of the case, either *Active* or *Closed*.
- **User**: The user for the case. If anonymization for usernames is enabled, anonymized information is displayed.
- **Time case opened**: The time that has passed since the case was opened.
- **Total policy alerts**: The number of policy matches included in the case. This number may increase if new alerts are added to the case.
- **Last updated**: The time that has passed since there has been an added case note or change in the case state.
- **Last updated by**: The name of the insider risk management analyst or investigator that last updated the case.

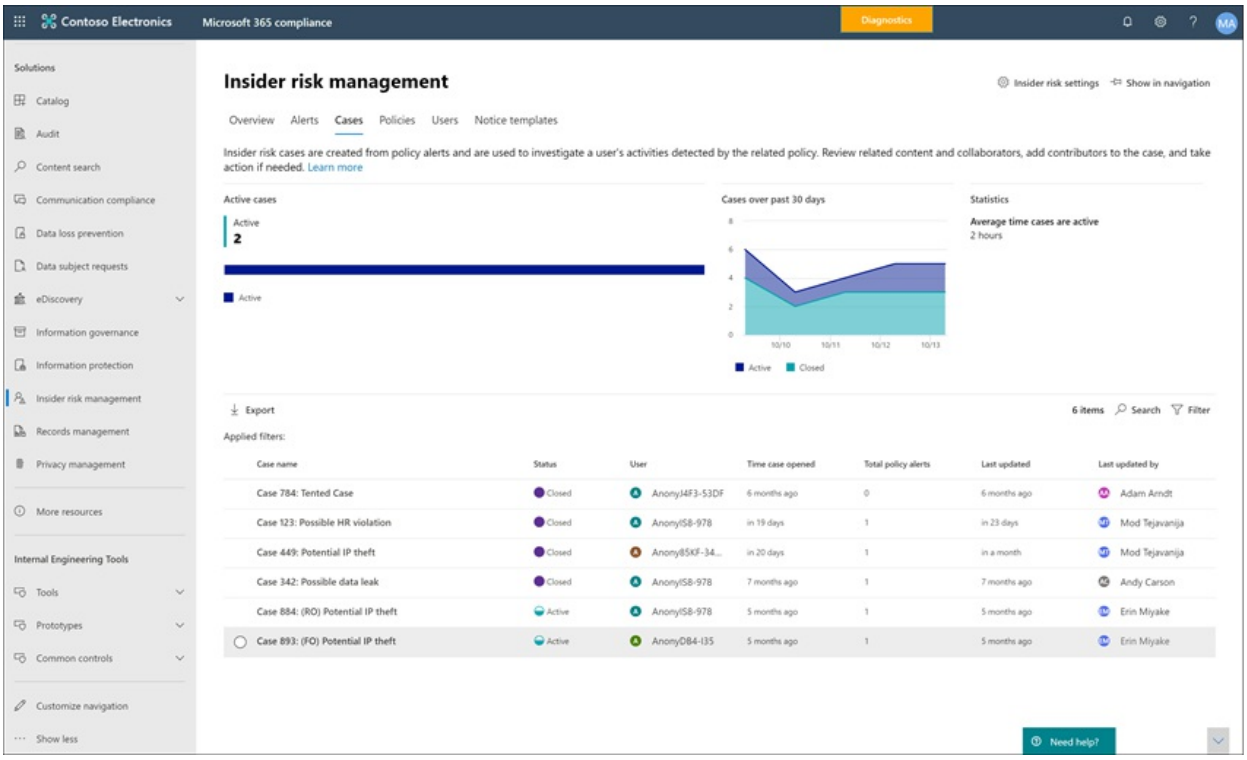

Use the **Search** control to search case names for specific text and use the case filter to sort cases by the following attributes:

- Status
- Time case opened, start date, and end date
- Last updated, start date, and end date

Filter cases

Depending on the number and type of active insider risk management policies in your organization, reviewing a large queue of cases can be challenging. Using case filters can help analysts and investigators sort cases by several attributes. To filter alerts on the **Cases dashboard**, select the **Filter** control. You can filter cases by one or more attributes:

- **Status**: Select one or more status values to filter the case list. The options are *Active* and *Closed*.
- **Time case opened**: Select the start and end dates for when the case was opened.
- **Last updated**: Select the start and end dates for when the case was updated.

Investigate a case

Deeper investigation into insider risk management alerts is critical to taking proper corrective actions. Insider risk management cases are the central management tool to dive deeper into user risk activity history and alert details, and to explore the content and messages exposed to risks. Risk analysts and investigators also use cases to centralize review feedback and notes and to process case resolution.

Selecting a case opens the case management tools and allows analysts and investigators to dig into the details of cases.

Case overview

The **Case overview** tab summarizes the alert activity and risk level history for the case.

- The **Alerts** widget shows the policy matches for the case, including the status of the alert, the alert risk severity, and when the alert was detected.
- The **Risk level history** chart displays the user risk level over the last 30 days. The line chart allows analysts

and investigators to quickly see the trend in overall user risk over time.

- The **Risk activity content** widget summarizes the types of data and content contained in alerts added to the case. This widget gives an all-up view of the entire data and content set at risk in the case.

The **Case details** pane is available on all case management tabs and summarizes the case details for risk analysts and investigators. It includes the following areas:

- **Case name**: The name of the case, prefixed with an autogenerated case sequence number and the name of the risk associated with the policy template that the first confirmed alert matches.
- **Case status**: The current status of the case, either *Active* or *Closed*.
- **User's risk score**: The current calculated risk level of the user for the case. This score is calculated every 24 hours and uses the alert risk scores from all active alerts associated to the user.
- **Alerts confirmed**: List of alerts for the user confirmed for the case.
- **Related content**: List of content, sorted by content sources and types. For example, for case alert content in SharePoint Online, you may see folder or file names listed that are associated with the risk activity for alerts in the case.

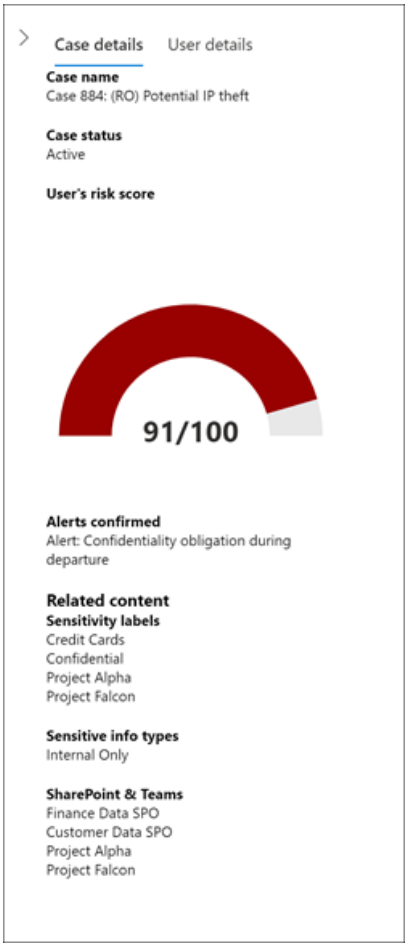

Alerts

The **Alerts** tab summarizes the current alerts included in the case. New alerts may be added to an existing case and they will be added to the **Alert** queue as they are assigned. The following alert attributes are listed the queue:

- Status
- Severity
- Time detected

Select an alert from the queue to display the **Alert detail** page.

Use the search control to search alert names for specific text and use the alert filter to sort cases by the

following attributes:

- Status
- Severity
- Time detected, start date, and end date

Use the filter control to filter alerts by several attributes, including:

- **Status**: Select one or more status values to filter the alert list. The options are *Confirmed*, *Dismissed*, *Needs review*, and *Resolved*.
- **Severity**: Select one or more alert risk severity levels to filter the alert list. The options are *High*, *Medium*, and *Low*.
- **Time detected**: Select the start and end dates for when the alert was created.
- **Policy**: Select one or more policies to filter the alerts generated by the selected policies.

User activity

The **User activity** tab is one of the most powerful tools for internal risk analysis and investigation for cases in the insider risk management solution. This tab is structured to enable quick review of a case, including a historical timeline of all alerts, all alerts details, the current risk score for the user in the case, and controls to take effective action to contain the risks in the case.

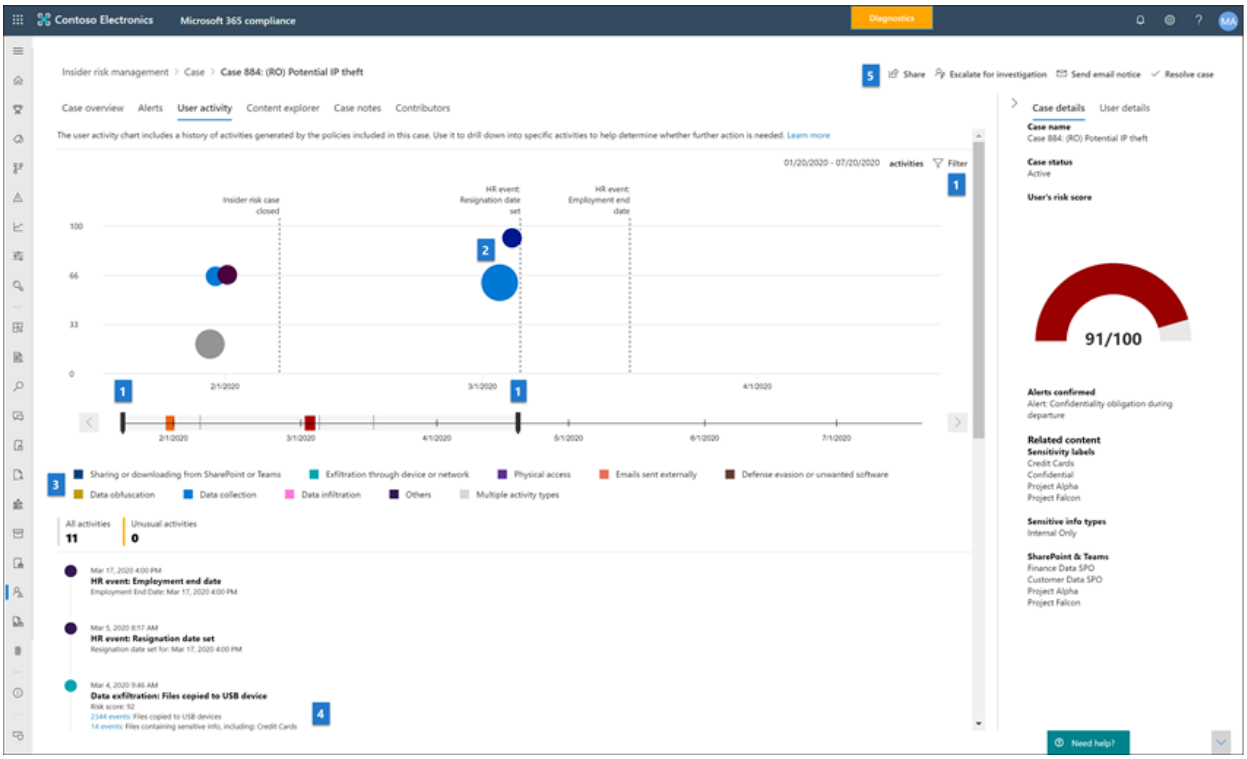

1. **Date and window time filters**: By default, the last six months of alerts confirmed in the case are displayed in the User activity chart. You can easily filter the chart view with either the slider controls at both ends of the chart window, or by defining specific start and end dates in the chart filter control.
2. **Risk alert activity and details**: Risk activities are visually displayed as colored bubbles in the User activity chart. Bubbles are created for different categories of risk and bubble size is proportional to the number of risk activities for the category. Select a bubble to display the details for each risk activity. Details include:
   - **Date** of the risk activity.
   - The **risk activity category**. For example, *Email(s) with attachments sent outside the organization* or *File(s) downloaded from SharePoint Online*.
   - **Risk score** for the alert. This score is the numerical score for the alert risk severity level.
   - Number of events associated with the alert. Links to each file or email associated with the risk activity is also available.

3. **Risk activity legend**: Across the bottom of the user activity chart, a color-coded legend helps you quickly determine risk category for each alert.
4. **Risk activity chronology**: The full chronology of all risk alerts associated with the case are listed, including all the details available in the corresponding alert bubble.
5. **Case actions**: Options for resolving the case are on the case action toolbar. You can resolve a case, send an email notice to the user, or escalate the case for a data or user investigation.

Activity explorer (preview)

> **IMPORTANT**
> The Activity explorer tab is available in the case management area for users with triggering events after this feature is available in your organization.

The **Activity explorer** tab allows risk analysts and investigators to review activity details associated with risk alerts. For example, as part of the case management actions, investigators and analysts may need to review all the risk activities associated with the case for more details. With the **Activity explorer**, reviewers can quickly review a timeline of detected risky activity and identify and filter all risk activities associated with alerts.

For more information about the Activity explorer, see the Insider risk management alerts article.

Content Explorer

The **Content Explorer** tab allows risk analysts and investigators to review copies of all individual files and email messages associated with risk alerts. For example, if an alert is created when a user downloads hundreds of files from SharePoint Online and the activity triggers a policy alert, all the downloaded files for the alert are captured and copied to the insider risk management case from original storage sources.

The Content Explorer is a powerful tool with basic and advanced search and filtering features. To learn more about using the Content Explorer, see Insider risk management Content Explorer.

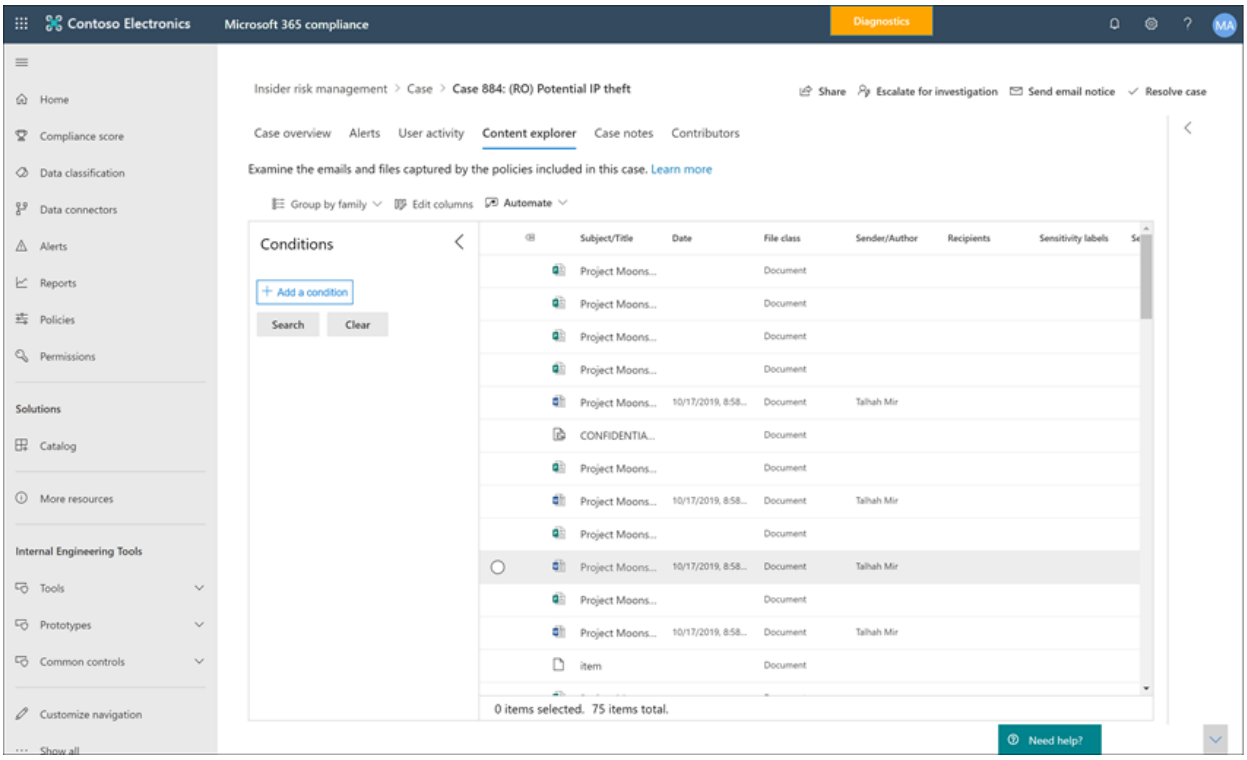

Case notes

The **Case notes** tab in the case is where risk analysts and investigators share comments, feedback, and insights about their work for the case. Notes are permanent additions to a case and cannot be edited or deleted after the note is saved. When a case is created from an alert, the comments entered in the **Confirm alert and create**

insider risk case dialog are automatically added as a case note.

The case notes dashboard displays notes by the user that created the note and the time that has passed since the note was saved. To search the case note text field for a specific keyword, use the **Search** button on the case dashboard and enter a specific keyword.

To add a note to a case:

1. In the Microsoft 365 compliance center, go to **Insider risk management** and select the **Cases** tab.
2. Select a case, then select the **Case notes** tab.
3. Select **Add case note**.
4. On the **Add case note** dialog, type your note for the case. Select **Save** to add the note to the case or select **Cancel** close without saving the note to the case.

Contributors

The **Contributors** tab in the case is where risk analysts and investigators can add other reviewers to the case. Be default, all users assigned the **Insider Risk Management Analysts** and **Insider Risk Management Investigators** roles are listed as contributors for each active and closed case. Only users assigned the **Insider Risk Management Investigators** role have permission to view files and messages in the Content Explorer.

Temporary access to a case can be granted by adding a user as a contributor. Contributors have all case management control on the specific case except:

- Permission to confirm or dismiss alerts
- Permission to edit the contributors for cases
- Permission to view files and messages in the Content Explorer

To add a contributor to a case:

1. In the Microsoft 365 compliance center, go to **Insider risk management** and select the **Cases** tab.
2. Select a case, then select the **Contributors** tab.
3. Select **Add contributor**.
4. On the **Add contributor** dialog, start typing the name of the user you want to add and then select the user from the suggested user list. This list is generated from the Azure Active Directory of your tenant subscription.
5. Select **Add** to add the user as a contributor or select **Cancel** close the dialog without adding the user as a contributor.

Case actions

Risk analysts and investigators can take action on a case in one of several methods, depending on the severity of the case, the history of risk of the user, and the risk guidelines of your organization. In some situations, you may need to escalate a case to a user or data investigation to collaborate with other areas of your organization and to dive deeper into risk activities. Insider risk management is tightly integrated with other Microsoft 365 compliance solutions to help you with end-to-end resolution management.

Send email notice

In most cases, user actions that create insider risk alerts are inadvertent or accidental. Sending a reminder notice to the user via email is an effective method for documenting case review and action, and is a method to remind users of corporate policies or point them to refresher training. Notices are generated from notice templates that you create for your insider risk management infrastructure.

It's important to remember that sending an email notice to a user ***does not*** resolve the case as *Closed*. In some cases, you may want to leave a case open after sending a notice to a user to look for more risk activities without opening a new case. If you want to resolve a case after sending a notice, you must select the **Resolve case** as a

follow-on step after sending a notice.

To send a notice to the user assigned to a case:

1. In the Microsoft 365 compliance center, go to **Insider risk management** and select the **Cases** tab.
2. Select a case, then select the **Send email notice** button on the case action toolbar.
3. On the **Send e-mail notice** dialog, select the **Choose a notice template** dropdown control to select the notice template for the notice. This selection pre-fills the other fields on the notice.
4. Review the notice fields and update as appropriate. The values entered here will override the values on the template.
5. Select **Send** to send the notice to the user or select **Cancel** close the dialog without sending the notice to the user. All sent notices are added to the case notes queue on the **Case notes** dashboard.

Escalate for investigation

Escalate the case for user investigation in situations where additional legal review is needed for the user's risk activity. This escalation opens a new Advanced eDiscovery case in your Microsoft 365 organization. Advanced eDiscovery provides an end-to-end workflow to preserve, collect, review, analyze, and export content that's responsive to your organization's internal and external legal investigations. It also lets your legal team manage the entire legal hold notification workflow to communicate with custodians involved in a case. Assigning a reviewer as a custodian in an Advanced eDiscovery case created from an insider risk management case helps your legal team take appropriate action and manage content preservation. To learn more about Advanced eDiscovery cases, see Overview of Advanced eDiscovery in Microsoft 365.

To escalate a case to a user investigation:

1. In the Microsoft 365 compliance center, go to **Insider risk management** and select the **Cases** tab.
2. Select a case, then select the **Escalate for investigation** button on the case action toolbar.
3. On the **Escalate for investigation** dialog, enter a name for the new user investigation. If needed, enter notes about the case and select **Escalate**.
4. Review the notice fields and update as appropriate. The values entered here will override the values on the template.
5. Select **Confirm** to create the user investigation case or select **Cancel** to close the dialog without creating a new user investigation case.

After the insider risk management case has been escalated to a new user investigation case, you can review the new case in the **eDiscovery** > **Advanced** area in the Microsoft 365 compliance center.

Run automated tasks with Power Automate flows for the case

Using recommended Power Automate flows, risk investigators and analysts can quickly take action to:

- Request information from HR or business about a user in an insider risk case
- Notify manager when a user has an insider risk alert
- Add calendar reminder to follow up on an insider risk case
- Create a record for an insider risk management case in ServiceNow

To run, manage, or create Power Automate flows for an insider risk management case:

1. Select **Automate** on the case action toolbar.
2. Choose the Power Automate flow to run, then select **Run flow**.
3. After the flow has completed, select **Done**.

To learn more about Power Automate flows for insider risk management, see Getting started with insider risk management settings.

View or create a Microsoft Teams team for the case

When Microsoft Teams integration for insider risk management is enabled in settings, a Microsoft Teams team is automatically created every time an alert is confirmed and a case is created. Risk investigators and analysts can quickly open Microsoft Teams and navigate directly to the team for a case by selecting **View Microsoft Teams team** on the case action toolbar.

For cases opened before enabling Microsoft Team integration, risk investigators and analysts can create a new Microsoft Teams team for a case by selecting **Create Microsoft Teams team** on the case action toolbar.

When a case is resolved, the associated Microsoft Team will be automatically archive (hidden and turned to read-only).

To learn more about Microsoft Teams for insider risk management, see Getting started with insider risk management settings.

Share the case

Sharing an insider risk management case allows risk investigators and analysts to easily collaborate with other compliance stakeholders in your organization. You can quickly share a link to an insider risk management case with external stakeholders from the case management area. To access the insider risk management case from the link, stakeholders must be included in any of the insider risk management role groups.

> **NOTE**
> Thank you for your feedback and support during the preview of the ServiceNow connector. We've decided to end the preview of ServiceNow connector and discontinue support in insider risk management on November 30, 2020. We are actively evaluating alternative methods to provide customers with ServiceNow integration in insider risk management.

The following sharing options are available:

- **ServiceNow**: After configuring the Microsoft 365 ServiceNow connector for your Microsoft 365 organization, you can easily share a link to the case, open an incident, or request a change with your ServiceNow organization. To share the case with ServiceNow, select **Share** > **ServiceNow** from the case action. ServiceNow integration with insider risk management supports includes the following case information and actions:
  - **Task name**: The name for the new ServiceNow task.
  - **Task description**: The description for the new ServiceNow task. This editable description field automatically includes a link to the insider risk management case.
  - **Task type**: The task type for the new ServiceNow task, either *Incident* or *Change request*.
  - **Priority**: The priority for the new ServiceNow task, either *Planning*, *Low*, *Moderate*, *High*, or *Critical*.
  - **Due date**: The requested date for completing the ServiceNow task.

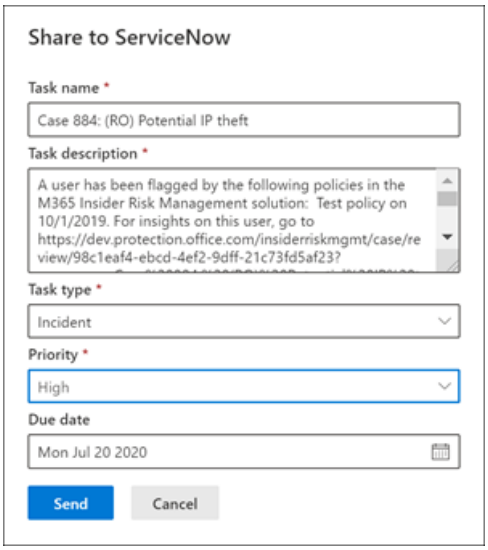

- **Email**: Shares a link to the insider risk management case in an email. You can choose any locally configured email client with this sharing option. To share the case link with email, select **Share** > **Email** from the case action toolbar.
- **Copy link**: Copies a link to the insider risk management case to your clipboard. To copy the case link to your clipboard, select **Share** > **Copy link** from the case action toolbar.

Resolve the case

After risk analysts and investigators have completed their review and investigation, a case can be resolved to act on all the alerts currently included in the case. Resolving a case adds a resolution classification, changes the case status to *Closed*, and the resolution action reasons are automatically added to the case notes queue on the **Case notes** dashboard. Cases are resolved as either:

- **Benign**: The classification for cases where policy match alerts are evaluated as low risk, non-serious, or false positive.
- **Confirmed policy violation**: The classification for cases where policy match alerts are evaluated as risky, serious, or the result of malicious intent.

To resolve a case:

1. In the Microsoft 365 compliance center, go to **Insider risk management** and select the **Cases** tab.
2. Select a case, then select the **Resolve case** button on the case action toolbar.
3. On the **Resolve case** dialog, select the **Resolve as** dropdown control to select the resolution classification for the case. The options are **Benign** or **Confirmed policy violation**.
4. On the **Resolve case** dialog, enter the reasons for the resolution classification in the **Action taken** text field.
5. Select **Resolve** to close the case or select **Cancel** close the dialog without resolving the case.

Insider risk management Content Explorer

2/18/2021 • 11 minutes to read • Edit Online

The insider risk management Content Explorer allows users assigned the *Insider Risk Management Investigators* role to examine the context and details of content associated with activity in alerts. For all alerts that are confirmed to a case, copies of data and message files are archived as a snapshot in time of the items, while maintaining the original files and messages in the storage sources. The copying of data and messages is transparent to the user associated with the alert and to the owner of the content. If the content includes Information Rights management permissions, these permissions are maintained for the copied content and users assigned the *Insider Risk Management Investigators* role will need these permissions and rights if they need to open and view the files. Each file and message are automatically assigned a unique file ID in the insider risk management case for management purposes. Documents associated with device indicator activities are not included in Content Explorer.

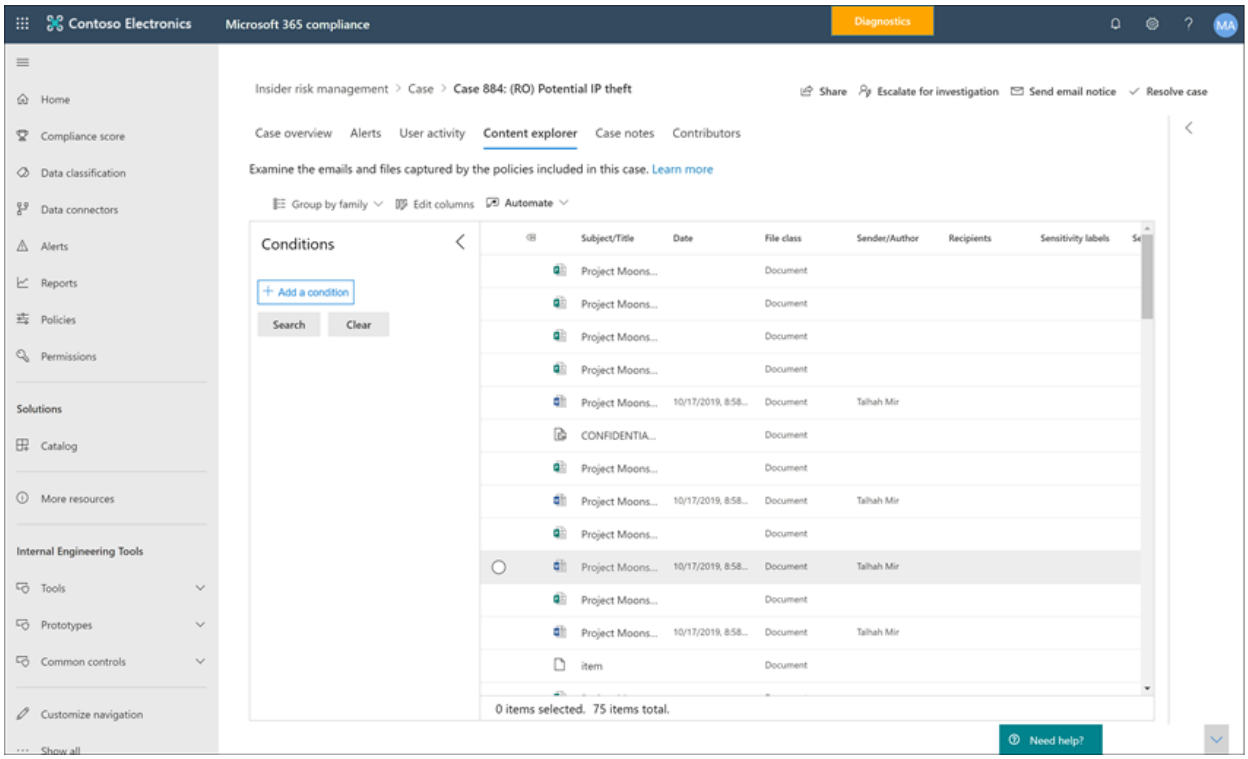

Column options

To make it easier for risk analysts and investigators to review captured data and messages and review the context to the case, several filtering and sorting tools are included in the Content Explorer. For basic sorting, the **Date** and **File class** columns support sorting using the column titles in the content queue pane. Other queue columns are available to add to the view to provide different pivots on the files and messages.

To add or remove column headings for the content queue, use the **Edit columns** control and select from the following column options. These columns map to the common, email, and document property conditions supported in the Content Explorer and listed later in this article.

COLUMN OPTION	DESCRIPTION

COLUMN OPTION	DESCRIPTION
Author	The author field from Office documents, which persists if a document is copied. For example, if a user creates a document and the emails it to someone else who then uploads it to SharePoint, the document will still retain the original author.
Bcc	Available for email messages, the users in the Bcc message field.
Cc	Available for email messages, the users in the Cc message field.
Compound path	Human readable path that describes the source of the item.
Conversation ID	Conversation Id from the message.
Conversation index	Conversation index from the message.
Created time	The time the file or email message was created.
Date	For email, the date a message was received by a recipient or sent by the sender. For documents, the date a document was last modified.
Dominant theme	Dominant theme as calculated for analytics.
Email set ID	Group ID for all messages in the same email set.
Family ID	Family Id groups together all items; for email, this column includes the message and all attachments; for documents, this column includes the document and any embedded items.
File class	For content from SharePoint and OneDrive: **Document**; for content from Exchange: **Email or **Attachment**.
File ID	Document identifier unique within the case.
File type icon	The extension of a file; for example, docx, one, pptx, or xlsx. This field is the same property as the FileExtension site property.
ID	The GUID identifier for the file.
Immutable ID	Immutable Id as stored in Office 365.
Inclusive type	Inclusive type calculated for analytics: **0** - not inclusive; **1** - inclusive; **2** - inclusive minus; **3** - inclusive copy.
Last modified	The date that a document was last changed.
Marked as representative	One document from each set of exact duplicates is marked as representatives.

COLUMN OPTION	DESCRIPTION
Message kind	The type of email message to search for. Possible values: contacts, docs, email, external data, faxes, im, journals, meetings, microsoft teams (returns items from chats, meetings, and calls in Microsoft Teams), notes, posts, rssfeeds, tasks, voicemail
Participants	List of all participants of a message; for example, Sender, To, Cc, Bcc.
Pivot ID	The ID of a pivot.
Received	The date that an email message was received by a recipient. This field is the same property as the Received email property.
Recipients	All recipient fields in an email message. These fields are To, Cc, and Bcc.
Representative ID	Numeric identifier of each set of exact duplicates.
Sender	The sender of an email message.
Sender/Author	For email, the person who sent a message. For documents, the person cited in the author field from Office documents. You can type more than one name, separated by commas. Two or more values are logically connected by the OR operator.
Sent	The date that an email message was sent by the sender. This field is the same property as the Sent email property.
Size	For both email and documents, the size of the item (in bytes).
Subject	The text in the subject line of an email message.
Subject/Title	For email, the text in the subject line of a message. For documents, the title of the document. As previously explained, the Title property is metadata specified in Microsoft Office documents. You can type the name of more than one subject/title, separated by commas. Two or more values are logically connected by the OR operator.
Themes list	Themes list as calculated for analytics.
Title	The title of the document. The Title property is metadata that's specified in Office documents. It's different than the file name of the document.
To	The recipient of an email message in the To field.

Advanced search conditions

You can add search conditions to narrow the scope of a search and return a more refined set of results. Each

condition adds a clause to the search query that is created and run when you start the search. A condition is logically connected to the keyword query (specified in the keyword box) by a logical operator (which is represented as c:c) that is similar in functionality to the AND operator. That means that items have to satisfy both the keyword query and one or more conditions to be included in the search results. This functionality is how conditions help to narrow your results.

For advanced filter and search tools, expand the **Filter** pane on the left side of the content queue. Select the **Add a condition** button to open the condition list:

Operators used with conditions

OPERATOR	QUERY EQUIVALENT	DESCRIPTION
After	`property>date`	Used with date conditions. Returns items that were sent, received, or modified after the specified date.
Before	`property<date`	Used with date conditions. Returns items that were sent, received, or modified before the specified date.
Between	`date..date`	Use with date and size conditions. When used with a date condition, returns items there were sent, received, or modified within the specified date range. When used with a size condition, returns items whose size is within the specified range.
Contains all of	`(property:value) OR (property:value)`	Used with conditions for properties that specify a string value. Returns items that contain all of one or more specified string values.
Contains any of	`(property:value) OR (property:value)`	Used with conditions for properties that specify a string value. Returns items that contain any part of one or more specified string values.
Contains none of	`-property:value` `NOT property:value`	Used with conditions for properties that specify a string value. Returns items that don't contain any part of the specified string value.
Doesn't equal any of	`-property=value` `NOT property=value`	Used with conditions for properties that specify a string value. Returns items that don't contain the specific string.
Equals	`size=value`	Returns items that are equal to the specified size.[1]
Equals any of	`(property=value) OR (property=value)`	Used with conditions for properties that specify a string value. Returns items that are an exact match of one or more specified string values.

OPERATOR	QUERY EQUIVALENT	DESCRIPTION
Equals none of	`(property=value) OR (property=value)`	Used with conditions for properties that specify a string value. Returns items that do not match one or more specified string values.
Greater than	`size>value`	Returns items where the specified property is greater than the specified value.[1]
Greater or equal	`size>=value`	Returns items where the specified property is greater than or equal to the specified value.[1]
Less than	`size<value`	Returns items that are greater than or equal to the specific value.[1]
Less or equal	`size<=value`	Returns items that are greater than or equal to the specific value.[1]
Not equal	`size<>value`	Returns items that don't equal the specified size.[1]

NOTE

[1] This operator is available only for conditions that use the Size property.

Common property conditions

CONDITION OPTION	DESCRIPTION
Date	For email, the date a message was received by a recipient or sent by the sender. For documents, the date a document was last modified.
Sender/Author	For email, the person who sent a message. For documents, the person cited in the author field from Office documents. You can type more than one name, separated by commas. Two or more values are logically connected by the **OR** operator.
Size	For both email and documents, the size of the item (in bytes).
Subject/Title	For email, the text in the subject line of a message. For documents, the title of the document. The Title property in documents is metadata specified in Microsoft Office documents. You can type the name of more than one subject/title, separated by commas. Two or more values are logically connected by the OR operator.

Email property conditions

The following table lists email message property conditions available the Content Explorer.

CONDITION OPTION	DESCRIPTION
Bcc	The Bcc field of an email message.
Cc	The Cc field of an email message.
Email security	Security setting of the message.
Email sensitivity	Sensitivity setting of the message.
Email set ID	Group ID for all messages in the same email set.
From	The sender of an email message.
Has attachment	Indicates whether a message has an attachment. Use the values **true** or **false**.
Importance	The importance of an email message, which a sender can specify when sending a message. By default, messages are sent with normal importance, unless the sender sets the importance as **high** or **low**.
Meeting end date	Meeting end date for meetings.
Meeting start date	Meeting start date for meetings.
Message kind	The type of email message to search for. Possible values: contacts, docs, email, external data, faxes, im, journals, meetings, microsoft teams (returns items from chats, meetings, and calls in Microsoft Teams), notes, posts, rssfeeds, tasks, voicemail
Participant domain	List of all domains of participants of a message.
Participants	All the people fields in an email message. These fields are From, To, Cc, and Bcc.
Received	The date that an email message was received by a recipient.
Recipient domains	List of all domains of recipients of a message.
Sender	Sender (From) field for message types. Format is **DisplayName <SmtpAddress>**.
Sender domain	Domain of the sender.
Subject	The text in the subject line of an email message. **Note:** When you use the Subject property in a query, the search returns all messages in which the subject line contains the text you're searching for. In other words, the query doesn't return only those messages that have an exact match. For example, if you search for `subject:"Quarterly Financials"`, your results will include messages with the subject "Quarterly Financials 2018".

CONDITION OPTION	DESCRIPTION
To	The To field of an email message.
Unique in email set	False if there's a duplicate of the attachment in its email set.

Document property conditions

The following table lists documents property conditions available the Content Explorer. Many of these property conditions are shared with review sets included in Advanced eDiscovery cases.

CONDITION OPTION	DESCRIPTION
Attorney-client privilege score	Attorney-client privilege model content score.
Author	The author field from Office documents, which persists if a document is copied. For example, if a user creates a document and the emails it to someone else who then uploads it to SharePoint, the document will still retain the original author.
Compliance labels	Compliance labels applied in Office 365.
Compound path	Human readable path that describes the source of the item.
Conversation ID	Conversation Id from the message.
Created time	The time the file or email message was created.
Custodian	Name of the custodian the item was associated with.
Dominant theme	Dominant theme as calculated for analytics.
Family ID	Family Id groups together all items; for email, this field includes the message and all attachments; for documents, this field includes the document and any embedded items.
File class	For content from SharePoint and OneDrive: **Document**; for content from Exchange: **Email or **Attachment**.
File types	The extension of a file; for example, docx, one, pptx, or xlsx.
Has attorney participant	True when at least one of the participants is found in the attorney list; otherwise, the value is False.
Immutable ID	Immutable Id as stored in Office 365.
Inclusive type	Inclusive type calculated for analytics: **0** - not inclusive; **1** - inclusive; **2** - inclusive minus; **3** - inclusive copy.
Item class	Item class supplied by exchange server; for example, **IPM.Note**
Last modified	The date that a document was last changed.

CONDITION OPTION	DESCRIPTION
Load ID	Load Id, in which the item was loaded into a review set.
Location name	String that identifies the source of the item. For exchange, this field will be the SMTP address of the mailbox. For SharePoint and OneDrive, the URL to the site collection.
Marked as representative	One document from each set of exact duplicates is marked as representatives.
Native file extension	Native extension of the item.
Native file name	Native file name of the item.
NdEtSortExclAttach	Concatenation of email set and ND set for efficient sorting at review time; D is added as a prefix to ND sets and E is added to email sets.
Pivot ID	The ID of a pivot.
Potentially privileged	True if attorney-client privilege detection model considers the document potentially privileged.
Processing status	Processing status after the item was added to a review set.
Read percentile	Read percentile for the document based on Relevance.
Relevance score	Relevance score of a document based on Relevance.
Relevance tag	Relevance score of a document based on Relevance.
Representative ID	Numeric identifier of each set of exact duplicates.
Tags	Tags applied in a review set.
Themes list	Themes list as calculated for analytics.
Title	The title of the document. The Title property is metadata that's specified in Office documents. It's different than the file name of the document.
Was remediated	True if the item was remediated, otherwise False.
Word count	The number of words in a file.

Insider risk management Users dashboard

11/2/2020 • 8 minutes to read • Edit Online

The **Users dashboard** is an important tool in the insider risk management workflow and helps investigators and analysts have a more complete understanding of risk activities. This dashboard offers views and management features to meet administrative needs between the creating insider risk management policies and managing insider risk management cases.

After users are added to insider risk management policies, background processes are automatically evaluating user activities for triggering indicators. After triggering indicators are present, user activities are assigned risk scores. Some of these activities may result in an insider risk alert, but some activities may not meet a minimum risk score level and an insider risk alert won't be created. The **Users dashboard** allows you to view users with these types of indicators and risk scores, as well users that have active insider risk alerts.

Additionally, there may be scenarios where you need to add temporarily users to insider risk policies after an unusual event is reported outside of the insider risk management workflow. The **Users dashboard** allows you to manually add a user to an insider risk policy for a specific amount of time and bypass the requirement for a user to have a triggering indicator. These users are always displayed in the Users dashboard when actively assigned to a policy.

Learn more about how the Users dashboard displays users in the following scenarios:

- Dashboard users with active insider risk policy alerts
- Dashboard users with triggering indicators
- Dashboard users added temporarily to policies

Dashboard users with active insider risk policy alerts

The **Users dashboard** automatically displays all users with active insider risk policy alerts. These users with alerts have both a triggering indicator and an activity risk score that meets the requirements for creating an insider risk alert. Activities for these users are viewed by selecting the user in the **Users dashboard** and navigating to the **User activity** tab.

Dashboard users with triggering indicators

The **Users dashboard** automatically displays all users with triggering indicators, but that don't have an activity risk score that would create an insider risk alert. For example, a user with a reported resignation date is displayed because this event is a triggering indicator but isn't an activity that has a risk score. Activities for these users are viewed by selecting the user in the **Users dashboard** and navigating to the **User activity** tab.

Dashboard users added temporarily to policies

The **Users dashboard** allows you to temporarily add users to an existing insider risk management policy after an unusual event outside of the insider risk management workflow. Temporarily adding users is also a way to add users to an insider risk management policy for testing the policy, even if a required connector isn't configured.

When a user is manually added to a policy, the user activities for the previous 90 days are scored and added to the **User activity** timeline. For example, you have a user not currently in-scope in an insider risk policy and the user has data leak activities reported to the legal department in your organization. The legal department recommends that you configure new short-term monitoring requirements for the user. You can temporarily

assign the user to your *Data leaks* policy for a designated length of time (activation window). All users added temporarily are displayed in the **Users dashboard** because triggering indicator requirements are waived.

> **NOTE**
>
> It may take several hours for new manually-added users to appear in the **Users dashboard**. Activities for the previous 90 days for these users may take up to 24 hours to display. To view activities for manually added users, select the user on the **Users dashboard** and open the **User activity** tab on the details pane.

The user is automatically removed from the insider policy and the **Users dashboard** when the time defined in the **Activation window** expires if:

- the user doesn't have any triggering indicators or insider risk policy alerts, and
- if the manually defined **Activation window** duration is longer than the global policy **Activation window** duration.

The **Activation window** setting with the longest duration always overrides the **Activation window** setting with a shorter duration. For example, you've configured the **Activation window** on the global **Policy timeframes** tab in the insider risk management global settings for 15 days, which is automatically applied to all your insider risk policies.

You temporarily add a user to your *Data leaks* insider risk policy and define 30 days as the **Activation window** for this user. The global **Activation window** setting of 15 days is overridden by defining the **Activation window** setting of 30 days for the temporarily added user. The temporarily added user will remain in the **Users dashboard** and be in-scope for the policy for 30 days.

In the opposite scenario where the global **Activation window** setting is longer than the **Activation window** setting defined for a temporarily added user, the global **Activation window** setting would override the **Activation window** setting for the temporarily added user. The temporarily added user will remain in the **Users dashboard** and be in-scope for the policy for the number of days defined in the global **Activation window** settings.

View user information on the Users dashboard

Each user displayed in the **Users dashboard** has the following information:

- **Users**: The username for a user. This field is anonymized if the global anonymization setting for insider risk management is enabled.
- **Risk level**: The current calculated risk level of the user. This score is calculated every 24 hours and uses the alert risk scores from all active alerts associated to the user. For users with only triggering indicators, the risk level is zero.
- **Active alerts**: The number of active alerts for all policies.
- **Confirmed violations**: The number of cases resolved as *confirmed policy violation* for the user.
- **Case**: The current active case for the user.

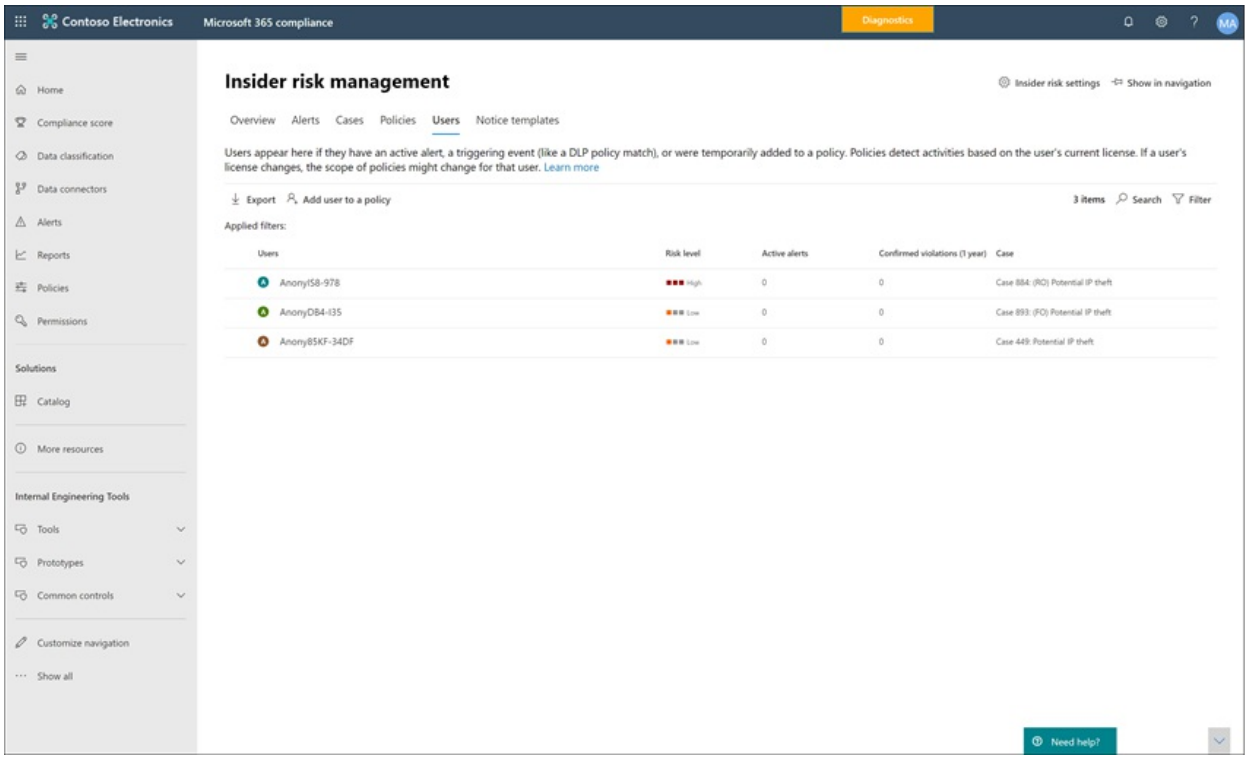

NOTE

The number of users displayed on the **Users dashboard** may be limited in some instances, depending on the volume of active alerts and matching policies. Users with active alerts are displayed on the **Users dashboard** as the alerts are generated, and there may be rare cases when the maximum number of displayed users is reached. If this happens, users with active alerts who aren't displayed will be added to the **Users dashboard** as existing user alerts are triaged.

View user details

To view more details about risk activity for a user, open the user details pane by double-clicking a user in the **Users dashboard**. On the details pane, you can view the following information:

- **User profile** tab
  - **Name and title**: The name and position title for the user from Azure Active Directory. These user fields will be anonymized or empty if the global anonymization setting for insider risk management is enabled.
  - **User email**: The email address for the user.
  - **Alias**: The network alias for the user.
  - **Organization or department**: The organization or department for the user.
- **User activity** tab
  - **History of recent user activity**: Lists both triggering indicators and insider risk indicators for user activities up to the last 180 days. All activities pertinent to insider risk indicators are also scored, though the activities may or may not have generated an insider risk alert. Triggering indicator examples may be a resignation date or the last scheduled date of work for the user. Insider risk indicators are activities determined to have an element of risk and are defined in policies that the user is included in. Event and risk activities are listed with the most recent item listed first.

Temporarily add a user to a policy

To temporarily add a user to an insider risk management policy, you'll use the **Users** tab in the **Insider risk management** solution in the Microsoft 365 compliance center. Users added manually bypass triggering

indicator requirements for the policy they are added to and are displayed in the **Users dashboard**. To permanently add a user to an insider risk management policy, you'll use the policy wizard.

Complete the following steps to add a user to an existing insider risk policy:

1. In the Microsoft 365 compliance center, go to **Insider risk management** and select the **Users** tab.
2. Select **Add a user to a policy** on the toolbar.
3. On the **Add a new user** dialog, start typing a user name in the **User** field. Select the user you want to add to a policy.
4. Select the dropdown arrow for the **Policy** field to display configured insider risk management policies. Select the policy to add the user to.
5. Use the **Activation window** slider control to define how long the user is included in a policy and displayed in the Users dashboard. The time you specify determines how long the policy is active for this user and starts when the first alert is generated or a triggering indicator (like a DLP policy match) is detected. The range for the **Activation window** is 5 to 30 days.
6. Select **Add** and then **Confirm** to add the user to the policy.

> **NOTE**
>
> It may take several hours for new manually-added users to appear in the **Users dashboard**. Activities for the previous 90 days for these users may take up to 24 hours to display. To view activities for manually added users, select the user on the **Users dashboard** and open the **User activity** tab on the details pane.

Run automated tasks with Power Automate flows for a user

Using recommended Power Automate flows, risk investigators and analysts can quickly take action to:

- Notify users when they're added to an insider risk policy

To run, manage, or create Power Automate flows for an insider risk management user:

1. Select **Automate** on the user action toolbar.
2. Choose the Power Automate flow to run, then select **Run flow**.
3. After the flow has completed, select **Done**.

To learn more about Power Automate flows for insider risk management, see Getting started with insider risk management settings.

Insider risk management notice templates

11/2/2020 • 4 minutes to read • Edit Online

Insider risk management notice templates allow you to send email messages to users when their activities generate a policy match and alert. In most cases, user actions that generate alerts are the result of mistakes or inadvertent activities without ill intent. Notices serve as simple reminders to users to be more careful, to provide links to information for refresher training, or to corporate policy resources. Notices can be an important part of your internal compliance training program and can help create a documented audit trail for users with recurring risk activities.

Create notice templates if you want to send users an email reminder notice for policy matches as part of the issue resolution process. Notices can only be sent to the user email address associated with the specific alert being reviewed. When selecting a notice template to apply to a policy match, you can choose to accept the field values defined in the template or overwrite the fields as needed.

Notice templates dashboard

The **Notices templates dashboard** displays a list of configured notice templates and allows you to create new notice templates. The notice templates are listed in reverse date order with the most recent notice template listed first.

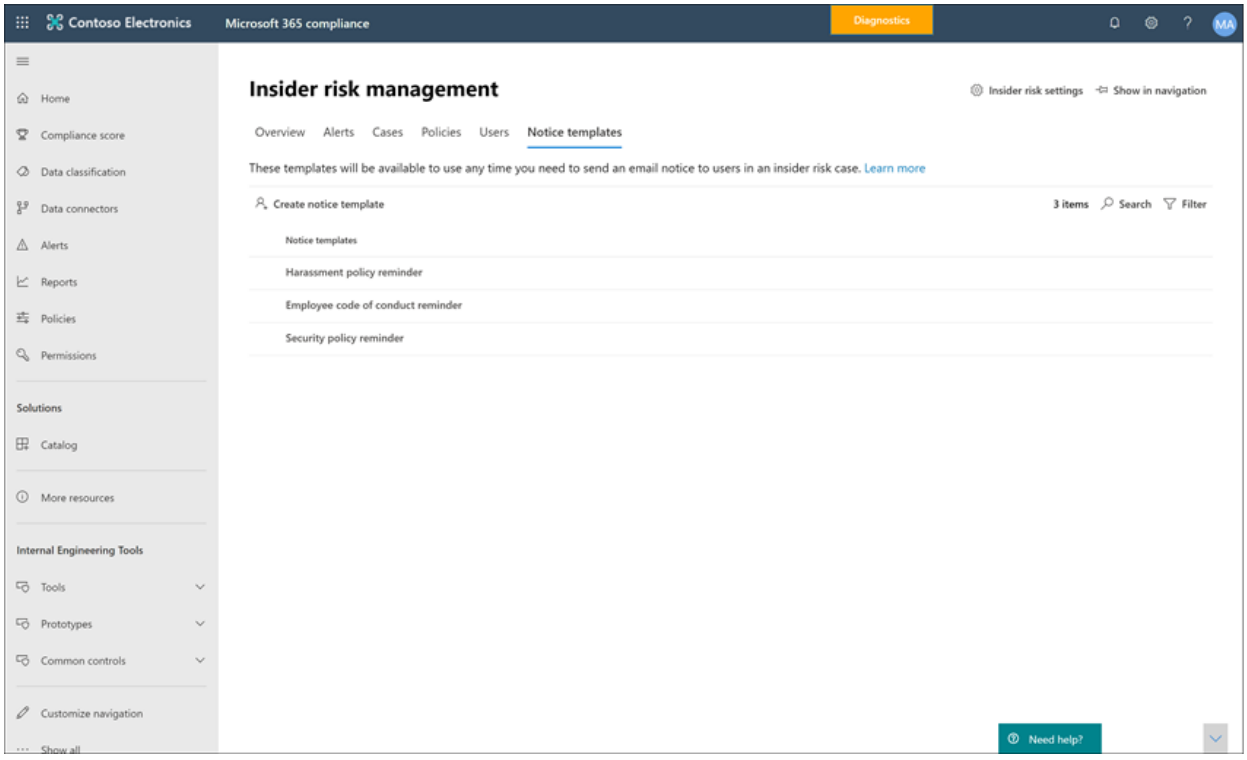

HTML for notices

If you'd like to create more than a simple text-based email message for notifications, you can create a more detailed message by using HTML in the message body field of a notice template. The following example provides the message body format for a basic HTML-based email notification template:

```
<!DOCTYPE html>
<html>
<body>
<h2>Action Required: Contoso User Code of Conduct Policy Training</h2>
<p>A recent activity you've performed has generated a risk alert prohibited by the Contoso User <a
href='https://www.contoso.com'>Code of Conduct Policy</a>.</p>
<p>You are required to attend the Contoso User Code of Conduct <a
href='https://www.contoso.com'>training</a> within the next 14 days. Please contact <a
href='mailto:hr@contoso.com'>Human Resources</a> with any questions about this training request.</p>
<p>Thank you,</p>
<p><em>Human Resources</em></p>
</body>
</html>
```

NOTE

HTML href attribute implementation in the insider risk management notice templates currently support only single quotation marks instead of double quotation marks for URL references.

Create a new notice template

To create a new insider risk management notice template, you'll use the notice wizard in **Insider risk management** solution in the Microsoft 365 compliance center.

Complete the following steps to create a new insider risk management notice template:

1. In the Microsoft 365 compliance center, go to **Insider risk management** and select the **Notice templates** tab.
2. Select **Create notice template** to open the notice wizard.
3. On the **Create a new notice template** page, complete the following fields:
   - **Template name**: Enter a friendly name for the notice. This name appears on the list of notices on the notice dashboard and in the notice selection list when sending notices from a case.
   - **Send from**: Enter the sender email address for the notice. This address will appear in the **From:** field in all notices sent to users unless changed when sending a notice from a case.
   - **Cc and Bcc** fields: Optional users or groups to be notified of the policy match, selected from the Active Directory for your subscription.
   - **Subject**: Information that appears in the subject line of the message, supports text characters.
   - **Message body**: Information that appears in the message body, supports text or HTML values.
4. Select **Create** to create and save the notice template or select **Cancel** to close without saving the notice template.

Update a notice template

To update an existing insider risk management notice template, complete the following steps:

1. In the Microsoft 365 compliance center, go to **Insider risk management** and select the **Notice templates** tab.
2. On the notice dashboard, select the notice template you want to manage.
3. On the notice details page, select **Edit**
4. On the **Edit** page, you can edit the following fields:
   - **Template name**: Enter a new friendly name for the notice. This name appears on the list of notices on the notice dashboard and in the notice selection list when sending notices from a case.
   - **Send from**: Update the sender email address for the notice. This address will appear in the **From:**

field in all notices sent to users unless changed when sending a notice from a case.

- **Cc and Bcc** fields: Update optional users or groups to be notified of the policy match, selected from the Active Directory for your subscription.
- **Subject**: Update information that appears in the subject line of the message, supports text characters.
- **Message body**: Update information that appears in the message body, supports text or HTML values.

5. Select **Save** to update and save the notice or select **Cancel** to close without saving the notice template.

Delete a notice template

To delete an existing insider risk management notice template, complete the following steps:

1. In the Microsoft 365 compliance center, go to **Insider risk management** and select the **Notice templates** tab.
2. On the notice dashboard, select the notice template you want to delete.
3. Select the **Delete** icon on the toolbar.
4. To delete the notice template, select **Yes** in the delete dialog. To cancel the deletion, select **Cancel**.

Information barriers in Microsoft 365

2/18/2021 • 2 minutes to read • Edit Online

Microsoft 365 enables communication and collaboration across groups and organizations and supports ways to restrict communication and collaboration among specific groups of users when necessary. This may include situations or scenarios where you want to restrict communication and collaboration between two groups to avoid a conflict of interest from occurring in your organization. This may also include situations when you need to restrict communication and collaboration between certain people inside your organization to safeguard internal information.

Information barriers are supported in Microsoft Teams, SharePoint Online, and OneDrive for Business. A compliance administrator or information barriers administrator can define policies to allow or prevent communications between groups of users in Microsoft Teams. Information barrier policies can be used for situations like these:

- User in the day trader group should not communicate or share files with the marketing team
- Finance personnel working on confidential company information should not communicate or share files with certain groups within their organization
- An internal team with trade secret material should not call or chat online with people in certain groups within their organization
- A research team should only call or chat online with a product development team

Configure information barriers for Microsoft 365

Use the following steps to configure information barriers for your organization:

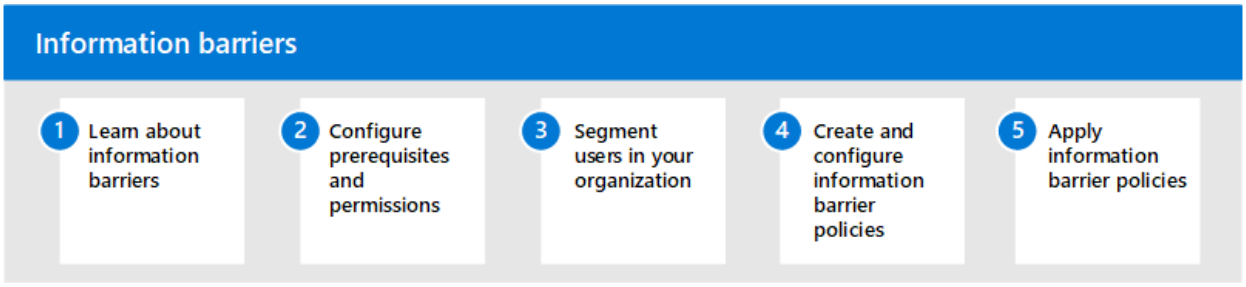

1. Learn about information barriers in Microsoft 365
2. Configure prerequisites and permissions
3. Segment users in your organization
4. Create and configure information barrier policies
5. Apply information barrier policies

More information about information barriers

- Attributes for information barrier policies
- Edit or remove information barrier policies

Learn about information barriers in Microsoft 365

2/18/2021 • 4 minutes to read • Edit Online

Microsoft cloud services include powerful communication and collaboration capabilities. But suppose that you want to restrict communication and collaboration between two groups to avoid a conflict of interest from occurring in your organization. Or, perhaps you want to restrict communication and collaboration between certain people inside your organization in order to safeguard internal information. Microsoft 365 enables communication and collaboration across groups and organizations, so is there a way to restrict communication and collaboration among specific groups of users when necessary? With information barriers, you can!

Microsoft Teams, SharePoint Online, and OneDrive for Business support information barriers. Assuming your subscription includes information barriers, a compliance administrator, or information barriers administrator can define policies to allow or prevent communications between groups of users in Microsoft Teams. Information barrier policies can be used for situations like these:

- User in the day trader group should not communicate or share files with the marketing team
- Finance personnel working on confidential company information should not communicate or share files with certain groups within their organization
- An internal team with trade secret material should not call or chat online with people in certain groups within their organization
- A research team should only call or chat online with a product development team
- A site for day trader group should not be shared or accessed by anyone outside the day trader group

> **IMPORTANT**
>
> Information barriers ***only supports*** two way restrictions. One way restrictions, such as marketing can communicate and collaborate with day traders, but day traders cannot communicate and collaborate with marketing ***is not supported***.

For all of these example scenarios (and more), information barrier policies can be defined to prevent or allow communications and collaboration in Microsoft Teams, SharePoint Online and OneDrive. Such policies can prevent people from calling or chatting with those they shouldn't, or enable people to communicate only with specific groups in Microsoft Teams. With information barrier policies in effect, whenever users who are covered by those policies attempt to communicate and collaborate with others in Microsoft Teams, SharePoint Online or OneDrive checks are done to prevent (or allow) communication and collaboration (as defined by information barrier policies).

To learn more about the user experience with information barriers, see:

- Information barriers in Microsoft Teams
- Information barriers in SharePoint Online
- Information barriers in OneDrive

> **IMPORTANT**
>
> Currently, information barriers do not apply to email communications. In addition, information barriers are independent from compliance boundaries.
>
> Before you define and apply information barrier policies, make sure your organization does not have Exchange address book policies in effect. (Information barriers are based on address book policies.)

What happens with information barriers

When information barrier policies are in place, people who should not communicate or share files with other specific users won't be able to find, select, chat, or call those users. With information barriers, checks are in place to prevent unauthorized communication and collaboration.

Information barriers applies to Microsoft Teams (chats and channels), SharePoint Online and OneDrive. In Microsoft Teams, information barrier policies determine and prevent the following kinds of unauthorized communications:

- Searching for a user
- Adding a member to a team
- Starting a chat session with someone
- Starting a group chat
- Inviting someone to join a meeting
- Sharing a screen
- Placing a call
- Sharing a file with another user
- Access to file through sharing link

If the people involved are included in an information barrier policy to prevent the activity, they will not be able to proceed. In addition, potentially, everyone included in an information barrier policy can be blocked from communicating with others in Microsoft Teams. When people affected by information barrier policies are part of the same team or group chat, they might be removed from those chat sessions and further communication with the group might not be allowed.

To learn more about the user experience with information barriers, see information barriers in Microsoft Teams.

In SharePoint Online and OneDrive, information barrier policies determine and prevent the following kinds of unauthorized collaborations:

- Adding a member to a site
- Accessing site or content by a user
- Sharing site or content with another user
- Searching a site

To learn more about the user experience with information barriers, see information barriers in SharePoint Online

Required licenses and permissions

Information barriers are rolling out now, and are included in subscriptions, such as:

- Microsoft 365 E5/A5
- Office 365 E5/A5
- Office 365 Advanced Compliance
- Microsoft 365 Compliance E5/A5
- Microsoft 365 Insider Risk Management

For more information, see Microsoft 365 licensing guidance for security & compliance.

To define or edit information barrier policies, you must be assigned one of the following roles:

- Microsoft 365 global administrator
- Office 365 global administrator
- Compliance administrator

- IB Compliance Management

(To learn more about roles and permissions, see Permissions in the Office 365 Security & Compliance Center.)

You must be familiar with PowerShell cmdlets in order to define, validate, or edit information barrier policies. Although we provide several examples of PowerShell cmdlets in the how-to article, you'll need to know other details, such as parameters, for your organization.

Next steps

- Learn more about information barriers in Microsoft Teams
- Learn more about information barriers in SharePoint Online
- Learn more about information barriers in OneDrive
- See the attributes that can be used for information barrier policies
- Define policies for information barriers
- Edit (or remove) information barrier policies

Define information barrier policies

2/18/2021 • 16 minutes to read • Edit Online

With information barriers, you can define policies that are designed to prevent certain segments of users from communicating with each other, or allow specific segments to communicate only with certain other segments. Information barrier policies can help your organization maintain compliance with relevant industry standards and regulations, and avoid potential conflicts of interest. To learn more, see Information barriers.

This article describes how to plan, define, implement, and manage information barrier policies. Several steps are involved, and the work flow is divided into several parts. Make sure to read through the prerequisites and the entire process before you begin defining (or editing) information barrier policies.

> **TIP**
> This article includes an example scenario and a downloadable Excel workbook to help you plan and define your information barrier policies.

Concepts of information barrier policies

When you define policies for information barriers, you'll work with user account attributes, segments, "block" and/or "allow" policies, and policy application.

- User account attributes are defined in Azure Active Directory (or Exchange Online). These attributes can include department, job title, location, team name, and other job profile details.
- Segments are sets of users that are defined in the Security & Compliance Center using a selected **user account attribute**. (See the list of supported attributes.)
- Information barrier policies determine communication limits or restrictions. When you define information barrier policies, you choose from two kinds of policies:
  - "Block" policies prevent one segment from communicating with another segment.
  - "Allow" policies allow one segment to communicate with only certain other segments.
- Policy application is done after all information barrier policies are defined, and you are ready to apply them in your organization.

The work flow at a glance

PHASE	WHAT'S INVOLVED
Make sure prerequisites are met	- Verify that you have the required licenses and permissions - Verify that your directory includes data for segmenting users - Enable scoped directory search for Microsoft Teams - Make sure audit logging is turned on - Make sure no Exchange address book policies are in place - Use PowerShell (examples are provided) - Provide admin consent for Microsoft Teams (steps are included)

PHASE	WHAT'S INVOLVED
Part 1: Segment users in your organization	- Determine what policies are needed - Make a list of segments to define - Identify which attributes to use - Define segments in terms of policy filters
Part 2: Define information barrier policies	- Define your policies (do not apply yet) - Choose from two kinds (block or allow)
Part 3: Apply information barrier policies	- Set policies to active status - Run the policy application - View policy status
(As needed) Edit a segment or a policy	- Edit a segment - Edit or remove a policy - Rerun the policy application - View policy status
(As needed) Troubleshooting	- Take action when things are not working as expected

Prerequisites

In addition to the required licenses and permissions, make sure that the following requirements are met:

- Directory data - Make sure that your organization's structure is reflected in directory data. To take this action, make sure that user account attributes, such as group membership, department name, etc. are populated correctly in Azure Active Directory (or Exchange Online). To learn more, see the following resources:
  - Attributes for information barrier policies
  - Add or update a user's profile information using Azure Active Directory
  - Configure user account properties with Office 365 PowerShell
- Scoped directory search - Before you define your organization's first information barrier policy, you must enable scoped directory search in Microsoft Teams. Wait at least 24 hours after enabling scoped directory search before you set up or define information barrier policies.
- EXO license - IB policies work only if the target users have been assigned an EXO license.
- Audit logging - In order to look up the status of a policy application, audit logging must be turned on. We recommend you enable auditing before you begin to define segments or policies. To learn more, see Turn the audit log search on or off.
- No address book policies - Before you define and apply information barrier policies, make sure no Exchange address book policies are in place. Information barriers are based on address book policies, but the two kinds of policies are not compatible. If you do have such policies, make sure to remove your address book policies first. Once information barrier policies are enabled and you have hierarchical address book enabled, all users ***who are not included*** in an information barrier segment will see the hierarchical address book in Exchange online.
- PowerShell - Currently, information barrier policies are defined and managed in the Office 365 Security & Compliance Center using PowerShell cmdlets. Although several examples are provided in this article, you'll need to be familiar with PowerShell cmdlets and parameters. You will also need the Azure PowerShell module.
  - Connect to Security & Compliance Center PowerShell

- Install the Azure PowerShell module

- Admin consent for information barriers in Microsoft Teams - When your IB policies are in place, they can remove non-IB compliance users from Groups (i.e. Teams channels, which are based on groups). This configuration helps ensure your organization remains compliant with policies and regulations. Use the following procedure to enable information barrier policies to work as expected in Microsoft Teams.

  1. Run the following PowerShell cmdlets:

     ```
     Connect-AzAccount
     $appId="bcf62038-e005-436d-b970-2a472f8c1982"
     $sp=Get-AzADServicePrincipal -ServicePrincipalName $appId
     if ($sp -eq $null) { New-AzADServicePrincipal -ApplicationId $appId }
     Start-Process  "https://login.microsoftonline.com/common/adminconsent?client_id=$appId"
     ```

  2. When prompted, sign in using your work or school account for Office 365.

  3. In the **Permissions requested** dialog box, review the information, and then choose **Accept**.

When all the prerequisites are met, proceed to the next section.

> **TIP**
>
> To help you prepare your plan, an example scenario is included in this article. See Contoso's departments, segments, and policies.
>
> In addition, a downloadable Excel workbook is available to help you plan and define your segments and policies (and create your PowerShell cmdlets). Get the workbook.

Part 1: Segment users

During this phase, you determine what information barrier policies are needed, make a list of segments to define, and then define your segments.

Determine what policies are needed

Considering legal and industry regulations, who are the groups within your organization who will need information barrier policies? Make a list. Are there any groups who should be prevented from communicating with another group? Are there any groups that should be allowed to communicate only with one or two other groups? Think about the policies you need as belonging to one of two groups:

- "Block" policies prevent one group from communicating with another group.
- "Allow" policies allow a group to communicate with only certain other, specific groups.

When you have your initial list of groups and policies, proceed to identify the segments you'll need.

Identify segments

In addition to your initial list of policies, make a list of segments for your organization. Users who will be included in information barrier policies should belong to a segment. Plan your segments carefully as a user can only be in one segment. Each segment can have only one information barrier policy applied.

> **IMPORTANT**
>
> A user can only be in one segment.

Determine which attributes in your organization's directory data you'll use to define segments. You can use *Department*, *MemberOf*, or any of the supported attributes. Make sure that you have values in the attribute you

select for users. See the list of supported attributes for information barriers.

> **IMPORTANT**
>
> **Before you proceed to the next section, make sure your directory data has values for attributes that you can use to define segments.** If your directory data does not have values for the attributes you want to use, then the user accounts must be updated to include that information before you proceed with information barriers. To get help with this, see the following resources:
>
> - Configure user account properties with Office 365 PowerShell
> - Add or update a user's profile information using Azure Active Directory

Define segments using PowerShell

Defining segments does not affect users; it just sets the stage for information barrier policies to be defined and then applied.

1. Use the **New-OrganizationSegment** cmdlet with the **UserGroupFilter** parameter that corresponds to the attribute you want to use.

SYNTAX	EXAMPLE
`New-OrganizationSegment -Name "segmentname" -UserGroupFilter "attribute -eq 'attributevalue'"`	`New-OrganizationSegment -Name "HR" -UserGroupFilter "Department -eq 'HR'"` In this example, a segment called *HR* is defined using *HR*, a value in the *Department* attribute. The **-eq** portion of the cmdlet refers to "equals." (Alternately, you can use **-ne** to mean "not equals". See Using "equals" and "not equals" in segment definitions.)

   After you run each cmdlet, you should see a list of details about the new segment. Details include the segment's type, who created or last modified it, and so on.

2. Repeat this process for each segment you want to define.

> **IMPORTANT**
>
> **Make sure that your segments do not overlap.** Each user who will be affected by information barriers should belong to one (and only one) segment. No user should belong to two or more segments. (See Example: Contoso's defined segments in this article.)

After you have defined your segments, proceed to define information barrier policies.

Using "equals" and "not equals" in segment definitions

In the following example, we are defining a segment such that "Department equals HR."

EXAMPLE	NOTE
`New-OrganizationSegment -Name "HR" -UserGroupFilter "Department -eq 'HR'"`	Notice that in this example, the segment definition includes an "equals" parameter denoted as **-eq**.

You can also define segments using a "not equals" parameter, denoted as **-ne**, as shown in the following table:

SYNTAX	EXAMPLE

SYNTAX	EXAMPLE
`New-OrganizationSegment -Name "NotSales" -UserGroupFilter "Department -ne 'Sales'"`	In this example, we defined a segment called *NotSales* that includes everyone who is not in *Sales*. The **-ne** portion of the cmdlet refers to "not equals".

In addition to defining segments using "equals" or "not equals", you can define a segment using both "equals" and "not equals" parameters. You can also define complex group filters using logical *AND* and *OR* operators.

SYNTAX	EXAMPLE
`New-OrganizationSegment -Name "LocalFTE" -UserGroupFilter "Location -eq 'Local'" -and "Position -ne 'Temporary'"`	In this example, we defined a segment called *LocalFTE* that includes people who are located locally and whose positions are not listed as *Temporary*.
`New-OrganizationSegment -Name "Segment1" -UserGroupFilter "MemberOf -eq 'group1@contoso.com'' -and MemberOf -ne 'group3@contoso.com'"`	In this example, we defined a segment called *Segment1* that includes people who are members of group1@contoso.com and not members of group3@contoso.com.
`New-OrganizationSegment -Name "Segment2" -UserGroupFilter "MemberOf -eq 'group2@contoso.com' -or MemberOf -ne 'group3@contoso.com'"`	In this example, we defined a segment called *Segment2* that includes people who are members of group2@contoso.com and not members of group3@contoso.com.
`New-OrganizationSegment -Name "Segment1and2" -UserGroupFilter "(MemberOf -eq 'group1@contoso.com' -or MemberOf -eq 'group2@contoso.com') -and MemberOf -ne 'group3@contoso.com'"`	In this example, we defined a segment called *Segment1and2* that includes people members of group1@contoso.com and group2@contoso.com and not members of group3@contoso.com.

TIP

If possible, use segment definitions that include "-eq" or "-ne". Try not to define complex segment definitions.

Part 2: Define information barrier policies

Determine whether you need to prevent communications between certain segments, or limit communications to certain segments. Ideally, you'll use the minimum number of policies to ensure your organization is compliant with legal and industry requirements.

With your list of user segments and the information barrier policies you want to define, select a scenario, and then follow the steps.

- Scenario 1: Block communications between segments
- Scenario 2: Allow a segment to communicate only with one other segment

IMPORTANT

Make sure that as you define policies, you do not assign more than one policy to a segment. For example, if you define one policy for a segment called *Sales*, do not define an additional policy for *Sales*.

In addition, as you define information barrier policies, make sure to set those policies to inactive status until you are ready to apply them. Defining (or editing) policies does not affect users until those policies are set to active status and then applied.

(See Example: Contoso's information barrier policies in this article.)

Scenario 1: Block communications between segments

When you want to block segments from communicating with each other, you define two policies: one for each direction. Each policy blocks communication one way only.

For example, suppose you want to block communications between Segment A and Segment B. In this case, you define one policy preventing Segment A from communicating with Segment B, and then define a second policy to prevent Segment B from communicating with Segment A.

1. To define your first blocking policy, use the **New-InformationBarrierPolicy** cmdlet with the **SegmentsBlocked** parameter.

SYNTAX	EXAMPLE
`New-InformationBarrierPolicy -Name "policyname" -AssignedSegment "segment1name" -SegmentsBlocked "segment2name"`	`New-InformationBarrierPolicy -Name "Sales-Research" -AssignedSegment "Sales" -SegmentsBlocked "Research" -State Inactive` In this example, we defined a policy called *Sales-Research* for a segment called *Sales*. When active and applied, this policy prevents people in *Sales* from communicating with people in a segment called *Research*.

2. To define your second blocking segment, use the **New-InformationBarrierPolicy** cmdlet with the **SegmentsBlocked** parameter again, this time with the segments reversed.

EXAMPLE	NOTE
`New-InformationBarrierPolicy -Name "Research-Sales" -AssignedSegment "Research" -SegmentsBlocked "Sales" -State Inactive`	In this example, we defined a policy called *Research-Sales* to prevent *Research* from communicating with *Sales*.

3. Proceed to one of the following actions:

   - (If needed) Define a policy to allow a segment to communicate only with one other segment
   - (After all your policies are defined) Apply information barrier policies

Scenario 2: Allow a segment to communicate only with one other segment

1. To allow one segment to communicate with only one other segment, use the **New-InformationBarrierPolicy** cmdlet with the **SegmentsAllowed** parameter.

SYNTAX	EXAMPLE
`New-InformationBarrierPolicy -Name "policyname" -AssignedSegment "segment1name" -SegmentsAllowed "segment2name","segment1name"`	`New-InformationBarrierPolicy -Name "Manufacturing-HR" -AssignedSegment "Manufacturing" -SegmentsAllowed "HR","Manufacturing" -State Inactive` In this example, we defined a policy called *Manufacturing-HR* for a segment called *Manufacturing*. When active and applied, this policy allows people in *Manufacturing* to communicate only with people in a segment called *HR*. (In this case, *Manufacturing* cannot communicate with users who are not part of *HR*.)

If needed, you can specify multiple segments with this cmdlet, as shown in the following example.

SYNTAX	EXAMPLE
`New-InformationBarrierPolicy -Name "policyname" -AssignedSegment "segment1name" -SegmentsAllowed "segment2name", "segment3name","segment1name"`	`New-InformationBarrierPolicy -Name "Research-HRManufacturing" -AssignedSegment "Research" -SegmentsAllowed "HR","Manufacturing","Research" -State Inactive` In this example, we defined a policy that allows the *Research* segment to communicate with only *HR* and *Manufacturing*.

Repeat this step for each policy you want to define to allow specific segments to communicate with only certain other specific segments.

2. Proceed to one of the following actions:

   - (If needed) Define a policy to block communications between segments
   - (After all your policies are defined) Apply information barrier policies

Part 3: Apply information barrier policies

Information barrier policies are not in effect until you set them to active status, and then apply the policies.

1. Use the **Get-InformationBarrierPolicy** cmdlet to see a list of policies that have been defined. Note the status and identity (GUID) of each policy.

   Syntax: `Get-InformationBarrierPolicy`

2. To set a policy to active status, use the **Set-InformationBarrierPolicy** cmdlet with an **Identity** parameter, and the **State** parameter set to **Active**.

SYNTAX	EXAMPLE
`Set-InformationBarrierPolicy -Identity GUID -State Active`	`Set-InformationBarrierPolicy -Identity 43c37853-ea10-4b90-a23d-ab8c93772471 -State Active` In this example, we set an information barrier policy that has the GUID *43c37853-ea10-4b90-a23d-ab8c93772471* to active status.

   Repeat this step as appropriate for each policy.

3. When you have finished setting your information barrier policies to active status, use the **Start-InformationBarrierPoliciesApplication** cmdlet in the Security & Compliance Center.

   Syntax: `Start-InformationBarrierPoliciesApplication`

   After you run `Start-InformationBarrierPoliciesApplication`, allow 30 minutes for the system to start applying the policies. The system applies policies user by user. The system processes about 5,000 user accounts per hour.

View status of user accounts, segments, policies, or policy application

With PowerShell, you can view status of user accounts, segments, policies, and policy application, as listed in the following table.

TO VIEW THIS INFORMATION	TAKE THIS ACTION

TO VIEW THIS INFORMATION	TAKE THIS ACTION

User accounts	Use the **Get-InformationBarrierRecipientStatus** cmdlet with Identity parameters.
	Syntax:
	`Get-InformationBarrierRecipientStatus -Identity <value> -Identity2 <value>`
	You can use any value that uniquely identifies each user, such as name, alias, distinguished name, canonical domain name, email address, or GUID.
	Example:
	`Get-InformationBarrierRecipientStatus -Identity meganb -Identity2 alexw`
	In this example, we refer to two user accounts in Office 365: *meganb* for *Megan*, and *alexw* for *Alex*.
	(You can also use this cmdlet for a single user:
	`Get-InformationBarrierRecipientStatus -Identity <value>`
	)
	This cmdlet returns information about users, such as attribute values and any information barrier policies that are applied.

Segments	Use the **Get-OrganizationSegment** cmdlet. Syntax: `Get-OrganizationSegment` This cmdlet will display a list of all segments defined for your organization.
Information barrier policies	Use the **Get-InformationBarrierPolicy** cmdlet. Syntax: `Get-InformationBarrierPolicy` This cmdlet will display a list of information barrier policies that were defined, and their status.
The most recent information barrier policy application	Use the **Get-InformationBarrierPoliciesApplicationStatus** cmdlet. Syntax: `Get-InformationBarrierPoliciesApplicationStatus` This cmdlet will display information about whether policy application completed, failed, or is in progress.
All information barrier policy applications	Use `Get-InformationBarrierPoliciesApplicationStatus -All` This cmdlet will display information about whether policy application completed, failed, or is in progress.

What if I need to remove or change policies?

Resources are available to help you manage your information barrier policies.

- If something goes wrong with information barriers, see Troubleshooting information barriers.

- To stop policies from being applied, see Stop a policy application.
- To remove an information barrier policy, see Remove a policy.
- To make changes to segments or policies, see Edit (or remove) information barrier policies.

Example: Contoso's departments, segments, and policies

To see how an organization might approach defining segments and policies, consider the following example.

Contoso's departments and plan

Contoso has five departments: HR, Sales, Marketing, Research, and Manufacturing. In order to remain compliant with industry regulations, people in some departments are not supposed to communicate with other departments, as listed in the following table:

SEGMENT	CAN TALK TO	CANNOT TALK TO
HR	Everyone	(no restrictions)
Sales	HR, Marketing, Manufacturing	Research
Marketing	Everyone	(no restrictions)
Research	HR, Marketing, Manufacturing	Sales
Manufacturing	HR, Marketing	Anyone other than HR or Marketing

For this structure, Contoso's plan includes three information barrier policies:

1. A policy designed to prevent Sales from communicating with Research (and another policy to prevent Research from communicating with Sales).
2. A policy designed to allow Manufacturing to communicate with HR and Marketing only.

For this scenario, it's not necessary to define policies for HR or Marketing.

Contoso's defined segments

Contoso will use the Department attribute in Azure Active Directory to define segments, as follows:

DEPARTMENT	SEGMENT DEFINITION
HR	`New-OrganizationSegment -Name "HR" -UserGroupFilter "Department -eq 'HR'"`
Sales	`New-OrganizationSegment -Name "Sales" -UserGroupFilter "Department -eq 'Sales'"`
Marketing	`New-OrganizationSegment -Name "Marketing" -UserGroupFilter "Department -eq 'Marketing'"`
Research	`New-OrganizationSegment -Name "Research" -UserGroupFilter "Department -eq 'Research'"`
Manufacturing	`New-OrganizationSegment -Name "Manufacturing" -UserGroupFilter "Department -eq 'Manufacturing'"`

With the segments defined, Contoso proceeds to define policies.

Contoso's information barrier policies

Contoso defines three policies, as described in the following table:

POLICY	POLICY DEFINITION
Policy 1: Prevent Sales from communicating with Research	`New-InformationBarrierPolicy -Name "Sales-Research" -AssignedSegment "Sales" -SegmentsBlocked "Research" -State Inactive` In this example, the information barrier policy is called *Sales-Research*. When this policy is active and applied, it will help prevent users who are in the Sales segment from communicating with users in the Research segment. This policy is a one-way policy; it won't prevent Research from communicating with Sales. For that, Policy 2 is needed.
Policy 2: Prevent Research from communicating with Sales	`New-InformationBarrierPolicy -Name "Research-Sales" -AssignedSegment "Research" -SegmentsBlocked "Sales" -State Inactive` In this example, the information barrier policy is called *Research-Sales*. When this policy is active and applied, it will help prevent users who are in the Research segment from communicating with users in the Sales segment.
Policy 3: Allow Manufacturing to communicate with HR and Marketing only	`New-InformationBarrierPolicy -Name "Manufacturing-HRMarketing" -AssignedSegment "Manufacturing" -SegmentsAllowed "HR","Marketing","Manufacturing" -State Inactive` In this case, the information barrier policy is called *Manufacturing-HRMarketing*. When this policy is active and applied, Manufacturing can communicate only with HR and Marketing. HR and Marketing are not restricted from communicating with other segments.

With segments and policies defined, Contoso applies the policies by running the **Start-InformationBarrierPoliciesApplication** cmdlet.

When the cmdlet finishes, Contoso is compliant with legal and industry requirements.

Resources

- Get an overview of information barriers
- Learn more about information barriers in Microsoft Teams
- Learn more about information barriers in SharePoint Online
- Learn more about information barriers in OneDrive

Attributes for information barrier policies

2/18/2021 • 2 minutes to read • Edit Online

Certain attributes in Azure Active Directory can be used to segment users. Once segments are defined, those segments can be used as filters for information barrier policies. For example, you might use **Department** to define segments of users by department within your organization (assuming no single employee works for two departments at the same time).

This article describes how to use attributes with information barriers, and it provides a list of attributes that can be used. To learn more about information barriers, see the following resources:

- Information barriers
- Define policies for information barriers in Microsoft Teams
- Edit (or remove) information barrier policies

How to use attributes in information barrier policies

The attributes listed in this article can be used to define or edit segments of users. Your defined segments serve as parameters (called *UserGroupFilter* values) in information barrier policies.

1. Determine which attribute you want to use to define segments. (See the Reference section in this article.)
2. Make sure the user accounts have values filled in for the attribute(s) you selected in Step 1. View user account details, and if necessary, edit user accounts to include attribute values.
   - To edit multiple accounts (or use PowerShell to edit a single account), see Configure user account properties with Office 365 PowerShell.
   - To edit a single account, see Add or update a user's profile information using Azure Active Directory.
3. Define segments using PowerShell, similar to the following examples:

EXAMPLE	CMDLET
Define a segment called Segment1 using the Department attribute	`New-OrganizationSegment -Name "Segment1" -UserGroupFilter "Department -eq 'Department1'"`
Define a segment called SegmentA using the MemberOf attribute (suppose this attribute contains group names, such as "BlueGroup")	`New-OrganizationSegment -Name "SegmentA" -UserGroupFilter "MemberOf -eq 'BlueGroup'"`
Define a segment called DayTraders using ExtensionAttribute1 (suppose this attribute contains job titles, such as "DayTrader")	`New-OrganizationSegment -Name "DayTraders" -UserGroupFilter "ExtensionAttribute1 -eq 'DayTrader'"`

TIP

When you define segments, use the same attribute for all your segments. For example, if you define some segments using *Department*, define all of the segments using *Department*. Don't define some segments using *Department* and others using *MemberOf*. Make sure your segments do not overlap; each user should be assigned to exactly one segment.

Reference

The following table lists the attributes that you can use with information barriers.

AZURE ACTIVE DIRECTORY PROPERTY NAME (LDAP DISPLAY NAME)	EXCHANGE PROPERTY NAME
Co	Co
Company	Company
Department	Department
ExtensionAttribute1	CustomAttribute1
ExtensionAttribute2	CustomAttribute2
ExtensionAttribute3	CustomAttribute3
ExtensionAttribute4	CustomAttribute4
ExtensionAttribute5	CustomAttribute5
ExtensionAttribute6	CustomAttribute6
ExtensionAttribute7	CustomAttribute7
ExtensionAttribute8	CustomAttribute8
ExtensionAttribute9	CustomAttribute9
ExtensionAttribute10	CustomAttribute10
ExtensionAttribute11	CustomAttribute11
ExtensionAttribute12	CustomAttribute12
ExtensionAttribute13	CustomAttribute13
ExtensionAttribute14	CustomAttribute14
ExtensionAttribute15	CustomAttribute15
MSExchExtensionCustomAttribute1	ExtensionCustomAttribute1
MSExchExtensionCustomAttribute2	ExtensionCustomAttribute2
MSExchExtensionCustomAttribute3	ExtensionCustomAttribute3
MSExchExtensionCustomAttribute4	ExtensionCustomAttribute4
MSExchExtensionCustomAttribute5	ExtensionCustomAttribute5

AZURE ACTIVE DIRECTORY PROPERTY NAME (LDAP DISPLAY NAME)	EXCHANGE PROPERTY NAME
MailNickname	Alias
PhysicalDeliveryOfficeName	Office
PostalCode	PostalCode
ProxyAddresses	EmailAddresses
StreetAddress	StreetAddress
TargetAddress	ExternalEmailAddress
UsageLocation	UsageLocation
UserPrincipalName	UserPrincipalName
Mail	WindowsEmailAddress
Description	Description
MemberOf	MemberOfGroup

Resources

- Define policies for information barriers in Microsoft Teams
- Troubleshooting information barriers
- Information barriers

Troubleshooting information barriers

2/18/2021 • 9 minutes to read • Edit Online

Information barriers can help your organization remain compliant with legal requirements and industry regulations. For example, with information barriers, you can restrict communication between specific groups of users to avoid a conflict of interest or other issues. (To learn more about how to set up information barriers, see Define policies for information barriers.)

In the event that people run into unexpected issues after information barriers are in place, there are some steps you can take to resolve those issues. Use this article as a guide.

> **IMPORTANT**
>
> To perform the tasks described in this article, you must be assigned an appropriate role, such as one of the following:
> - Microsoft 365 Enterprise Global Administrator
> - global administrator
> - Compliance Administrator
> - IB Compliance Management (this is a new role!)
>
> To learn more about prerequisites for information barriers, see Prerequisites (for information barrier policies).
>
> Make sure to connect to Security & Compliance Center PowerShell.

Issue: Users are unexpectedly blocked from communicating with others in Microsoft Teams

In this case, people are reporting unexpected issues communicating with others in Microsoft Teams. Some examples are:

- A user searches for, but is unable to find, another user in Microsoft Teams.
- A user can find, but cannot select, another user in Microsoft Teams.
- A user can see another user, but cannot send messages to that other user in Microsoft Teams.

What to do

Determine whether the users are affected by an information barrier policy. Depending on how policies are configured, information barriers might be working as expected. Or, you might have to refine your organization's policies.

1. Use the **Get-InformationBarrierRecipientStatus** cmdlet with the Identity parameter.

SYNTAX	EXAMPLE
`Get-InformationBarrierRecipientStatus -Identity` You can use any identity value that uniquely identifies each recipient, such as Name, Alias, Distinguished name (DN), Canonical DN, Email address, or GUID.	`Get-InformationBarrierRecipientStatus -Identity meganb` In this example, we are using an alias (*meganb*) for the Identity parameter. This cmdlet will return information that indicates whether the user is affected by an information barrier policy. (Look for *ExoPolicyId: <GUID>.)

   If the users are not included in information barrier policies, contact support. Otherwise, proceed to the next step.

2. Find out which segments are included in an information barrier policy. To do this, use the `Get-InformationBarrierPolicy` cmdlet with the Identity parameter.

SYNTAX	EXAMPLE
`Get-InformationBarrierPolicy` Use details, such as the policy GUID (ExoPolicyId) you received during the previous step, as an identity value.	`Get-InformationBarrierPolicy -Identity b42c3d0f-49e9-4506-a0a5-bf2853b5df6f` In this example, we are getting detailed information about the information barrier policy that has ExoPolicyId *b42c3d0f-49e9-4506-a0a5-bf2853b5df6f*.

   After you run the cmdlet, in the results, look for **AssignedSegment**, **SegmentsAllowed**, and **SegmentsBlocked** values.

   For example, after running the `Get-InformationBarrierPolicy` cmdlet, we saw the following in our list of results:

```
AssignedSegment : Sales
SegmentsAllowed : {}
SegmentsBlocked : {Research}
```

   In this case, we can see that an information barrier policy affects people who are in the Sales and Research segments. In this case, people in Sales are prevented from communicating with people in Research.

   If this seems correct, then information barriers are working as expected. If not, proceed to the next step.

3. Make sure your segments are defined correctly. To do this, use the `Get-OrganizationSegment` cmdlet, and review the list of results.

SYNTAX	EXAMPLE
`Get-OrganizationSegment` Use this cmdlet with an Identity parameter.	`Get-OrganizationSegment -Identity c96e0837-c232-4a8a-841e-ef45787d8fcd` In this example, we are getting information about the segment that has GUID *c96e0837-c232-4a8a-841e-ef45787d8fcd*.

   Review the details for the segment. If necessary, edit a segment, and then re-use the `Start-InformationBarrierPoliciesApplication` cmdlet.

   If you are still having issues with your information barrier policy, contact support.

Issue: Communications are allowed between users who should be blocked in Microsoft Teams

In this case, although information barriers are defined, active, and applied, people who should be prevented from communicating with each other are somehow able to chat with and call each other in Microsoft Teams.

What to do

Verify that the users in question are included in an information barrier policy.

1. Use the **Get-InformationBarrierRecipientStatus** cmdlet with Identity parameters.

SYNTAX*	EXAMPLE
`Get-InformationBarrierRecipientStatus -Identity <value> -Identity2 <value>` You can use any value that uniquely identifies each user, such as name, alias, distinguished name, canonical domain name, email address, or GUID.	`Get-InformationBarrierRecipientStatus -Identity meganb -Identity2 alexw` In this example, we refer to two user accounts in Office 365: *meganb* for *Megan*, and *alexw* for *Alex*.

> **TIP**
> You can also use this cmdlet for a single user: `Get-InformationBarrierRecipientStatus -Identity <value>`

2. Review the findings. The **Get-InformationBarrierRecipientStatus** cmdlet returns information about users, such as attribute values and any information barrier policies that are applied.

   Review the results, and then take your next steps, as described in the following table:

RESULTS	WHAT TO DO NEXT
No segments are listed for the selected user(s)	Do one of the following: - Assign users to an existing segment by editing their user profiles in Azure Active Directory. (See Configure user account properties with Office 365 PowerShell.) - Define a segment using a supported attribute for information barriers. Then, either define a new policy or edit an existing policy to include that segment.
Segments are listed but no information barrier policies are assigned to those segments	Do one of the following: - Define a new information barrier policy for each segment in question - Edit an existing information barrier policy to assign it to the correct segment
Segments are listed and each is included in an information barrier policy	- Run the `Get-InformationBarrierPolicy` cmdlet to verify that information barrier policies are active - Run the `Get-InformationBarrierPoliciesApplicationStatus` cmdlet to confirm the policies are applied - Run the `Start-InformationBarrierPoliciesApplication` cmdlet to apply all active information barrier policies

Issue: I need to remove a single user from an information barrier policy

In this case, information barrier policies are in effect, and a one or more users are unexpectedly blocked from communicating with others in Microsoft Teams. Rather than remove information barrier policies altogether, you can remove one or more individual users from information barrier policies.

What to do

Information barrier policies are assigned to segments of users. Segments are defined by using certain attributes in user account profiles. If you must remove a policy from a single user, consider editing that user's profile in Azure Active Directory such that the user is no longer included in a segment affected by information barriers.

1. Use the **Get-InformationBarrierRecipientStatus** cmdlet with Identity parameters. This cmdlet returns

information about users, such as attribute values and any information barrier policies that are applied.

SYNTAX	EXAMPLE
`Get-InformationBarrierRecipientStatus -Identity <value> -Identity2 <value>` You can use any value that uniquely identifies each user, such as name, alias, distinguished name, canonical domain name, email address, or GUID.	`Get-InformationBarrierRecipientStatus -Identity meganb -Identity2 alexw` In this example, we refer to two user accounts in Office 365: *meganb* for *Megan*, and *alexw* for *Alex*.
`Get-InformationBarrierRecipientStatus -Identity <value>` You can use any value that uniquely identifies the user, such as name, alias, distinguished name, canonical domain name, email address, or GUID.	`Get-InformationBarrierRecipientStatus -Identity jeanp` In this example, we refer to a single account in Office 365: *jeanp*.

2. Review the results to see if information barrier policies are assigned, and to which segment(s) the user(s) belong.
3. To remove a user from a segment affected by information barriers, update the user's profile information in Azure Active Directory.
4. Wait about 30 minutes for FwdSync to occur. Or, run the `Start-InformationBarrierPoliciesApplication` cmdlet to apply all active information barrier policies.

Issue: The information barrier application process is taking too long

After running the **Start-InformationBarrierPoliciesApplication** cmdlet, the process is taking a really long time to finish.

What to do

Keep in mind that when you run the policy application cmdlet, information barrier policies are being applied (or removed), user by user, for all accounts in your organization. If you have a lot of users, it will take a while to process. (As a general guideline, it takes about an hour to process 5,000 user accounts.)

1. Use the **Get-InformationBarrierPoliciesApplicationStatus** cmdlet to verify status of the most recent policy application.

TO VIEW THE MOST RECENT POLICY APPLICATION	TO VIEW STATUS FOR ALL POLICY APPLICATIONS
`Get-InformationBarrierPoliciesApplicationStatus`	`Get-InformationBarrierPoliciesApplicationStatus -All $true`

   This will display information about whether policy application completed, failed, or is in progress.

2. Depending on the results of the previous step, take one of the following steps:

STATUS	NEXT STEP
Not started	If it has been more than 45 minutes since the **Start-InformationBarrierPoliciesApplication** cmdlet has been run, review your audit log to see if there are any errors in policy definitions, or some other reason why the application has not started.

STATUS	NEXT STEP
Failed	If the application has failed, review your audit log. Also review your segments and policies. Are any users assigned to more than one segment? Are any segments assigned more than one poliicy? If necessary, edit segments and/or edit policies, and then run the **Start-InformationBarrierPoliciesApplication** cmdlet again.
In progress	If the application is still in progress, allow more time for it to complete. If it has been several days, gather your audit logs, and then contact support.

Issue: Information barrier policies are not being applied at all

In this case, you have defined segments, defined information barrier policies, and have attempted to apply those policies. However, when you run the `Get-InformationBarrierPoliciesApplicationStatus` cmdlet, you can see that policy application has failed.

What to do

Make sure that your organization does not have Exchange address book policies in place. Such policies will prevent information barrier policies from being applied.

1. Connect to Exchange Online PowerShell.
2. Run the Get-AddressBookPolicy cmdlet, and review the results.

RESULTS	NEXT STEP
Exchange address book policies are listed	Remove address book policies
No address book policies exist	Review your audit logs to find out why policy application is failing

3. View status of user accounts, segments, policies, or policy application.

Issue: Information barrier policy not applied to all designated users

After you have defined segments, defined information barrier policies, and have attempted to apply those policies, you may find that the policy is applying to some recipients, but not to others. When you run the `Get-InformationBarrierPoliciesApplicationStatus` cmdlet, search the output for text like this.

> Identity: `<application guid>`
>
> Total Recipients: 81527
>
> Failed Recipients: 2
>
> Failure Category: None
>
> Status: Complete

What to do

1. Search in the audit log for `<application guid>`. You can copy this PowerShell code and modify for your variables.

```
$DetailedLogs = Search-UnifiedAuditLog -EndDate <yyyy-mm-ddThh:mm:ss>  -StartDate <yyyy-mm-ddThh:mm:ss> -
RecordType InformationBarrierPolicyApplication -ResultSize 1000 |?{$_.AuditData.Contains(<application
guid>)}
```

2. Check the detailed output from the audit log for the values of the `"UserId"` and `"ErrorDetails"` fields. This will give you the reason for the failure. You can copy this PowerShell code and modify for your variables.

```
   $DetailedLogs[1] |fl
```

For example:

> "UserId": User1
>
> "ErrorDetails":"Status: IBPolicyConflict. Error: IB segment "segment id1" and IB segment "segment id2" has conflict and cannot be assigned to the recipient.

3. Usually, you will find that a user has been included in more than one segment. You can fix this by updating the `-UserGroupFilter` value in `OrganizationSegments`.
4. Re-apply information barrier policies using these procedures Information Barriers policies.

Resources

- Define policies for information barriers in Microsoft Teams
- Information barriers

Manage information barrier policies

2/18/2021 • 6 minutes to read • Edit Online

After you have defined information barrier policies, you might need to make changes to those policies or to your user segments, as part of troubleshooting or as regular maintenance. Use this article as a guide.

What do you want to do?

ACTION	DESCRIPTION
Edit user account attributes	Fill in attributes in Azure Active Directory that can be used to define segments. Edit user account attributes when users are not included in segments they should be, to change which segments users are in, or to define segments using different attributes.
Edit a segment	Edit segments when you want to change how a segment is defined. For example, you might have originally defined segments using *Department* and now want to use another attribute, such as *MemberOf*.
Edit a policy	Edit an information barrier policy when you want to change how a policy works. For example, instead of blocking communications between two segments, you might decide you want to allow communications to occur only between certain segments.
Set a policy to inactive status	Set a policy to inactive status when you want to make changes to a policy, or when you don't want a policy to be in effect.
Remove a policy	Remove an information barrier policy when you no longer need a particular policy in place.
Stop a policy application	Take this action when you want to stop the process of applying information barrier policies. Stopping a policy application is not instant, and it does not undo policies that are already applied to users.
Define policies for information barriers	Define an information barrier policy when you do not already have such policies in place, and you must restrict or limit communications between specific groups of users.
Troubleshooting information barriers	Refer to this article when you run into unexpected issues with information barriers.

IMPORTANT

To perform the tasks described in this article, you must be assigned an appropriate role, such as one of the following:
- Microsoft 365 Enterprise Global Administrator
- Global Administrator
- Compliance Administrator
- IB Compliance Management (this is a new role!)

To learn more about prerequisites for information barriers, see Prerequisites (for information barrier policies).

Make sure to connect to the Security & Compliance Center PowerShell.

Edit user account attributes

Use this procedure to edit attributes that are used for segmenting users. For example, if you are using a Department attribute, and one or more user accounts do not currently have any values listed for Department, you must edit those user accounts to include Department information. User account attributes are used for defining segments so that information barrier policies can be assigned.

1. To view details for a specific user account, such as attribute values and assigned segment(s), use the **Get-InformationBarrierRecipientStatus** cmdlet with Identity parameters.

SYNTAX	EXAMPLE
`Get-InformationBarrierRecipientStatus -Identity <value> -Identity2 <value>` You can use any value that uniquely identifies each user, such as name, alias, distinguished name, canonical domain name, email address, or GUID. (You can also use this cmdlet for a single user: `Get-InformationBarrierRecipientStatus -Identity <value>`)	`Get-InformationBarrierRecipientStatus -Identity meganb -Identity2 alexw` In this example, we refer to two user accounts in Office 365: *meganb* for *Megan*, and *alexw* for *Alex*.

2. Determine which attribute you want to edit for your user account profile(s). For more information, see Attributes for information barrier policies.

3. Edit one or more user accounts to include values for the attribute you selected in the previous step. To take this action, use one of the following procedures:

   - To edit a single account, see Add or update a user's profile information using Azure Active Directory.
   - To edit multiple accounts (or use PowerShell to edit a single account), see Configure user account properties with Office 365 PowerShell.

Edit a segment

Use this procedure edit the definition of a user segment. For example, you might change the name of a segment, or the filter that is used to determine who's included in the segment.

1. To view all existing segments, use the **Get-OrganizationSegment** cmdlet.

   Syntax: `Get-OrganizationSegment`

   You will see a list of segments and details for each, such as segment type, its UserGroupFilter value, who

created or last modified it, GUID, and so on.

> **TIP**
> Print or save your list of segments for reference later. For example, if you want to edit a segment, you will need to know its name or identify value (this is used with the Identity parameter).

2. To edit a segment, use the **Set-OrganizationSegment** cmdlet with the **Identity** parameter and relevant details.

SYNTAX	EXAMPLE
`Set-OrganizationSegment -Identity GUID -UserGroupFilter "attribute -eq 'attributevalue'"`	`Set-OrganizationSegment -Identity c96e0837-c232-4a8a-841e-ef45787d8fcd -UserGroupFilter "Department -eq 'HRDept'"` In this example, for the segment that has the GUID *c96e0837-c232-4a8a-841e-ef45787d8fcd*, we updated the department name to "HRDept".

When you have finished editing segments for your organization, you can either define or edit information barrier policies.

Edit a policy

1. To view a list of current information barrier policies, use the **Get-InformationBarrierPolicy** cmdlet.

   Syntax: `Get-InformationBarrierPolicy`

   In the list of results, identify the policy that you want to change. Note the policy's GUID and name.

2. Use the **Set-InformationBarrierPolicy** cmdlet with an **Identity** parameter, and specify the changes you want to make.

   Example: Suppose a policy was defined to block the *Research* segment from communicating with the *Sales* and *Marketing* segments. The policy was defined by using this cmdlet:

   ```
   New-InformationBarrierPolicy -Name "Research-SalesMarketing" -AssignedSegment "Research" -SegmentsBlocked "Sales","Marketing"
   ```

   Suppose we want to change it so that people in the *Research* segment can only communicate with people in the *HR* segment. To make this change, we use this cmdlet:

   ```
   Set-InformationBarrierPolicy -Identity 43c37853-ea10-4b90-a23d-ab8c93772471 -SegmentsAllowed "HR"
   ```

   In this example, we changed "SegmentsBlocked" to "SegmentsAllowed" and specified the *HR* segment.

3. When you are finished editing a policy, make sure to apply your changes. (See Apply information barrier policies.)

Set a policy to inactive status

1. To view a list of current information barrier policies, use the **Get-InformationBarrierPolicy** cmdlet.

   Syntax: `Get-InformationBarrierPolicy`

   In the list of results, identify the policy that you want to change (or remove). Note the policy's GUID and name.

2. To set the policy's status to inactive, use the **Set-InformationBarrierPolicy** cmdlet with an Identity parameter and the State parameter set to Inactive.

SYNTAX	EXAMPLE
`Set-InformationBarrierPolicy -Identity GUID -State Inactive`	`Set-InformationBarrierPolicy -Identity 43c37853-ea10-4b90-a23d-ab8c9377247 -State Inactive` In this example, we set an information barrier policy that has GUID *43c37853-ea10-4b90-a23d-ab8c9377247* to an inactive status.

3. To apply your changes, use the **Start-InformationBarrierPoliciesApplication** cmdlet.

   Syntax: `Start-InformationBarrierPoliciesApplication`

   Changes are applied, user by user, for your organization. If your organization is large, it can take 24 hours (or more) for this process to complete. (As a general guideline, it takes about an hour to process 5,000 user accounts.)

At this point, one or more information barrier policies are set to inactive status. From here, you can do any of the following actions:

- Keep it as is (a policy set to inactive status has no effect on users)
- Edit a policy
- Remove a policy

Remove a policy

1. To view a list of current information barrier policies, use the **Get-InformationBarrierPolicy** cmdlet.

   Syntax: `Get-InformationBarrierPolicy`

   In the list of results, identify the policy that you want to remove. Note the policy's GUID and name. Make sure the policy is set to inactive status.

2. Use the **Remove-InformationBarrierPolicy** cmdlet with an Identity parameter.

SYNTAX	EXAMPLE
`Remove-InformationBarrierPolicy -Identity GUID`	`Remove-InformationBarrierPolicy -Identity 43c37853-ea10-4b90-a23d-ab8c93772471` In this example, we are removing the policy that has GUID *43c37853-ea10-4b90-a23d-ab8c93772471*.

   When prompted, confirm the change.

3. Repeat steps 1-2 for each policy you want to remove.

4. When you are finished removing policies, apply your changes. To take this action, use the **Start-InformationBarrierPoliciesApplication** cmdlet.

   Syntax: `Start-InformationBarrierPoliciesApplication`

   Changes are applied, user by user, for your organization. If your organization is large, it can take 24 hours (or more) for this process to complete.

Stop a policy application

After you have started applying information barrier policies, if you want to stop those policies from being applied, use the following procedure. It will take approximately 30-35 minutes for the process to begin.

1. To view the status of the most recent information barrier policy application, use the **Get-InformationBarrierPoliciesApplicationStatus** cmdlet.

   Syntax: `Get-InformationBarrierPoliciesApplicationStatus`

   Note the application's GUID.

2. Use the **Stop-InformationBarrierPoliciesApplication** cmdlet with an Identity parameter.

SYNTAX	EXAMPLE
`Stop-InformationBarrierPoliciesApplication -Identity GUID`	`Stop-InformationBarrierPoliciesApplication -Identity 46237888-12ca-42e3-a541-3fcb7b5231d1` In this example, we are stopping information barrier policies from being applied.

Resources

- Get an overview of information barriers
- Define policies for information barriers
- Learn more about information barriers in Microsoft Teams
- Learn more about information barriers in SharePoint Online
- Learn more about information barriers in OneDrive
- Attributes for information barrier policies
- Troubleshooting information barriers

Privileged access management in Microsoft 365

2/18/2021 • 2 minutes to read • Edit Online

Having standing access by some users to sensitive information or critical network configuration settings in Microsoft Exchange Online is a potential pathway for compromised accounts or internal threat activities. Privileged access management helps protect your organization from breaches and helps to meet compliance best practices by limiting standing access to sensitive data or access to critical configuration settings. Instead of administrators having constant access, just-in-time access rules are implemented for tasks that need elevated permissions. Enabling privileged access management for Exchange Online in Microsoft 365 allows your organization to operate with zero standing privileges and provide a layer of defense against standing administrative access vulnerabilities.

Configure privileged access management for Microsoft 365

Use the following steps to configure privileged access management for your organization:

1. Learn about privileged access management in Microsoft 365
2. Create an approver's group
3. Enable privileged access management
4. Create an access policy
5. Submit/approve privileged access requests

More information about privileged access management

- Frequently asked questions about privileged access management

Learn about privileged access management

2/18/2021 • 4 minutes to read • Edit Online

Privileged access management allows granular access control over privileged admin tasks in Office 365. It can help protect your organization from breaches that use existing privileged admin accounts with standing access to sensitive data or access to critical configuration settings. Privileged access management requires users to request just-in-time access to complete elevated and privileged tasks through a highly scoped and time-bounded approval workflow. This configuration gives users just-enough-access to perform the task at hand, without risking exposure of sensitive data or critical configuration settings. Enabling privileged access management in Microsoft 365 allows your organization to operate with zero standing privileges and provide a layer of defense against standing administrative access vulnerabilities.

For a quick overview of the integrated Customer Lockbox and privileged access management workflow, see this Customer Lockbox and privileged access management video.

Layers of protection

Privileged access management complements other data and access feature protections within the Microsoft 365 security architecture. Including privileged access management as part of an integrated and layered approach to security provides a security model that maximizes protection of sensitive information and Microsoft 365 configuration settings. As shown in the diagram, privileged access management builds on the protection provided with native encryption of Microsoft 365 data and the role-based access control security model of Microsoft 365 services. When used with Azure AD Privileged Identity Management, these two features provide access control with just-in-time access at different scopes.

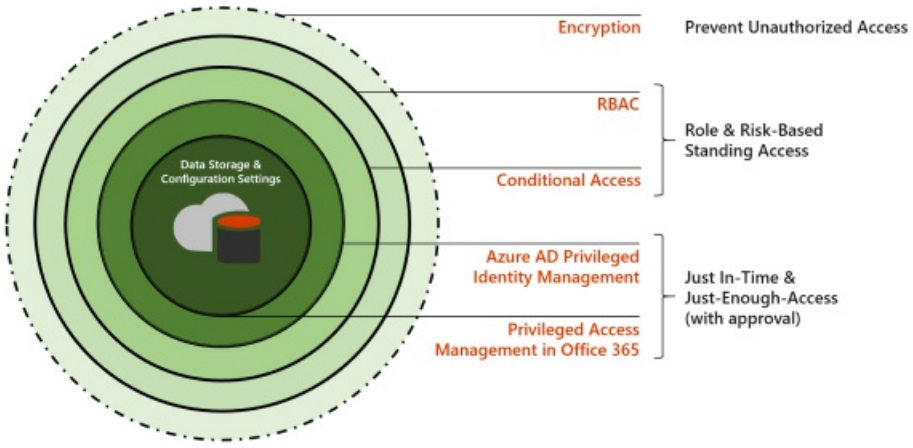

Privileged access management is defined and scoped at the **task** level, while Azure AD Privileged Identity Management applies protection at the **role** level with the ability to execute multiple tasks. Azure AD Privileged Identity Management primarily allows managing accesses for AD roles and role groups, while privileged access management in Microsoft 365 applies only at the task level.

- **Enabling privileged access management while already using Azure AD Privileged Identity Management:** Adding privileged access management provides another granular layer of protection and audit capabilities for privileged access to Microsoft 365 data.
- **Enabling Azure AD Privileged Identity Management while already using privileged access management in Office 365:** Adding Azure AD Privileged Identity Management to privileged access management can extend privileged access to data outside of Microsoft 365 that's primarily defined by

user roles or identity.

Privileged access management architecture and process flow

Each of the following process flows outline the architecture of privileged access and how it interacts with the Microsoft 365 substrate, auditing, and the Exchange Management runspace.

Step 1: Configure a privileged access policy

When you configure a privileged access policy with the Microsoft 365 admin center or the Exchange Management PowerShell, you define the policy and the privileged access feature processes and the policy attributes in the Microsoft 365 substrate. The activities are logged in the Security & Compliance Center. The policy is now enabled and ready to handle incoming requests for approvals.

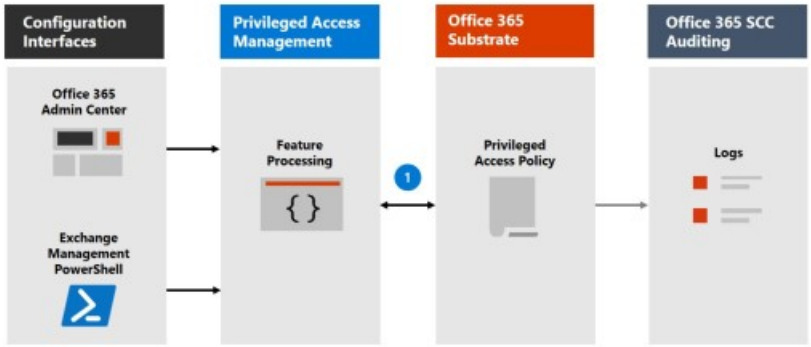

Step 2: Access request

In the Microsoft 365 admin center or with the Exchange Management PowerShell, users can request access to elevated or privileged tasks. The privileged access feature sends the request to the Microsoft 365 substrate for processing against the configured privilege access policy and records the Activity in the Security & Compliance Center logs.

Step 3: Access approval

An approval request is generated and the pending request notification is emailed to approvers. If approved, the privileged access request is processed as an approval and the task is ready to be completed. If denied, the task is blocked and no access is granted to the requestor. The requestor is notified of the request approval or denial via email message.

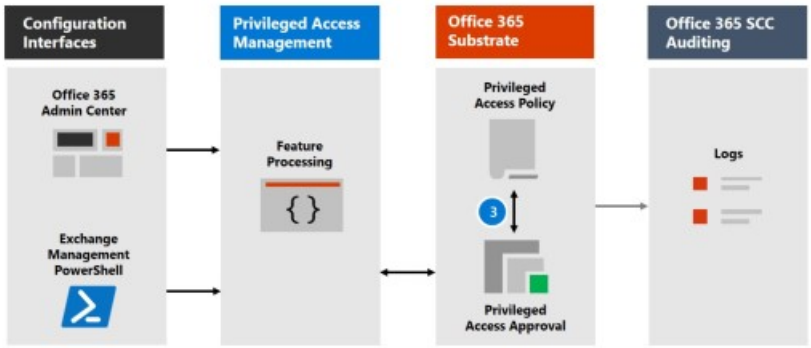

Step 4: Access processing

For an approved request, the task is processed by the Exchange Management runspace. The approval is checked against the privileged access policy and processed by the Microsoft 365 substrate. All activity for the task is logged in the Security & Compliance Center.

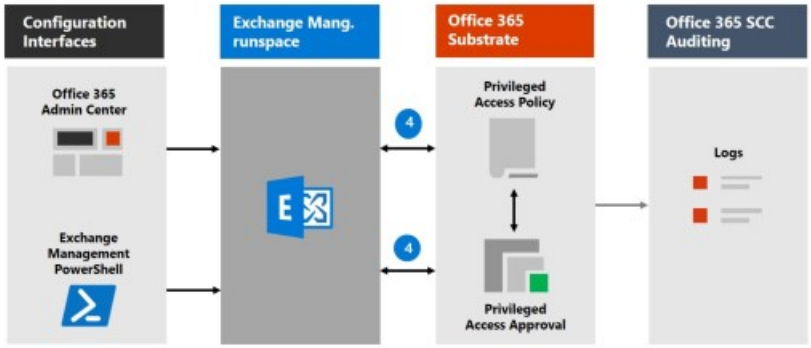

Frequently asked questions

What SKUs can use privileged access in Office 365?

Privileged access management is available for customers for a wide selection of Microsoft 365 and Office 365 subscriptions and add-ons. See Get started with privileged access management for details.

When will privileged access support Office 365 workloads beyond Exchange?

Privileged access management will be available in other Office 365 workloads soon. Visit the Microsoft 365 Roadmap for more details.

My organization needs more than 30 privileged access policies, will this limit be increased?

Yes, raising the current limit of 30 privileged access policies per organization is on the feature roadmap.

Do I need to be a Global Admin to manage privileged access in Office 365?

No, you need the Exchange Role Management role assigned to accounts that manage privileged access in Office 365. If you don't want to configure the Role Management role as a stand-alone account permission, the Global Administrator role includes this role by default and can manage privileged access. Users included in an approvers' group don't need to be a Global Admin or have the Role Management role assigned to review and approve requests with PowerShell.

How is privileged access management related to Customer Lockbox?

Customer Lockbox allows a level of access control for organizations when Microsoft accesses data. Privileged access management allows granular access control within an organization for all Microsoft 365 privileged tasks.

Ready to get started?

Start configuring your organization for privileged access management.

Learn more

Interactive guide: Monitor and control administrator tasks with privileged access management

Get started with privileged access management

2/18/2021 • 8 minutes to read • Edit Online

This topic guides you through enabling and configuring privileged access management in your organization. You can use either the Microsoft 365 admin center or Exchange Management PowerShell to manage and use privileged access.

Before you begin

Before you get started with privileged access management, you should confirm your Microsoft 365 subscription and any add-ons. To access and use privileged access management, your organization must have one of the following subscriptions or add-ons:

- Microsoft 365 E5 subscription (paid or trial version)
- Microsoft 365 E3 subscription (or Office 365 E3 subscription + Enterprise Mobility and Security E3 subscription) + the Microsoft 365 E5 Compliance add-on
- Any Microsoft 365, Office 365, Exchange, SharePoint, or OneDrive for Business subscription + the Microsoft 365 E5 Insider Risk Management add-on
- Microsoft 365 A5 subscription (paid or trial version)
- Microsoft 365 A3 subscription (or Office 365 A3 subscription + Enterprise Mobility and Security A3 subscription) + the Microsoft A5 Compliance add-on
- Any Microsoft 365, Office 365, Exchange, SharePoint, or OneDrive for Education subscription + the Microsoft 365 A5 Insider Risk Management add-on
- Office 365 Enterprise E5 subscription (paid or trial version)
- Office 365 Enterprise E3 subscription + the Office 365 Advanced Compliance add-on (no longer available for new subscriptions, see note)

Users submitting and responding to privileged access management requests must be assigned one of the licenses above.

> **IMPORTANT**
> Office 365 Advanced Compliance is no longer sold as a standalone subscription. When current subscriptions expire, customers should transition to one of the subscriptions above, which contain the same or additional compliance features.

If you don't have an existing Office 365 Enterprise E5 plan and want to try privileged access management, you can add Microsoft 365 to your existing Office 365 subscription or sign up for a trial of Microsoft 365 Enterprise E5.

Enable and configure privileged access management

Follow these steps to set up and use privileged access in your organization:

- Step 1: Create an approver's group

  Before you start using privilege access, determine who needs approval authority for incoming requests for access to elevated and privileged tasks. Any user who is part of the Approvers' group is able to approve access requests. This group is enabled by creating a mail-enabled security group in Office 365.

- Step 2: Enable privileged access

Privileged access must be explicitly enabled in Office 365 with the default approver group, including a set of system accounts that you want excluded from the privileged access management access control.

- Step 3: Create an access policy

  Creating an approval policy allows you to define the specific approval requirements scoped at individual tasks. The approval type options are **Auto** or **Manual**.

- Step 4: Submit/approve privileged access requests

  Once enabled, privileged access requires approvals for any task that has an associated approval policy defined. For tasks included in an approval policy, users must request and be granted access approval to have permissions necessary to execute the task.

After approval is granted, the requesting user can execute the intended task and privileged access will authorize and execute the task on behalf of the user. The approval remains valid for the requested duration (default duration is 4 hours), during which the requester can execute the intended task multiple times. All such executions are logged and made available for security and compliance auditing.

> **NOTE**
> If you want to use Exchange Management PowerShell to enable and configure privileged access, follow the steps in Connect to Exchange Online PowerShell using Multi-Factor authentication to connect to Exchange Online PowerShell with your Office 365 credentials. You do not need to enable multi-factor authentication for your organization to use the steps to enable privileged access while connecting to Exchange Online PowerShell. Connecting with multi-factor authentication creates an OAuth token that is used by privileged access for signing your requests.

Step 1: Create an approver's group

1. Sign into the Microsoft 365 admin center using credentials for an admin account in your organization.
2. In the Admin Center, go to **Groups** > **Add a group**.
3. Select **mail-enabled security group** and then complete the **Name**, **Group email address**, and **Description** fields for the new group.
4. Save the group. It may take a few minutes for the group to be fully configured and to appear in the Microsoft 365 admin center.
5. Select the new approver's group and select **edit** to add users to the group.
6. Save the group.

Step 2: Enable privileged access

In the Microsoft 365 Admin Center

1. Sign into the Microsoft 365 Admin Center using credentials for an admin account in your organization.
2. In the Admin Center, go to **Settings** > **Org Settings** > **Security & Privacy** > **Privileged access**.
3. Enable the **Require approvals for privileged tasks** control.
4. Assign the approver's group you created in Step 1 as the **Default approvers group**.
5. **Save** and **Close**.

In Exchange Management PowerShell

To enable privileged access and to assign the approver's group, run the following command in Exchange Online

PowerShell:

```
Enable-ElevatedAccessControl -AdminGroup '<default approver group>' -SystemAccounts
@('<systemAccountUPN1>','<systemAccountUPN2>')
```

Example:

```
Enable-ElevatedAccessControl -AdminGroup 'pamapprovers@fabrikam.onmicrosoft.com' -SystemAccounts
@('sys1@fabrikamorg.onmicrosoft.com', 'sys2@fabrikamorg.onmicrosoft.com')
```

NOTE

System accounts feature is made available to ensure certain automations within your organizations can work without dependency on privileged access, however it is recommended that such exclusions be exceptional and those allowed should be approved and audited regularly.

Step 3: Create an access policy

You can create and configure up to 30 privileged access policies for your organization.

In the Microsoft 365 Admin Center

1. Sign into the Microsoft 365 Admin Center using credentials for an admin account in your organization.
2. In the Admin Center, go to **Settings** > **Org Settings** > **Security & Privacy** > **Privileged access**.
3. Select **Manage access policies and requests**.
4. Select **Configure policies** and select **Add a policy**.
5. From the drop-down fields, select the appropriate values for your organization:

   Policy type: Task, Role, or Role Group

   Policy scope: Exchange

   Policy name: Select from the available policies

   Approval type: Manual or Auto

   Approval group: Select the approvers group created in Step 1
6. Select **Create** and then **Close**. It may take a few minutes for the policy to be fully configured and enabled.

In Exchange Management PowerShell

To create and define an approval policy, run the following command in Exchange Online PowerShell:

```
New-ElevatedAccessApprovalPolicy -Task 'Exchange\<exchange management cmdlet name>' -ApprovalType <Manual,
Auto> -ApproverGroup '<default/custom approver group>'
```

Example:

```
New-ElevatedAccessApprovalPolicy -Task 'Exchange\New-MoveRequest' -ApprovalType Manual -ApproverGroup
'mbmanagers@fabrikamorg.onmicrosoft.com'
```

Step 4: Submit/approve privileged access requests

Requesting elevation authorization to execute privileged tasks

Requests for privileged access are valid for up to 24 hours after the request is submitted. If not approved or denied, the requests expire and access is not approved.

In the Microsoft 365 Admin Center

1. Sign into the Microsoft 365 Admin Center using your credentials.
2. In the Admin Center, go to **Settings** > **Org Settings** > **Security & Privacy** > **Privileged access**.
3. Select **Manage access policies and requests**.
4. Select **New request**. From the drop-down fields, select the appropriate values for your organization:

   Request type: Task, Role, or Role Group

   Request scope: Exchange

   Request for: Select from the available policies

   Duration (hours): Number of hours of requested access. There isn't a limit on the number of hours that can be requested.

   Comments: Text field for comments related to your access request
5. Select **Save** and then **Close**. Your request will be sent to the approver's group via email.

In Exchange Management PowerShell

Run the following command in Exchange Online PowerShell to create and submit an approval request to the approver's group:

```
New-ElevatedAccessRequest -Task 'Exchange\<exchange management cmdlet name>' -Reason '<appropriate reason>'
-DurationHours <duration in hours>
```

Example:

```
New-ElevatedAccessRequest -Task 'Exchange\New-MoveRequest' -Reason 'Attempting to fix the user mailbox
error' -DurationHours 4
```

View status of elevation requests

After an approval request is created, elevation request status can be reviewed in the admin center or in Exchange Management PowerShell using the associated with request ID.

In the Microsoft 365 admin center

1. Sign into the Microsoft 365 admin center with your credentials.
2. In the admin center, go to **Settings** > **Org Settings** > **Security & Privacy** > **Privileged access**.
3. Select **Manage access policies and requests**.
4. Select **View** to filter submitted requests by **Pending**, **Approved**, **Denied**, or **Customer Lockbox** status.

In Exchange Management PowerShell

Run the following command in Exchange Online PowerShell to view an approval request status for a specific request ID:

```
Get-ElevatedAccessRequest -Identity <request ID> | select RequestStatus
```

Example:

```
Get-ElevatedAccessRequest -Identity 28560ed0-419d-4cc3-8f5b-603911cbd450 | select RequestStatus
```

Approving an elevation authorization request

When an approval request is created, members of the relevant approver group receive an email notification and can approve the request associated with the request ID. The requestor is notified of the request approval or denial via email message.

In the Microsoft 365 admin center

1. Sign into the Microsoft 365 admin center with your credentials.
2. In the admin center, go to **Settings** > **Org Settings** > **Security & Privacy** > **Privileged access**.
3. Select **Manage access policies and requests**.
4. Select a listed request to view the details and to take action on the request.
5. Select **Approve** to approve the request or select **Deny** to deny the request. Previously approved requests can have access revoked by selecting **Revoke**.

In Exchange Management PowerShell

To approve an elevation authorization request, run the following command in Exchange Online PowerShell:

```
Approve-ElevatedAccessRequest -RequestId <request id> -Comment '<approval comment>'
```

Example:

```
Approve-ElevatedAccessRequest -RequestId a4bc1bdf-00a1-42b4-be65-b6c63d6be279 -Comment '<approval comment>'
```

To deny an elevation authorization request, run the following command in Exchange Online PowerShell:

```
Deny-ElevatedAccessRequest -RequestId <request id> -Comment '<denial comment>'
```

Example:

```
Deny-ElevatedAccessRequest -RequestId a4bc1bdf-00a1-42b4-be65-b6c63d6be279 -Comment '<denial comment>'
```

Delete a privileged access policy in Office 365

If it is no longer needed in your organization, you can delete a privileged access policy.

In the Microsoft 365 admin center

1. Sign into the Microsoft 365 admin center using credentials for an admin account in your organization.
2. In the admin center, go to **Settings** > **Org Settings** > **Security & Privacy** > **Privileged access**.
3. Select **Manage access policies and requests**.
4. Select **Configure policies**.
5. Select the policy you want to delete, then select **Remove Policy**.
6. Select **Close**.

In Exchange Management PowerShell

To delete a privileged access policy, run the following command in Exchange Online Powershell:

```
Remove-ElevatedAccessApprovalPolicy -Identity <identity GUID of the policy you want to delete>
```

Disable privileged access in Office 365

If needed, you can disable privileged access management for your organization. Disabling privileged access does not delete any associated approval policies or approver groups.

In the Microsoft 365 admin center

1. Sign into the Microsoft 365 admin center with credentials for an admin account in your organization.
2. In the Admin Center, go to **Settings** > **Org Settings** > **Security & Privacy** > **Privileged access**.
3. Enable the **Require approvals for privileged access** control.

In Exchange Management PowerShell

To disable privileged access, run the following command in Exchange Online Powershell:

```
Disable-ElevatedAccessControl
```

Protect user and device access

11/2/2020 • 2 minutes to read • Edit Online

Protecting access to your Microsoft 365 data and services is crucial to defending against cyberattacks and guarding against data loss. The same protections can be applied to other SaaS applications in your environment and even to on-premises applications published with Azure Active Directory Application Proxy.

Step 1: Review recommendations

Recommended capabilities for protecting identities and devices that access Office 365, other SaaS services, and on-premises applications published with Azure AD Application Proxy.

PDF | Visio | More languages

Step 2: Protect administrator accounts and access

The administrative accounts you use to administer your Microsoft 365 environment include elevated privileges. These are valuable targets for hackers and cyberattackers.

Begin by using administrator accounts only for administration. Admins should have a separate user account for regular, non-administrative use and only use their administrative account when necessary to complete a task associated with their job function.

Protect your administrator accounts with multi-factor authentication and conditional access. For more information, see Protecting administrator accounts.

Next, configure privileged access management in Office 365. Privileged access management allows granular access control over privileged admin tasks in Office 365. It can help protect your organization from breaches that may use existing privileged admin accounts with standing access to sensitive data or access to critical configuration settings.

- Overview of privileged access management
- Configure privileged access management

Another top recommendation is to use workstations specifically configured for administrative work. These are dedicated devices that are only used for administrative tasks. See Securing privileged access.

Finally, you can mitigate the impact of inadvertent lack of administrative access by creating two or more emergency access accounts in your tenant. See Manage emergency access accounts in Azure AD.

Step 3: Configure recommended identity and device access policies

Multi-factor authentication (MFA) and conditional access policies are powerful tools for mitigating against compromised accounts and unauthorized access. We recommend implementing a set of policies that have been tested together. For more information, including deployment steps, see Identity and device access configurations.

These policies implement the following capabilities:

- Mult-factor authentication
- Conditional access
- Intune app protection (app and data protection for devices)

- Intune device compliance
- Azure AD Identity Protection

Implementing Intune device compliance requires device enrollment. Managing devices allows you to ensure that they are healthy and compliant before allowing them access to resources in your environment. See Enroll devices for management in Intune

Step 4: Configure SharePoint device access policies

Microsoft recommends you protect content in SharePoint sites with sensitive and highly-regulated content with device access controls. For more information, see Policy recommendations for securing SharePoint sites and files.

Customer Lockbox in Office 365

11/2/2020 • 12 minutes to read • Edit Online

This article provides deployment and configuration guidance for Customer Lockbox. Customer Lockbox supports requests to access data in Exchange Online, SharePoint Online, and OneDrive for Business. To recommend support for other services, please submit a request at Office 365 UserVoice.

To see the options for licensing your users to benefit from Microsoft 365 compliance offerings, including this one, as of April 1, 2020, see the Microsoft 365 licensing guidance for security & compliance.

Customer Lockbox ensures that Microsoft cannot access your content to perform a service operation without your explicit approval. Customer Lockbox brings you into the approval workflow for requests to access your content.

Occasionally, Microsoft engineers help troubleshoot and fix customer reported issues in the support process. Usually, issues are fixed through extensive telemetry and debugging tools Microsoft has in place for its services. However, some cases require a Microsoft engineer to access customer content to determine the root cause and fix the issue. Customer Lockbox requires the engineer to request access from the customer as a final step in the approval workflow. This gives organizations the option to approve or deny these requests, and provide direct-access control to the customer.

Customer Lockbox overview video

Customer Lockbox workflow

The following steps outline the typical workflow when a Microsoft engineer initiates a Customer Lockbox request:

1. Someone at an organization experiences an issue with their Microsoft 365 mailbox.
2. After the user troubleshoots the issue, but can't fix it, they open a support request with Microsoft Support.
3. A Microsoft support engineer reviews the service request and determines a need to access the organization's tenant to repair the issue in Exchange Online.
4. The Microsoft support engineer logs into the Customer Lockbox request tool and makes a data access request that includes the organization's tenant name, service request number, and the estimated time the engineer needs access to the data.
5. After a Microsoft Support manager approves the request, Customer Lockbox sends the designated approver at the organization an email notification about the pending access request from Microsoft.

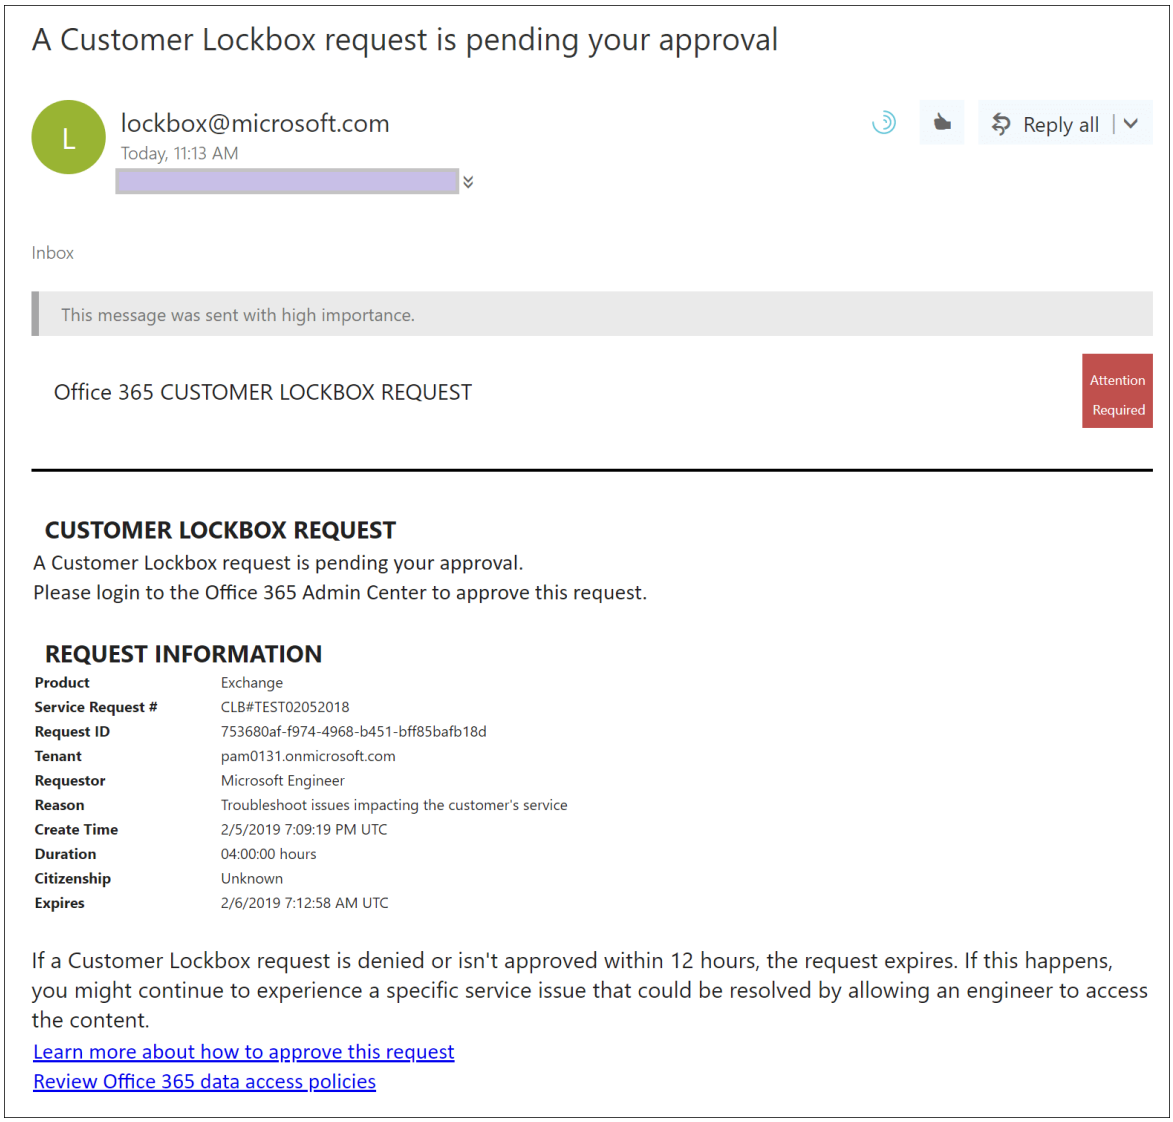

Anyone who is assigned the Customer Lockbox access approver admin role in Microsoft 365 admin center can approve Customer Lockbox requests.

6. The approver signs in to the Microsoft 365 admin center and approves the request. This step triggers the creation of an audit record available by searching the audit log. For more information, see Auditing Customer Lockbox requests.

   If the customer rejects the request or doesn't approve the request within 12 hours, the request expires and no access is granted to the Microsoft engineer.

   > **IMPORTANT**
   > Microsoft does not include any links in Customer Lockbox email notifications requiring you to sign in to Office 365.

7. After the approver from the organization approves the request, the Microsoft engineer receives the approval message, logs into the tenant in Exchange Online, and fixes the customer's issue. Microsoft engineers have the requested duration to fix the issue after which the access is automatically revoked.

> **NOTE**
> All actions performed by a Microsoft engineer are logged in the audit log. You can search for and review these audit records.

Turn Customer Lockbox requests on or off

You can turn on Customer Lockbox controls in the Microsoft 365 admin center. When you turn on Customer Lockbox, Microsoft must obtain your organization's approval before accessing any of your tenant's content.

1. Using a work or school account that has either the global administrator or the **Customer Lockbox access approver** role assigned, go to https://admin.microsoft.com and sign in.
2. Choose **Settings** > **Org Settings**.
3. Select **Security & Privacy** > **Customer Lockbox** > **Edit**, and then move the toggle to **On** or **Off** to turn the feature on or off.

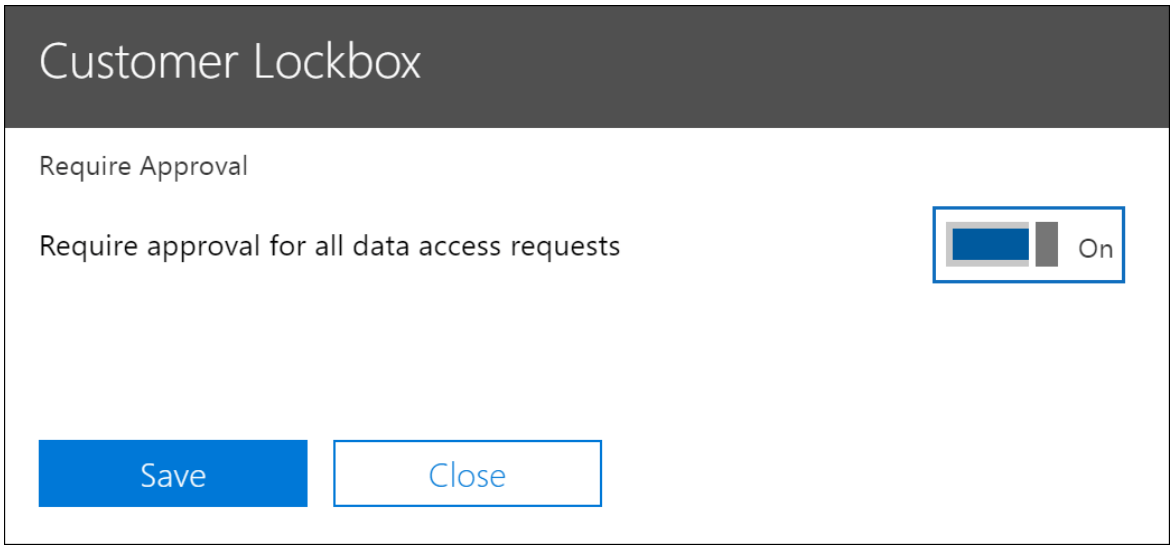

Approve or deny a Customer Lockbox request

1. Using a work or school account that has either the global administrator or the **Customer Lockbox access approver** role assigned, go to https://admin.microsoft.com and sign in.
2. Choose **Support** > **Customer Lockbox Requests**.

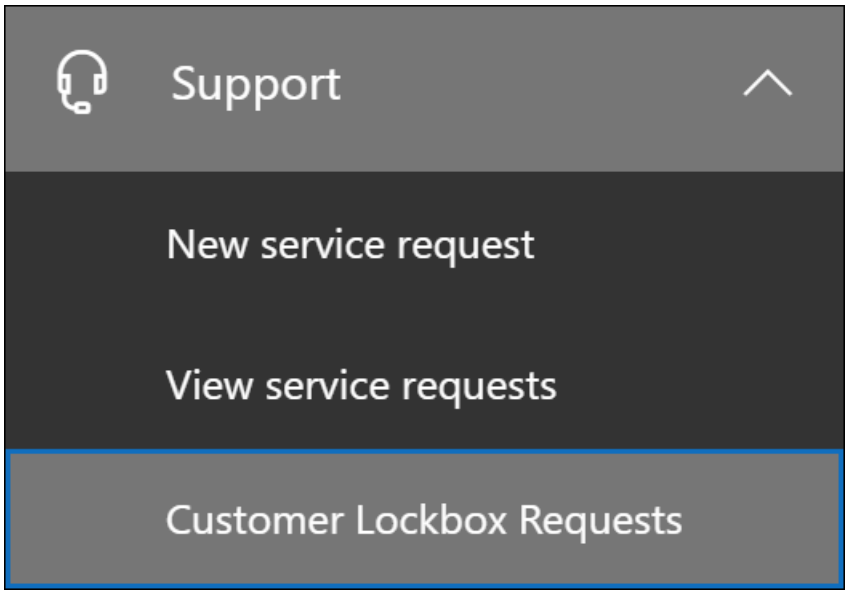

A list of Customer Lockbox requests displays.

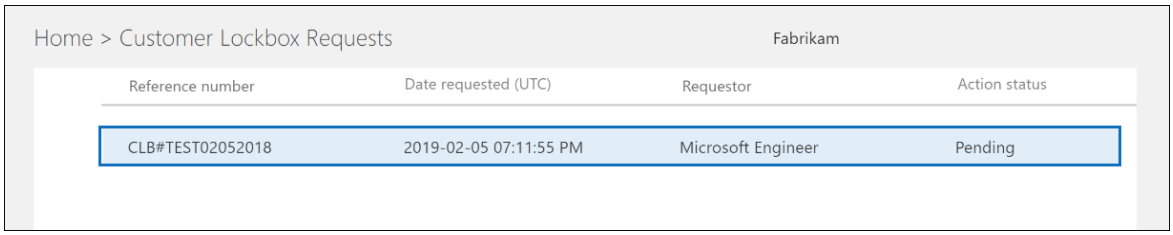

3. Select a Customer Lockbox request, and then choose **Approve** or **Deny**.

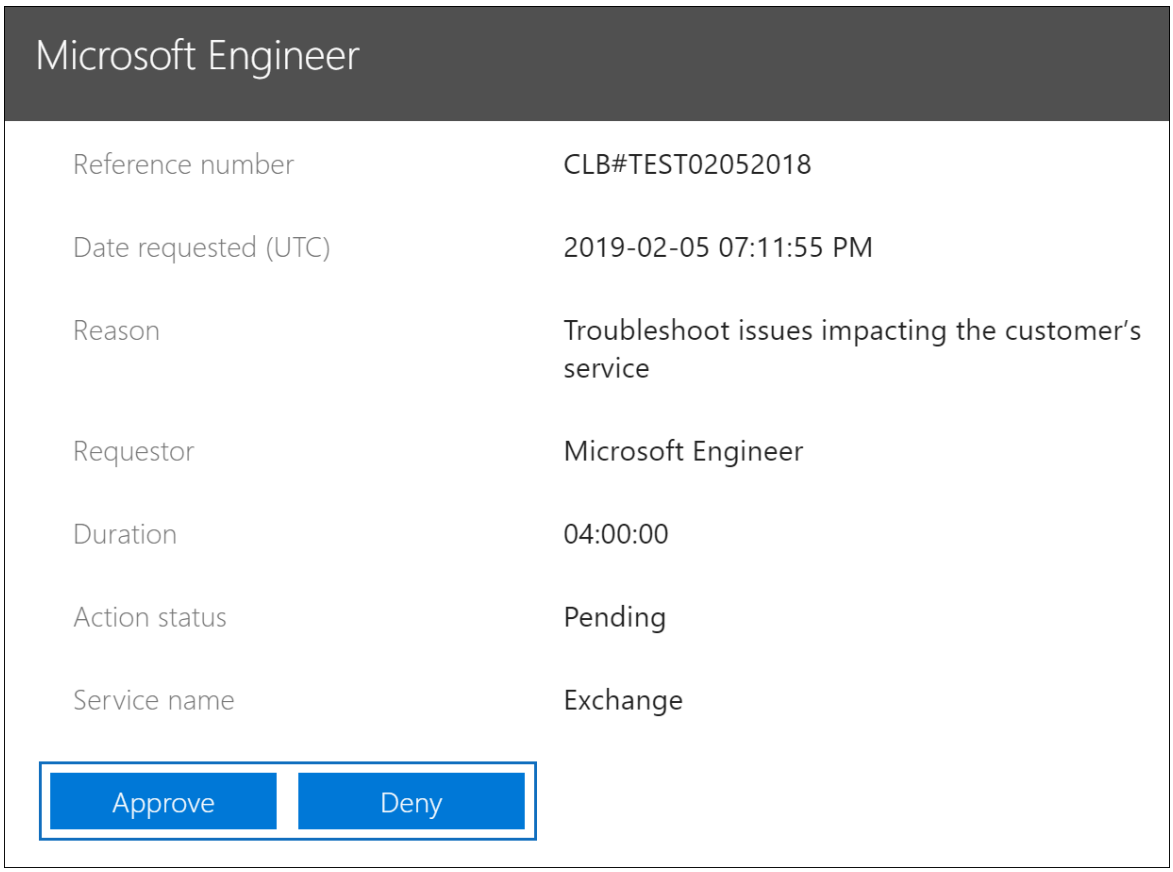

A confirmation message about the approval of the Customer Lockbox request displays.

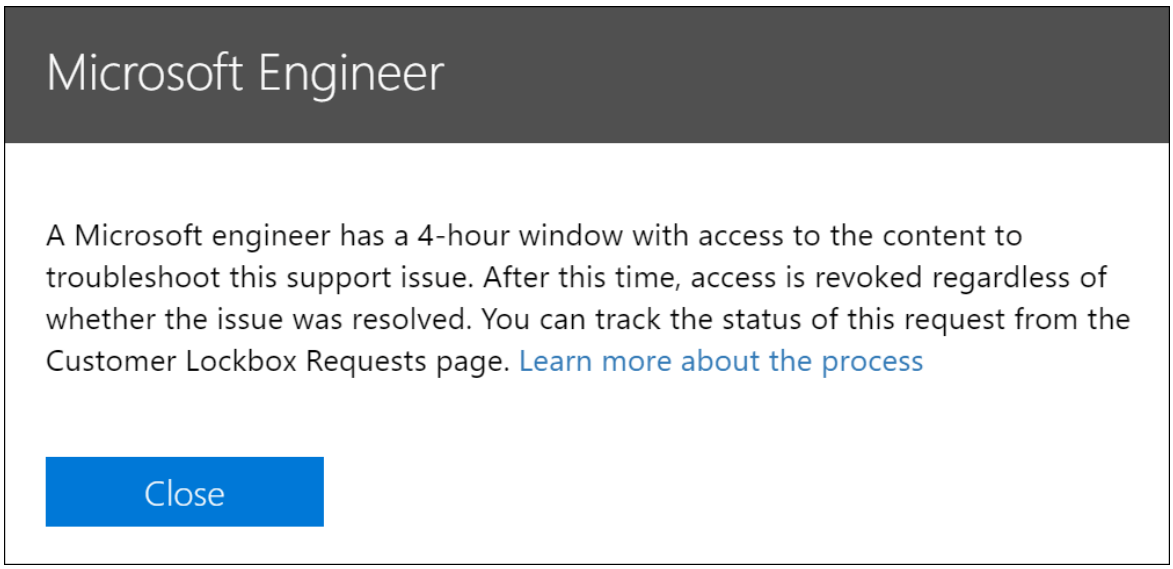

NOTE

Use the Set-AccessToCustomerDataRequest cmdlet to approve, deny, or cancel Microsoft 365 customer lockbox requests that control access to your data by Microsoft support engineers. For more information, see Set-AccessToCustomerDataRequest.

Auditing Customer Lockbox requests

Audit records that correspond to the Customer Lockbox requests are logged in the audit log. You can access these logs by using the audit log search tool in the Security & Compliance Center. Actions related to accepting or denying a Customer Lockbox request and actions performed by Microsoft engineers (when access requests are approved) are also logged in the audit log. You can search for and review these audit records.

Search the audit log for activity related to Customer Lockbox requests

Before you can use the audit log to track requests for Customer Lockbox, there are some steps you need to take to set up audit logging. For more information, see Search the audit log in the Security & Compliance Center. Once you've completed setup, use these steps to create an audit log search query to return audit records related to Customer Lockbox:

1. Go to https://protection.office.com.
2. Sign in using your work or school account.
3. In the left pane of the Security & Compliance Center, choose **Search & investigation** > **Audit log search**.

   The **Audit log search** page displays.

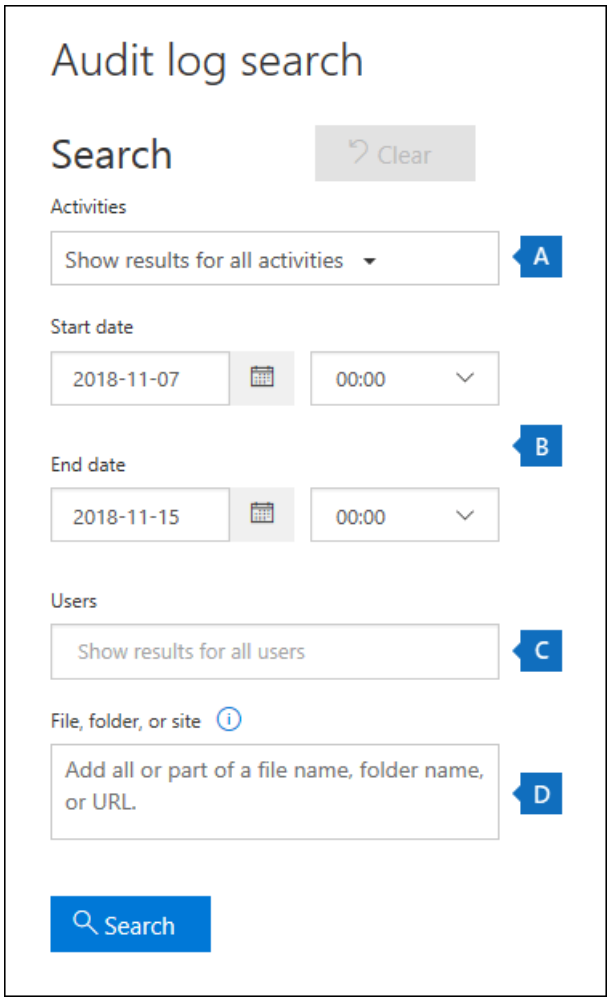

4. Configure the following search criteria:

    a. **Activities** - Leave this field blank so that the search returns audit records for all activities. This is necessary to return any audit records related to Customer Lockbox requests and corresponding activity performed by Microsoft engineers.

    b. **Start date** and **End date** - Select a date and time range to display the events that occurred within that period.

    c. **Users** - Leave this field blank.

    d. **File, folder, or site** - Leave this field blank.

5. Click **Search** to run the search using your search criteria.

    The search results are loaded, and after a few moments they are displayed under **Results** on the **Audit log search** page.

6. Click **Filter results** on the search results page, and do one of the following things:

    - To display audit records related to an approver in your organization approving or denying a Customer Lockbox request: In the box under the **Activity** column, type **Set-AccessToCustomerDataRequest**.
    - To display audit records related to a Microsoft engineer performing actions in response to an approved Customer Lockbox request: In the box under the **User** column, type **Microsoft Operator**. The **Activity** column displays the action performed by the engineer.

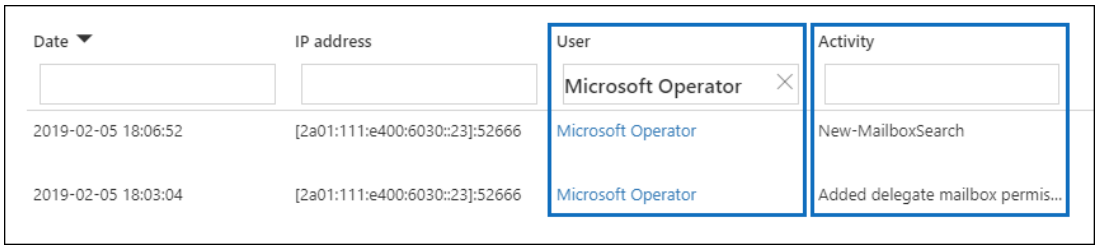

7. In the list of results, click an audit record to display it.

Audit record for a Customer Lockbox access request

When a person in your organization approves or denies a Customer Lockbox request, an audit record is logged in the audit log. This record contains the following information.

AUDIT RECORD PROPERTY	DESCRIPTION
Date	The date and time when the Customer Lockbox request was approved or denied.
IP address	The IP address of the machine the approver used to approve or deny a request.
User	The service account BOXServiceAccount@[customerforest].prod.outlook.com.
Activity	Set-AccessToCustomerDataRequest; this is the auditing activity that is logged when you approve or deny a Customer Lockbox request.
Item	The Guid of the Customer Lockbox request

The following screenshot shows an example of an audit log record that corresponds to an approved Customer Lockbox request. If a Customer Lockbox request was denied, then the value of **ApprovalDecision** parameter would be **Deny**.

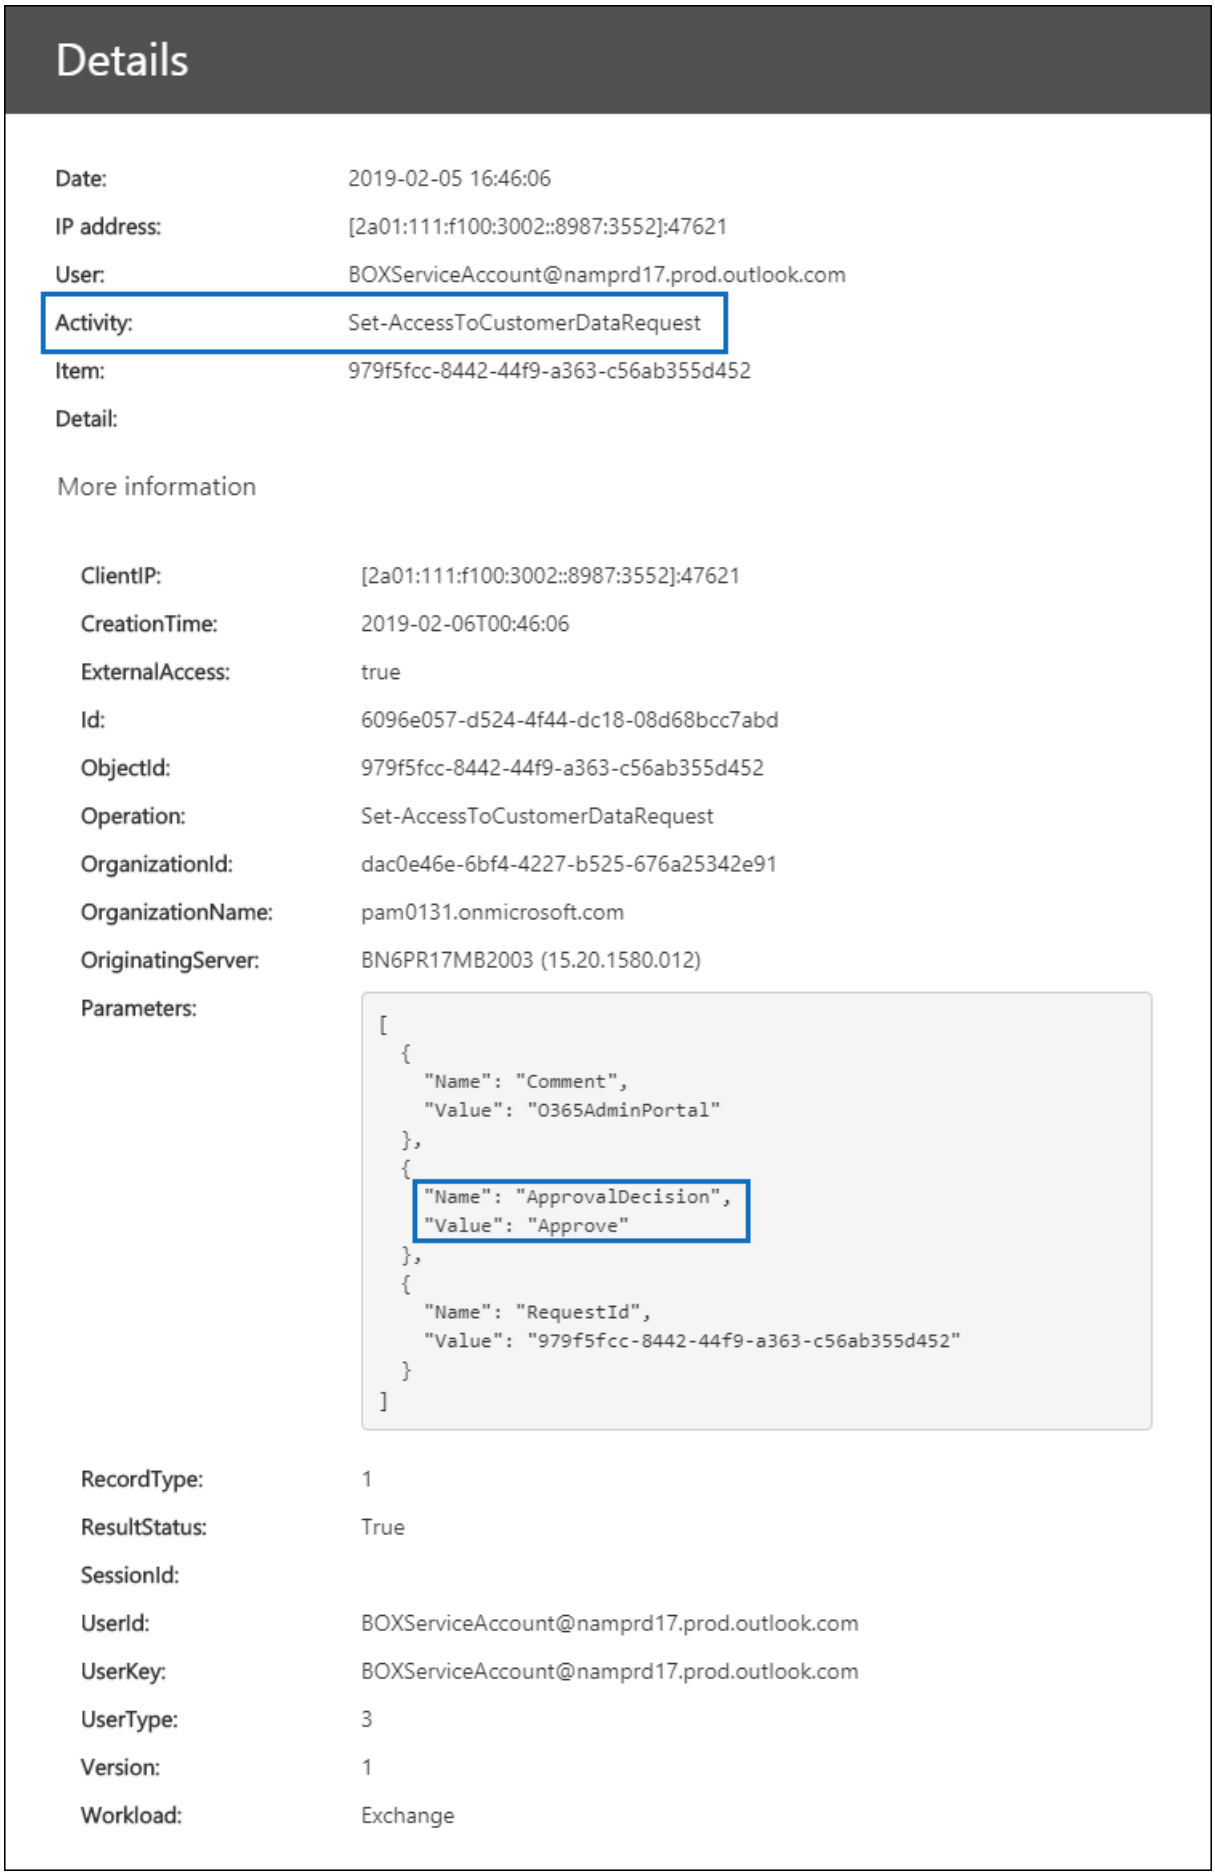

TIP

To display more detailed information in an audit record, click **More information**.

Audit record for an action performed by a Microsoft engineer

The actions performed by a Microsoft engineer after a Customer Lockbox request is approved (and that may result in accessing customer content) are logged in the audit log. These records contain the following information.

AUDIT RECORD PROPERTY	DESCRIPTION
Date	Date time when the action was performed. Note that the time that this action was performed will be within 4 hours of when the Customer Lockbox request was approved.
IP address	The IP Address of the machine Microsoft engineer used.
User	Microsoft Operator; this value indicates that this record is related to a Customer Lockbox request.
Activity	Name of the activity performed by the Microsoft engineer.
Item	<empty>

Frequently asked questions

Which Microsoft 365 services does Customer Lockbox apply to?

Customer Lockbox is currently supported in Exchange Online, SharePoint Online, and OneDrive for Business.

Is Customer Lockbox available to all customers?

Customer Lockbox is included with the Microsoft 365 or Office 365 E5 subscriptions and can be added to other plans with an Information Protection and Compliance or an Advanced Compliance add-on subscription. Please see Plans and pricing for more information.

What is customer content?

Customer content is the data created by users of Microsoft 365 services and applications. Examples of customer content include:

- Email body or email attachments
- SharePoint site contents
- Information in the body of a SharePoint file
- Skype for Business presentation file body
- Instant messages (IM) or voice conversations
- Customer-generated blob or structured storage data (for example, SQL Containers)
- Customer-owned security information (for example, certificates, encryption keys, and passwords)
- Inferences, and all subsequent inferences, if customer content remains

For additional information about customer content in Office 365, see the Office 365 Trust Center.

Who is notified when there is a request to access my content?

Global administrators and anyone assigned the Customer Lockbox access approver admin role are notified.

These are also the same users who can approve for Customer Lockbox requests.

Who can approve or reject these requests in my organization?

Global administrators and anyone assigned the Customer Lockbox access approver admin role can approve Customer Lockbox requests. Customers control these role assignments in their organizations.

How do I opt in to Customer Lockbox?

A global administrator can enable and configure Customer Lockbox in the Microsoft 365 or Microsoft 365 admin center.

If I approve a Customer Lockbox request, what can the engineer do and how will I know what the Microsoft engineer did?

After you approve a Customer Lockbox request, the Microsoft engineer granted these necessary privileges to access customer content by using pre-approved cmdlets. Actions taken by Microsoft engineers in response to Customer Lockbox requests are logged and accessible in the audit log in the Security & Compliance Center.

How do I know that Microsoft follows the approval process?

You can cross-reference the email approval notifications sent to admins and approvers in your organization with the Customer Lockbox request history in the Microsoft 365 admin center.

Customer Lockbox is included in the latest SOC 1 SSAE 16 audit report. For more details, you can find the latest reports in the Microsoft Service Trust Portal.

Can Microsoft modify the list of approvers for my tenant? If not, how is it prevented?

Only a global administrator in your organization can specify who can approve Customer Lockbox requests. That means only the members of the Global administrator group in Azure Active Directory can specify who can approve request. Membership of the Global administrator group in Azure Active Directory is managed only by your organization.

What if I need more information about a content access request to approve it?

Each Customer Lockbox request contains a Microsoft 365 service request number. You can contact Microsoft Support and reference this service number to get more information about the request.

When a Customer Lockbox request is approved, how long are the permissions valid?

Currently, the maximum period for the access permissions granted to the Microsoft engineer is 4 hours. The Microsoft engineer can also request a shorter period.

How can I get a history of all Customer Lockbox requests?

All Customer Lockbox requests are viewed in the Microsoft 365 admin center.

How do I correlate the content access requests with the related audit logs?

The Compliance Center Activity Feed contains log activities of Customer Lockbox. Customers can cross-reference the Customer Lockbox log activities from the activity feed against the email request they receive.

What happens when a customer doesn't respond to a Customer Lockbox request?

Customer Lockbox requests have a default duration of 12 hours. If you don't respond to a request within 12 hour, the request expires.

What does Microsoft do when a customer rejects a Customer Lockbox request?

If a customer rejects a Customer Lockbox request, no access to customer content occurs. If a user in your organization continues to experience a service issue requiring Microsoft to access customer content to resolve the issue, then the service issue might persist and Microsoft will inform the user about this.

Does Customer Lockbox protect against data requests from law enforcement agencies or other third parties?

No. Microsoft takes third-party requests for customer data seriously. As a cloud service provider, Microsoft always advocates for the privacy of customer data. In the event we get a subpoena, Microsoft always attempts to redirect the third party to the customer to obtain the information. (Read Brad Smith's blog: Protecting customer data from government snooping). We periodically publish detailed information about the law enforcement requests that Microsoft receives.

See the Microsoft Trust Center regarding third-party data requests and the "Disclosure of Customer Data" section in the Online Services Terms for more information.

How does Microsoft ensure that a member of its staff doesn't have standing access to customer content in Office 365 applications?

Microsoft implements extensive preventive measures through access control systems, and detective measures to identify and address attempts to circumvent these access control systems. Microsoft 365 operates with the principles of least privilege and just-in-time access. Therefore, no Microsoft personnel have permission to access customer content on an ongoing basis. If permission is granted, it is for a limited duration.

Microsoft 365 uses an access control system called *Lockbox* to process requests for permissions that grant the ability to perform operational and administrative functions within the service. An operator must request access to customer content using Lockbox, which then requires a second person to take action on the request (e.g., approve it) before access is granted. That second person can't be the requestor and must be designated to approve access to customer content. Only if the request is approved does the operator acquire temporary access to customer content. After the elevation period expires, Lockbox revokes access.

Please refer to the Online Services Terms for more details about Microsoft general security practices.

Under what circumstances do Microsoft engineers need access to my content?

The most common scenario where Microsoft engineers need access customer content is when the customer makes a support request requiring access for troubleshooting. A foundational principle of Microsoft 365 is that the service operates without Microsoft access to customer content. Nearly all service operations performed by Microsoft are fully automated and human involvement is highly controlled and abstracted away from customer content. The goal for Microsoft 365 is access to customer content to support the service isn't needed until the customer approves a specific request for Microsoft access.

I already thought my data was secure with the Microsoft cloud, so why do I need Customer Lockbox?

Customer Lockbox provides an extra layer of control by offering customers the ability to give explicit access authorization for service operations. By demonstrating that procedures are in place for explicit data access authorization, Customer Lockbox also helps customers meet certain compliance obligations such as HIPAA and FEDRAMP.

Microsoft Information Protection in Microsoft 365

2/18/2021 • 3 minutes to read • Edit Online

Licensing for Microsoft 365 Security & Compliance

Implement Microsoft Information Protection (MIP) to help you discover, classify, and protect sensitive information wherever it lives or travels.

MIP capabilities are included with Microsoft 365 Compliance and give you the tools to know your data, protect your data, and prevent data loss.

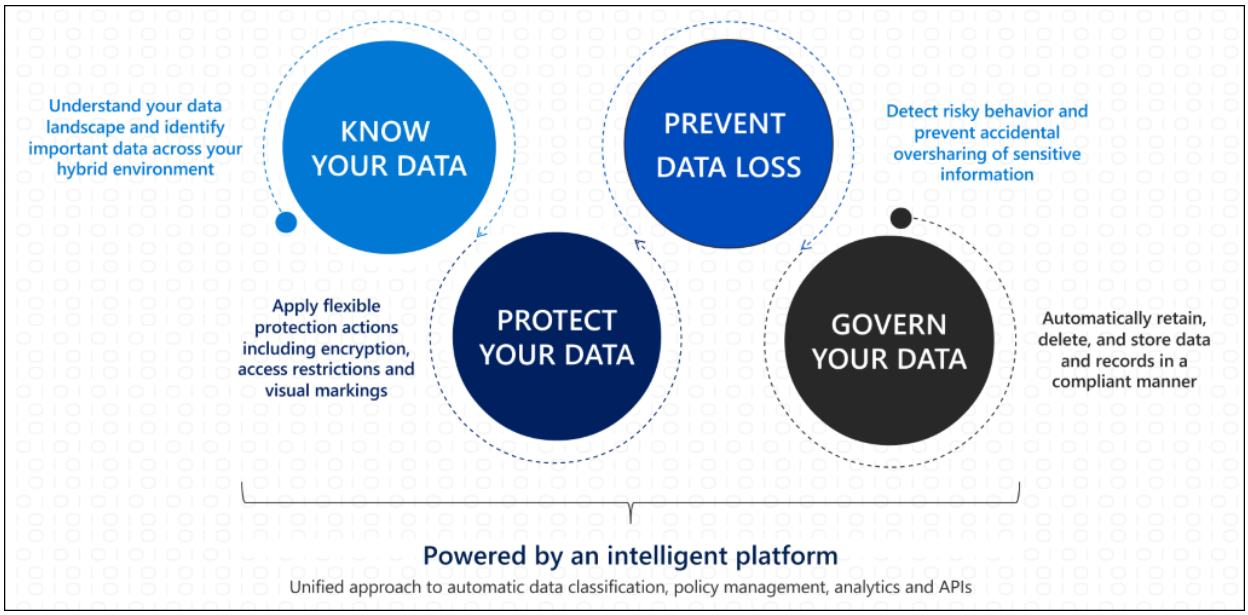

For information about governing your data, see Microsoft Information Governance in Microsoft 365.

Know your data

NOTE

For information about classifying and labeling data in Azure Purview, currently in preview, see Automatically label your content in Azure Purview.

For information about this recent release, see the blog post Microsoft Information Protection and Microsoft Azure Purview: Better Together.

To understand your data landscape and identify important data across your hybrid environment, use the following capabilities:

CAPABILITY	WHAT PROBLEMS DOES IT SOLVE?	GET STARTED
Sensitive information types	Identifies sensitive data by using built-in or custom regular expressions or a function. Corroborative evidence includes keywords, confidence levels, and proximity.	Customize a built-in sensitive information type

CAPABILITY	WHAT PROBLEMS DOES IT SOLVE?	GET STARTED
Trainable classifiers	Identifies sensitive data by using examples of the data you're interested in rather than identifying elements in the item (pattern matching). You can use built-in classifiers or train a classifier with your own content.	Get started with trainable classifiers
Data classification	A graphical identification of items in your organization that have a sensitivity label, a retention label, or have been classified. You can also use this information to gain insights into the actions that your users are taking on these items.	Get started with content explorer Get started with activity explorer

Protect your data

To apply flexible protection actions that include encryption, access restrictions, and visual markings, use the following capabilities:

CAPABILITY	WHAT PROBLEMS DOES IT SOLVE?	GET STARTED
Sensitivity labels	A single solution across apps, services, and devices to label and protect your data as it travels inside and outside your organization. Example scenarios: Manage sensitivity labels for Office apps Encrypt documents and emails Apply and view labels in Power BI For a comprehensive list of scenarios for sensitivity labels, see the Get started documentation.	Get started with sensitivity labels
Azure Information Protection unified labeling client	For Windows computers, extends sensitivity labels for additional features and functionality that includes labeling and protecting all file types from File Explorer and PowerShell Example additional features: Custom configurations for the Azure Information Protection unified labeling client	Azure Information Protection unified labeling client administrator guide
Double Key Encryption	Under all circumstances, only your organization can ever decrypt protected content or for regulatory requirements, you must hold encryption keys within a geographical boundary.	Deploy Double Key Encryption

CAPABILITY	WHAT PROBLEMS DOES IT SOLVE?	GET STARTED
Office 365 Message Encryption (OME)	Encrypts email messages and attached documents that are sent to any user on any device, so only authorized recipients can read emailed information. Example scenario: Revoke email encrypted by Advanced Message Encryption	Set up new Message Encryption capabilities
Service encryption with Customer Key	Protects against viewing of data by unauthorized systems or personnel, and complements BitLocker disk encryption in Microsoft datacenters.	Set up Customer Key for Office 365
SharePoint Information Rights Management (IRM)	Protects SharePoint lists and libraries so that when a user checks out a document, the downloaded file is protected so that only authorized people can view and use the file according to policies that you specify.	Set up Information Rights Management (IRM) in SharePoint admin center
Rights Management connector	Protection-only for existing on-premises deployments that use Exchange or SharePoint Server, or file servers that run Windows Server and File Classification Infrastructure (FCI).	Steps to deploy the RMS connector
Azure Information Protection unified labeling scanner	Discovers, labels, and protects sensitive information that resides in data stores that are on premises.	Configuring and installing the Azure Information Protection unified labeling scanner
Microsoft Cloud App Security	Discovers, labels, and protects sensitive information that resides in data stores that are in the cloud.	Discover, classify, label, and protect regulated and sensitive data stored in the cloud
Microsoft Information Protection SDK	Extends sensitivity labels to third-party apps and services. Example scenario: Set and get a sensitivity label (C++)	Microsoft Information Protection (MIP) SDK setup and configuration

Prevent data loss

To help prevent accidental oversharing of sensitive information, use the following capabilities:

CAPABILITY	WHAT PROBLEMS DOES IT SOLVE?	GET STARTED
Data loss prevention (DLP)	Helps prevent unintentional sharing of sensitive items. Example scenario: Protect sensitive information in Microsoft Teams chat and channel messages	Get started with the default DLP policy

CAPABILITY	WHAT PROBLEMS DOES IT SOLVE?	GET STARTED
Learn about Endpoint data loss prevention	Extends DLP capabilities to items that are used and shared on Windows 10 computers.	Get started with Endpoint data loss prevention

Learn about sensitive information types

2/18/2021 • 4 minutes to read • Edit Online

Identifying and classifying sensitive items that are under your organizations control is the first step in the Information Protection discipline. Microsoft 365 provides three ways of identifying items so that they can be classified:

- manually by users
- automated pattern recognition, like sensitive information types
- machine learning

Sensitive information types are pattern-based classifiers. They detect sensitive information like social security, credit card, or bank account numbers to identify sensitive items, see Sensitive information types entity definitions

Sensitive information types are used in

- Data loss prevention policies
- Sensitivity labels
- Retention labels
- Communication compliance
- Auto-labelling policies

Fundamental parts of a sensitive information type

Every sensitive information type entity is defined by these fields:

- name: how the sensitive information type is referred to
- description: describes what the sensitive information type is looking for
- pattern: A pattern defines what a sensitive information type detects. It consists of the following components
  - Primary element – the main element that the sensitive information type is looking for. It can be a **regular expression** with or without a checksum validation, a **keyword list**, a **keyword dictionary**, or a **function**.
  - Supporting element – elements that act as supporting evidence that help in increasing the confidence of the match. For example, keyword "SSN" in proximity of an SSN number. It can be a regular expression with or without a checksum validation, keyword list, keyword dictionary.
  - Confidence Level - Confidence levels (high, medium, low) reflect how much supporting evidence was detected along with the primary element. The more supporting evidence an item contains, the higher the confidence that a matched item contains the sensitive info you're looking for.
  - Proximity – Number of characters between primary and supporting element

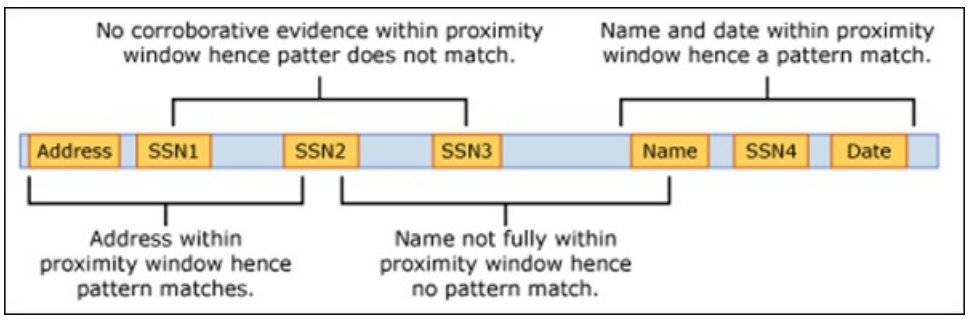

Learn more about confidence levels in this video

Example sensitive information type

Argentina national identity (DNI) number

Format

Eight digits separated by periods

Pattern

Eight digits:

- two digits
- a period
- three digits
- a period
- three digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_argentina_national_id finds content that matches the pattern.
- A keyword from Keyword_argentina_national_id is found.

```
<!-- Argentina National Identity (DNI) Number -->
<Entity id="eefbb00e-8282-433c-8620-8f1da3bffdb2" recommendedConfidence="75" patternsProximity="300">
   <Pattern confidenceLevel="75">
      <IdMatch idRef="Regex_argentina_national_id"/>
      <Match idRef="Keyword_argentina_national_id"/>
  </Pattern>
</Entity>
```

Keywords

Keyword_argentina_national_id

- Argentina National Identity number
- Identity
- Identification National Identity Card
- DNI
- NIC National Registry of Persons
- Documento Nacional de Identidad
- Registro Nacional de las Personas
- Identidad
- Identificación

More on confidence levels

In a sensitive information type entity definition, **confidence level** reflects how much supporting evidence is detected in addition to the primary element. The more supporting evidence an item contains, the higher the confidence that a matched item contains the sensitive info you're looking for. For example, matches with a high

confidence level will contain more supporting evidence in close proximity of the primary element, whereas matches with a low confidence level would contain little to no supporting evidence in close proximity.

A high confidence level returns the fewest false positives but might result in more false negatives. Low or medium confidence levels returns more false positives but few to zero false negatives.

- **low confidence**: value of 65, matched items will contain the fewest false negatives but the most false positives.
- **medium confidence**: value of 75, matched items will contain an average amount of false positives and false negatives.
- **high confidence**: value of 85, matched items will contain the fewest false positives but the most false negatives.

You should use high confidence level patterns with low counts, say five to ten, and low confidence patterns with higher counts, say 20 or more.

Creating custom sensitive information types

To create custom sensitive information types in the Security & Compliance Center, you can choose from several options:

- **Use the UI** You can set up a custom sensitive information type using the Security & Compliance Center UI. With this method, you can use regular expressions, keywords, and keyword dictionaries. To learn more, see Create a custom sensitive information type.
- **Use EDM** You can set up custom sensitive information types using Exact Data Match (EDM)-based classification. This method enables you to create a dynamic sensitive information type using a secure database that you can refresh periodically. See Create a custom sensitive information type with Exact Data Match based classification.
- **Use PowerShell** You can set up custom sensitive information types using PowerShell. Although this method is more complex than using the UI, you have more configuration options. See Create a custom sensitive information type in Security & Compliance Center PowerShell.

> **NOTE**
> Improved confidence levels are available for immediate use within Data Loss Prevention for Microsoft 365 services, Microsoft Information Protection for Microsoft 365 services, Communication Compliance, Information Governance, and Records Management.

> Microsoft 365 Information Protection now supports in preview double byte character set languages for:
>
> - Chinese (simplified)
> - Chinese (traditional)
> - Korean
> - Japanese

> This support is available for sensitive information types. See, Information protection support for double byte character sets release notes (preview) for more information.

For further information

- Sensitive information type entity definitions
- Create a custom sensitive information type

- Create a custom sensitive information type in PowerShell

Sensitive information type entity definitions

2/18/2021 • 192 minutes to read • Edit Online

Data loss prevention (DLP) in the Compliance Center includes many sensitive information types that are ready to use in your DLP policies. This article lists all of these sensitive information types and shows what a DLP policy looks for when it detects each type. A sensitive information type is defined by a pattern that can be identified by a regular expression or a function. Corroborative evidence, like keywords and checksums, can be used to identify a sensitive information type. Confidence level and proximity are also used in the evaluation process.

Sensitive information types require one of these subscriptions:

- Microsoft 365 E3
- Microsoft 365 E5

Sensitive information types are used in:

- Data loss prevention policies
- Sensitivity labels
- Retention labels
- Communication compliance
- Auto-labelling policies

ABA routing number

Format

nine digits that may be in a formatted or unformatted pattern

Pattern

Formatted:

- four digits beginning with 0, 1, 2, 3, 6, 7, or 8
- a hyphen
- four digits
- a hyphen
- a digit

Unformatted: nine consecutive digits beginning with 0, 1, 2, 3, 6, 7, or 8

Checksum

No

Definition

A policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_aba_routing finds content that matches the pattern.
- A keyword from Keyword_ABA_Routing is found.

A DLP policy has low confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_aba_routing finds content that matches the pattern.

```
<!-- ABA Routing Number -->
<Entity id="cb353f78-2b72-4c3c-8827-92ebe4f69fdf" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_aba_routing" />
    <Match idRef="Keyword_ABA_Routing" />
  </Pattern>
  <Pattern confidenceLevel="65">
    <IdMatch idRef="Func_aba_routing" />
  </Pattern>
</Entity>
```

Keywords

Keyword_aba_routing

- aba number
- aba#
- aba
- abarouting#
- abaroutingnumber
- americanbankassociationrouting#
- americanbankassociationroutingnumber
- bankrouting#
- bankroutingnumber
- routing #
- routing no
- routing number
- routing transit number
- routing#
- RTN

Argentina national identity (DNI) number

Format

Eight digits with or without periods

Pattern

Eight digits:

- two digits
- an optional period
- three digits
- an optional period
- three digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_argentina_national_id finds content that matches the pattern.
- A keyword from Keyword_argentina_national_id is found.

```
<!-- Argentina National Identity (DNI) Number -->
<Entity id="eefbb00e-8282-433c-8620-8f1da3bffdb2" recommendedConfidence="75" patternsProximity="300">
   <Pattern confidenceLevel="75">
      <IdMatch idRef="Regex_argentina_national_id"/>
      <Match idRef="Keyword_argentina_national_id"/>
  </Pattern>
</Entity>
```

Keywords

Keyword_argentina_national_id

- Argentina National Identity number
- cedula
- cédula
- dni
- documento nacional de identidad
- documento número
- documento numero
- registro nacional de las personas
- rnp

Australia bank account number

Format

six to ten digits with or without a bank state branch number

Pattern

Account number is 6 to 10 digits.

Australia bank state branch number:

- three digits
- a hyphen
- three digits

Checksum

No

Definition

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_australia_bank_account_number finds content that matches the pattern.
- A keyword from Keyword_australia_bank_account_number is found.
- The regular expression Regex_australia_bank_account_number_bsb finds content that matches the pattern.

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_australia_bank_account_number finds content that matches the pattern.
- A keyword from Keyword_australia_bank_account_number is found.

```
<!-- Australia Bank Account Number -->
<Entity id="74a54de9-2a30-4aa0-a8aa-3d9327fc07c7" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
        <IdMatch idRef="Regex_australia_bank_account_number" />
        <Match idRef="Keyword_australia_bank_account_number" />
        <Match idRef="Regex_australia_bank_account_number_bsb" />
  </Pattern>
  <Pattern confidenceLevel="75">
        <IdMatch idRef="Regex_australia_bank_account_number" />
        <Match idRef="Keyword_australia_bank_account_number" />
  </Pattern>
 </Entity>
```

Keywords

Keyword_australia_bank_account_number

- swift bank code
- correspondent bank
- base currency
- usa account
- holder address
- bank address
- information account
- fund transfers
- bank charges
- bank details
- banking information
- full names
- iaea

Australia business number

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

11 digits with optional delimiters

Pattern

11 digits with optional delimiters:

- two digits
- an optional hyphen or space
- three digits
- an optional hyphen or space
- three digits
- an optional hyphen or space
- three digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_australian_business_number finds content that matches the pattern.
- A keyword from Keywords_australian_business_number is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_australian_business_number finds content that matches the pattern.

```
<!-- Australia Business Number -->
<Entity id="76e83b3b-01ee-4530-aced-e667a6609f49" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_australian_business_number" />
    <Match idRef="Keywords_australian_business_number" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_australian_business_number" />
  </Pattern>
</Entity>
```

Keywords

Keyword_australia_business_number

- australia business no
- business number
- abn#
- businessid#
- business id
- abn
- businessno#

Australia company number

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

nine digits with delimiters

Pattern

nine digits with delimiters:

- three digits
- a space
- three digits

- a space
- three digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_Australian_Company_Number finds content that matches the pattern.
- A keyword from Keyword_Australian_Company_Number is found.

A DLP policy has low confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_Australian_Company_Number finds content that matches the pattern.

```
<!-- Australia Company Number -->
<Entity id="b1fba4f7-7b3e-4bb9-8f9a-9366df604dbb" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_Australian_Company_Number" />
    <Match idRef="Keyword_Australian_Company_Number" />
  </Pattern>
  <Pattern confidenceLevel="65">
    <IdMatch idRef="Func_Australian_Company_Number" />
  </Pattern>
</Entity>
```

Keywords

Keyword_australia_company_number

- acn
- australia company no
- australia company no#
- australia company number
- australian company no
- australian company no#
- australian company number

Australia driver's license number

Format

nine letters and digits

Pattern

nine letters and digits:

- two digits or letters (not case-sensitive)
- two digits
- five digits or letters (not case-sensitive)

OR

- one to two optional letters (not case-sensitive)
- four to nine digits

OR

- nine digits or letters (not case-sensitive)

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_australia_drivers_license_number finds content that matches the pattern.
- A keyword from Keyword_australia_drivers_license_number is found.
- No keyword from Keyword_australia_drivers_license_number_exclusions is found.

```
<!-- Australia Drivers License Number -->
<Entity id="1cbbc8f5-9216-4392-9eb5-5ac2298d1356" patternsProximity="300" recommendedConfidence="75">
   <Pattern confidenceLevel="75">
        <IdMatch idRef="Regex_australia_drivers_license_number" />
        <Match idRef="Keyword_australia_drivers_license_number" />
        <Any minMatches="0" maxMatches="0">
          <Match idRef="Keyword_australia_drivers_license_number_exclusions" />
        </Any>
  </Pattern>
</Entity>
```

Keywords

Keyword_australia_drivers_license_number

- international driving permits
- australian automobile association
- international driving permit
- DriverLicence
- DriverLicences
- Driver Lic
- Driver Licence
- Driver Licences
- DriversLic
- DriversLicence
- DriversLicences
- Drivers Lic
- Drivers Lics
- Drivers Licence
- Drivers Licences
- Driver'Lic
- Driver'Lics
- Driver'Licence
- Driver'Licences
- Driver' Lic
- Driver' Lics
- Driver' Licence
- Driver' Licences
- Driver'sLic

- Driver'sLics
- Driver'sLicence
- Driver'sLicences
- Driver's Lic
- Driver's Lics
- Driver's Licence
- Driver's Licences
- DriverLic#
- DriverLics#
- DriverLicence#
- DriverLicences#
- Driver Lic#
- Driver Lics#
- Driver Licence#
- Driver Licences#
- DriversLic#
- DriversLics#
- DriversLicence#
- DriversLicences#
- Drivers Lic#
- Drivers Lics#
- Drivers Licence#
- Drivers Licences#
- Driver'Lic#
- Driver'Lics#
- Driver'Licence#
- Driver'Licences#
- Driver' Lic#
- Driver' Lics#
- Driver' Licence#
- Driver' Licences#
- Driver'sLic#
- Driver'sLics#
- Driver'sLicence#
- Driver'sLicences#
- Driver's Lic#
- Driver's Lics#
- Driver's Licence#
- Driver's Licences#

Keyword_australia_drivers_license_number_exclusions

- aaa
- DriverLicense
- DriverLicenses
- Driver License
- Driver Licenses
- DriversLicense

- DriversLicenses
- Drivers License
- Drivers Licenses
- Driver'License
- Driver'Licenses
- Driver' License
- Driver' Licenses
- Driver'sLicense
- Driver'sLicenses
- Driver's License
- Driver's Licenses
- DriverLicense#
- DriverLicenses#
- Driver License#
- Driver Licenses#
- DriversLicense#
- DriversLicenses#
- Drivers License#
- Drivers Licenses#
- Driver'License#
- Driver'Licenses#
- Driver' License#
- Driver' Licenses#
- Driver'sLicense#
- Driver'sLicenses#
- Driver's License#
- Driver's Licenses#

Australia medical account number

Format

10-11 digits

Pattern

10-11 digits:

- First digit is in the range 2-6
- Nine digit is a check digit
- Tenth digit is the issue digit
- Eleventh digit (optional) is the individual number

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_australian_medical_account_number finds content that matches the pattern.
- A keyword from Keyword_Australia_Medical_Account_Number is found.

- The checksum passes.

```
  <!-- Australia Medical Account Number -->
<Entity id="104a99a0-3d3b-4542-a40d-ab0b9e1efe63" recommendedConfidence="85" patternsProximity="300">
    <Pattern confidenceLevel="85">
     <IdMatch idRef="Func_australian_medical_account_number"/>
     <Match idRef="Keyword_Australia_Medical_Account_Number"/>
    </Pattern>
</Entity>
```

Keywords

Keyword_Australia_Medical_Account_Number

- bank account details
- medicare payments
- mortgage account
- bank payments
- information branch
- credit card loan
- department of human services
- local service
- medicare

Australia passport number

Format

A letter followed by seven digits

Pattern

A letter (not case-sensitive) followed by seven digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_australia_passport_number finds content that matches the pattern.
- A keyword from Keyword_passport or Keyword_australia_passport_number is found.

```
<!-- Australia Passport Number -->
<Entity id="29869db6-602d-4853-ab93-3484f905df50" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
        <IdMatch idRef="Regex_australia_passport_number" />
        <Any minMatches="1">
          <Match idRef="Keyword_passport" />
          <Match idRef="Keyword_australia_passport_number" />
        </Any>
   </Pattern>
</Entity>
```

Keywords

Keyword_passport

- Passport Number

- Passport No
- Passport #
- Passport#
- PassportID
- Passportno
- passportnumber
- パスポート
- パスポート番号
- パスポートのNum
- パスポート ＃
- Numéro de passeport
- Passeport n °
- Passeport Non
- Passeport #
- Passeport#
- PasseportNon
- Passeportn °

Keyword_australia_passport_number

- passport
- passport details
- immigration and citizenship
- commonwealth of australia
- department of immigration
- residential address
- department of immigration and citizenship
- visa
- national identity card
- passport number
- travel document
- issuing authority

Australia tax file number

Format

eight to nine digits

Pattern

eight to nine digits typically presented with spaces as follows:

- three digits
- an optional space
- three digits
- an optional space
- two to three digits where the last digit is a check digit

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_australian_tax_file_number finds content that matches the pattern.
- No keyword from Keyword_Australia_Tax_File_Number or Keyword_number_exclusions is found.
- The checksum passes.

```
<!-- Australia Tax File Number -->
 <Entity id="e29bc95f-ff70-4a37-aa01-04d17360a4c5" patternsProximity="300" recommendedConfidence="85">
   <Pattern confidenceLevel="85">
     <IdMatch idRef="Func_australian_tax_file_number" />
     <Match idRef="Keyword_Australia_Tax_File_Number" />
   </Pattern>
 </Entity>
```

Keywords

Keyword_australia_tax_file_number

- australian business number
- marginal tax rate
- medicare levy
- portfolio number
- service veterans
- withholding tax
- individual tax return
- tax file number
- tfn

Austria driver's license number

Format

eight digits without spaces and delimiters

Pattern

eight digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_austria_eu_driver's_license_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_driver's_license_number` or `Keywords_austria_eu_driver's_license_number` is found.

```
<!-- Austria Driver's License Number -->
<Entity id="682f18ce-44eb-482b-8198-2bcb96a0761e" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_austria_eu_driver's_license_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_driver's_license_number" />
      <Match idRef="Keywords_austria_eu_driver's_license_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_driver's_license_number

- driverlic
- driverlics
- driverlicense
- driverlicenses
- driverlicence
- driverlicences
- driver lic
- driver lics
- driver license
- driver licenses
- driver licence
- driver licences
- driverslic
- driverslics
- driverslicence
- driverslicences
- driverslicense
- driverslicenses
- drivers lic
- drivers lics
- drivers license
- drivers licenses
- drivers licence
- drivers licences
- driver'lic
- driver'lics
- driver'license
- driver'licenses
- driver'licence
- driver'licences
- driver' lic
- driver' lics
- driver' license
- driver' licenses
- driver' licence
- driver' licences

- driver'slic
- driver'slics
- driver'slicense
- driver'slicenses
- driver'slicence
- driver'slicences
- driver's lic
- driver's lics
- driver's license
- driver's licenses
- driver's licence
- driver's licences
- dl#
- dls#
- driverlic#
- driverlics#
- driverlicense#
- driverlicenses#
- driverlicence#
- driverlicences#
- driver lic#
- driver lics#
- driver license#
- driver licenses#
- driver licences#
- driverslic#
- driverslics#
- driverslicense#
- driverslicenses#
- driverslicence#
- driverslicences#
- drivers lic#
- drivers lics#
- drivers license#
- drivers licenses#
- drivers licence#
- drivers licences#
- driver'lic#
- driver'lics#
- driver'license#
- driver'licenses#
- driver'licence#
- driver'licences#
- driver' lic#
- driver' lics#
- driver' license#

- driver' licenses#
- driver' licence#
- driver' licences#
- driver'slic#
- driver'slics#
- driver'slicense#
- driver'slicenses#
- driver'slicence#
- driver'slicences#
- driver's lic#
- driver's lics#
- driver's license#
- driver's licenses#
- driver's licence#
- driver's licences#
- driving licence
- driving license
- dlno#
- driv lic
- driv licen
- driv license
- driv licenses
- driv licence
- driv licences
- driver licen
- drivers licen
- driver's licen
- driving lic
- driving licen
- driving licenses
- driving licence
- driving licences
- driving permit
- dl no
- dlno
- dl number

Keywords_austria_eu_driver's_license_number

- fuhrerschein
- führerschein
- Führerscheine
- Führerscheinnummer
- Führerscheinnummern

Austria identity card

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

A 24-character combination of letters, digits, and special characters

Pattern

24 characters:

- 22 letters (not case-sensitive), digits, backslashes, forward slashes, or plus signs
- two letters (not case-sensitive), digits, backslashes, forward slashes, plus signs, or equal signs

Checksum

Not applicable

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_austria_eu_national_id_card` finds content that matches the pattern.
- A keyword from `Keywords_austria_eu_national_id_card` is found.

```
<!-- Austria Identity Card -->
<Entity id="5ec06c3b-007e-4820-8343-7ff73b889735" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_austria_eu_national_id_card" />
    <Match idRef="Keywords_austria_eu_national_id_card" />
  </Pattern>
</Entity>
```

Keywords

Keywords_austria_eu_national_id_card

- identity number
- national id
- personalausweis republik österreich

Austria passport number

Format

One letter followed by an optional space and seven digits

Pattern

A combination of one letter, seven digits, and one space:

- one letter (not case-sensitive)
- one space (optional)
- seven digits

Checksum

not applicable

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_austria_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_austria_eu_passport_number` is found.
- The regular expression `Regex_eu_passport_date1` finds date in the format DD.MM.YYYY or a keyword from `Keywords_eu_passport_date` is found

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_austria_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_austria_eu_passport_number` is found.

```
<!-- Austria Passport Number -->
<Entity id="1c96ae4e-303b-447d-86c7-77113ac266bf" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Regex_austria_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_austria_eu_passport_number" />
    </Any>
    <Any minMatches="1">
      <Match idRef="Regex_eu_passport_date1" />
      <Match idRef="Keywords_eu_passport_date" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_austria_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_austria_eu_passport_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_passport_number

- passport#
- passport #
- passportid
- passports
- passportno
- passport no
- passportnumber
- passport number
- passportnumbers
- passport numbers

Keywords_austria_eu_passport_number

- reisepassnummer
- reisepasse
- No-Reisepass
- Nr-Reisepass
- Reisepass-Nr

- Passnummer
- reisepässe

Keywords_eu_passport_date

- date of issue
- date of expiry

Austria social security number

Format

10 digits in the specified format

Pattern

10 digits:

- three digits that correspond to a serial number
- one check digit
- six digits that correspond to the birth date (DDMMYY)

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_austria_eu_ssn_or_equivalent` finds content that matches the pattern.
- a keyword from `Keywords_austria_eu_ssn_or_equivalent` is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_austria_eu_ssn_or_equivalent` finds content that matches the pattern.

```
<!-- Austria Social Security Number -->
<Entity id="6896a906-86c9-4d19-a2da-6e43ccd19b7b" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_austria_eu_ssn_or_equivalent" />
    <Match idRef="Keywords_austria_eu_ssn_or_equivalent" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_austria_eu_ssn_or_equivalent" />
    <Any minMatches="0" maxMatches="0">
      <Match idRef="Keywords_austria_eu_telephone_number" />
      <Match idRef="Keywords_austria_eu_mobile_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_austria_eu_ssn_or_equivalent

- austrian ssn
- ehic number
- ehic no
- insurance code
- insurancecode#

- insurance number
- insurance no
- krankenkassennummer
- krankenversicherung
- socialsecurityno
- socialsecurityno#
- social security no
- social security number
- social security code
- sozialversicherungsnummer
- sozialversicherungsnummer#
- soziale sicherheit kein
- sozialesicherheitkein#
- ssn#
- ssn
- versicherungscode
- versicherungsnummer
- zdravstveno zavarovanje

Austria tax identification number

Format

nine digits with optional hyphen and forward slash

Pattern

nine digits with optional hyphen and forward slash:

- two digits
- a hyphen (optional)
- three digits
- a forward slash (optional)
- four digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_austria_eu_tax_file_number` finds content that matches the pattern.
- A keyword from `Keywords_austria_eu_tax_file_number` is found.

A DLP policy has low confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_austria_eu_tax_file_number` finds content that matches the pattern.

```
<!-- Austria Tax Identification Number -->
<Entity id="4fd58d22-af28-4451-b18a-6f722430a56d" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_austria_eu_tax_file_number" />
    <Match idRef="Keywords_austria_eu_tax_file_number" />
  </Pattern>
  <Pattern confidenceLevel="65">
    <IdMatch idRef="Func_austria_eu_tax_file_number" />
  </Pattern>
</Entity>
```

Keywords

Keywords_austria_eu_tax_file_number

- österreich
- st.nr.
- steuernummer
- tax id
- tax identification no
- tax identification number
- tax no#
- tax no
- tax number
- tax registration number
- taxid#
- taxidno#
- taxidnumber#
- taxno#
- taxnumber#
- taxnumber
- tin id
- tin no
- tin#
- tax number

Austria value added tax

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

11-character alphanumeric pattern

Pattern

11-character alphanumeric pattern:

- A or a
- T or t

- Optional space
- U or u
- optional space
- two or three digits
- optional space
- four digits
- optional space
- one or two digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_Austria_Value_Added_Tax finds content that matches the pattern.
- A keyword from Keyword_Austria_Value_Added_Tax is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_Austria_Value_Added_Tax finds content that matches the pattern.

```
<!-- Austria Value Added Tax -->
<Entity id="b6a3eda2-c56c-4b69-a5f7-dca34db00f48" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_Austria_Value_Added_Tax" />
    <Match idRef="Keyword_Austria_Value_Added_Tax" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_Austria_Value_Added_Tax" />
  </Pattern>
</Entity>
```

Keywords

Keyword_austria_value_added_tax

- vat number
- vat#
- austrian vat number
- vat no.
- vatno#
- value added tax number
- austrian vat
- mwst
- umsatzsteuernummer
- mwstnummer
- ust.-identifikationsnummer
- umsatzsteuer-identifikationsnummer
- vat identification number
- atu number
- uid number

Azure DocumentDB auth key

Format

The string "DocumentDb" followed by the characters and strings outlined in the pattern below.

Pattern

- The string "DocumentDb"
- Any combination of between 3-200 lower- or uppercase letters, digits, symbols, special characters, or spaces
- A greater than symbol (>), an equal sign (=), a quotation mark ("), or an apostrophe (')
- Any combination of 86 lower- or uppercase letters, digits, forward slash (/), or plus sign (+)
- Two equal signs (=)

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression CEP_Regex_AzureDocumentDBAuthKey finds content that matches the pattern.
- The regular expression CEP_CommonExampleKeywords doesn't find content that matches the pattern.

```
<!-- Azure Document DB Auth Key -->
<Entity id="0f587d92-eb28-44a9-bd1c-90f2892b47aa" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
        <IdMatch idRef="CEP_Regex_AzureDocumentDBAuthKey" />
        <Any minMatches="0" maxMatches="0">
            <Match idRef="CEP_CommonExampleKeywords" />
          </Any>
  </Pattern>
</Entity>
```

Keywords

CEP_CommonExampleKeywords

(Technically, this sensitive information type identifies these keywords by using a regular expression, not a keyword list.)

- contoso
- fabrikam
- northwind
- sandbox
- onebox
- localhost
- 127.0.0.1
- testacs.com
- s-int.net

Azure IAAS database connection string and Azure SQL connection string

Format

The string "Server", "server", or "data source" followed by the characters and strings outlined in the pattern below, including the string "cloudapp.azure.com" or "cloudapp.azure.net" or "database.windows.net", and the

string "Password" or "password" or "pwd".

Pattern

- the string "Server", "server", or "data source"
- zero to two whitespace characters
- an equal sign (=)
- zero to two whitespace characters
- any combination of between 1-200 lower- or uppercase letters, digits, symbols, special characters, or spaces
- The string "cloudapp.azure.com", "cloudapp.azure.net", or "database.windows.net"
- any combination of between 1-300 lower- or uppercase letters, digits, symbols, special characters, or spaces
- the string "Password", "password", or "pwd"
- zero to two whitespace characters
- an equal sign (=)
- zero to two whitespace characters
- one or more characters that aren't a semicolon (;), quotation mark ("), or apostrophe (')
- a semicolon (;), quotation mark ("), or apostrophe (')

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression CEP_Regex_AzureConnectionString finds content that matches the pattern.
- The regular expression CEP_CommonExampleKeywords doesn't find content that matches the pattern.

```
<!--Azure IAAS Database Connection String and Azure SQL Connection String-->
<Entity id="ce1a126d-186f-4700-8c0c-486157b953fd" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
        <IdMatch idRef="CEP_Regex_AzureConnectionString" />
        <Any minMatches="0" maxMatches="0">
            <Match idRef="CEP_CommonExampleKeywords" />
        </Any>
    </Pattern>
</Entity>
```

Keywords

CEP_common_example_keywords

(Technically, this sensitive information type identifies these keywords by using a regular expression, not a keyword list.)

- contoso
- fabrikam
- northwind
- sandbox
- onebox
- localhost
- 127.0.0.1
- testacs.com
- s-int.net

Azure IoT connection string

Format

The string "HostName" followed by the characters and strings outlined in the pattern below, including the strings "azure-devices.net" and "SharedAccessKey".

Pattern

- the string "HostName"
- zero to two whitespace characters
- an equal sign (=)
- zero to two whitespace characters
- any combination of between 1-200 lower- or uppercase letters, digits, symbols, special characters, or spaces
- the string "azure-devices.net"
- any combination of between 1-200 lower- or uppercase letters, digits, symbols, special characters, or spaces
- the string "SharedAccessKey"
- zero to two whitespace characters
- an equal sign (=)
- zero to two whitespace characters
- any combination of 43 lower- or uppercase letters, digits, forward slash (/), or plus sign (+)
- an equal sign (=)

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression CEP_Regex_AzureIoTConnectionString finds content that matches the pattern.
- The regular expression CEP_CommonExampleKeywords doesn't find content that matches the pattern.

```
<!--Azure IoT Connection String-->
<Entity id="0b34bec3-d5d6-4974-b7b0-dcdb5c90c29d" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
        <IdMatch idRef="CEP_Regex_AzureIoTConnectionString" />
        <Any minMatches="0" maxMatches="0">
            <Match idRef="CEP_CommonExampleKeywords" />
        </Any>
  </Pattern>
</Entity>
```

Keywords

CEP_common_example_keywords

(Technically, this sensitive information type identifies these keywords by using a regular expression, not a keyword list.)

- contoso
- fabrikam
- northwind
- sandbox
- onebox
- localhost

- 127.0.0.1
- testacs.com
- s-int.net

Azure publish setting password

Format

The string "userpwd=" followed by an alphanumeric string.

Pattern

- the string "userpwd="
- any combination of 60 lowercase letters or digits
- a quotation mark (")

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression CEP_Regex_AzurePublishSettingPasswords finds content that matches the pattern.
- The regular expression CEP_CommonExampleKeywords doesn't find content that matches the pattern.

```
<!--Azure Publish Setting Password-->
<Entity id="75f4cc8a-a68e-49e5-89ce-fa8f03d286a5" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
       <IdMatch idRef="CEP_Regex_AzurePublishSettingPasswords" />
       <Any minMatches="0" maxMatches="0">
           <Match idRef="CEP_CommonExampleKeywords" />
       </Any>
  </Pattern>
</Entity>
```

Keywords

CEP_common_example_keywords

(Technically, this sensitive information type identifies these keywords by using a regular expression, not a keyword list.)

- contoso
- fabrikam
- northwind
- sandbox
- onebox
- localhost
- 127.0.0.1
- testacs.com
- s-int.net

Azure Redis cache connection string

Format

The string "redis.cache.windows.net" followed by the characters and strings outlined in the pattern below,

including the string "password" or "pwd".

Pattern

- the string "redis.cache.windows.net"
- any combination of between 1-200 lower- or uppercase letters, digits, symbols, special characters, or spaces
- the string "password" or "pwd"
- zero to two whitespace characters
- an equal sign (=)
- zero to two whitespace characters
- any combination of 43 characters that are lower- or uppercase letters, digits, forward slash (/), or plus sign (+)
- an equal sign (=)

Checksum

No

Definition

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression CEP_Regex_AzureRedisCacheConnectionString finds content that matches the pattern.
- The regular expression CEP_CommonExampleKeywords doesn't find content that matches the pattern.

```
<!--Azure Redis Cache Connection String-->
<Entity id="095a7e6c-efd8-46d5-af7b-5298d53a49fc" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
        <IdMatch idRef="CEP_Regex_AzureRedisCacheConnectionString" />
        <Any minMatches="0" maxMatches="0">
            <Match idRef="CEP_CommonExampleKeywords" />
        </Any>
  </Pattern>
</Entity>
```

Keywords

CEP_common_example_keywords

(Technically, this sensitive information type identifies these keywords by using a regular expression, not a keyword list.)

- contoso
- fabrikam
- northwind
- sandbox
- onebox
- localhost
- 127.0.0.1
- testacs.com
- s-int.net

Azure SAS

Format

The string "sig" followed by the characters and strings outlined in the pattern below.

Pattern

- the string "sig"
- zero to two whitespace characters
- an equal sign (=)
- zero to two whitespace characters
- any combination of between 43-53 characters that are lower- or uppercase letters, digits, or the percent sign (%)
- the string "%3d"
- any character that isn't a lower- or uppercase letter, digit, or percent sign (%)

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression CEP_Regex_AzureSAS finds content that matches the pattern.

```
<!--Azure SAS-->
<Entity id="4d235014-e564-47f4-a6fb-6ebb4a826834" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
        <IdMatch idRef="CEP_Regex_AzureSAS" />
  </Pattern>
</Entity>
```

Azure service bus connection string

Format

The string "EndPoint" followed by the characters and strings outlined in the pattern below, including the strings "servicebus.windows.net" and "SharedAccesKey".

Pattern

- the string "EndPoint"
- zero to two whitespace characters
- an equal sign (=)
- zero to two whitespace characters
- any combination of between 1-200 lower- or uppercase letters, digits, symbols, special characters, or spaces
- the string "servicebus.windows.net"
- any combination of between 1-200 lower- or uppercase letters, digits, symbols, special characters, or spaces
- the string "SharedAccessKey"
- zero to two whitespace characters
- an equal sign (=)
- zero to two whitespace characters
- any combination of 43 characters that are lower- or uppercase letters, digits, forward slash (/), or plus sign (+)
- an equal sign (=)

Checksum

No

Definition

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression CEP_Regex_AzureServiceBusConnectionString finds content that matches the pattern.
- The regular expression CEP_CommonExampleKeywords doesn't find content that matches the pattern.

```
<!--Azure Service Bus Connection String-->
<Entity id="b9a6578f-a83f-4fcd-bf44-2130bae49a6f" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
        <IdMatch idRef="CEP_Regex_AzureServiceBusConnectionString" />
        <Any minMatches="0" maxMatches="0">
            <Match idRef="CEP_CommonExampleKeywords" />
        </Any>
  </Pattern>
</Entity>
```

Keywords

CEP_common_example_keywords

(Technically, this sensitive information type identifies these keywords by using a regular expression, not a keyword list.)

- contoso
- fabrikam
- northwind
- sandbox
- onebox
- localhost
- 127.0.0.1
- testacs.com
- s-int.net

Azure storage account key

Format

The string "DefaultEndpointsProtocol" followed by the characters and strings outlined in the pattern below, including the string "AccountKey".

Pattern

- the string "DefaultEndpointsProtocol"
- zero to two whitespace characters
- an equal sign (=)
- zero to two whitespace characters
- any combination of between 1-200 lower- or uppercase letters, digits, symbols, special characters, or spaces
- the string "AccountKey"
- zero to two whitespace characters
- an equal sign (=)
- zero to two whitespace characters
- any combination of 86 characters that are lower- or uppercase letters, digits, forward slash (/), or plus sign (+)
- two equal signs (=)

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression CEP_Regex_AzureStorageAccountKey finds content that matches the pattern.
- The regular expression CEP_AzureEmulatorStorageAccountFilter doesn't find content that matches the pattern.
- The regular expression CEP_CommonExampleKeywords doesn't find content that matches the pattern.

```
<!--Azure Storage Account Key-->
<Entity id="c7bc98e8-551a-4c35-a92d-d2c8cda714a7" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
        <IdMatch idRef="CEP_Regex_AzureStorageAccountKey" />
        <Any minMatches="0" maxMatches="0">
            <Match idRef="CEP_AzureEmulatorStorageAccountFilter" />
            <Match idRef="CEP_CommonExampleKeywords" />
        </Any>
  </Pattern>
</Entity>
```

Keywords

CEP_azure_emulator_storage_account_filter

(Technically, this sensitive information type identifies these keywords by using a regular expression, not a keyword list.)

- Eby8vdM02xNOcqFlqUwJPLlmEtlCDXJ1OUzFT50uSRZ6IFsuFq2UVErCz4I6tq/K1SZFPTOtr/KBHBeksoGMGw==

CEP_common_example_keywords

(Technically, this sensitive information type identifies these keywords by using a regular expression, not a keyword list.)

- contoso
- fabrikam
- northwind
- sandbox
- onebox
- localhost
- 127.0.0.1
- testacs.com
- s-int.net

Azure Storage account key (generic)

Format

Any combination of 86 lower- or uppercase letters, digits, the forward slash (/), or plus sign (+), preceded or followed by the characters outlined in the pattern below.

Pattern

- zero to one of the greater than symbol (>), apostrophe ('), equal sign (=), quotation mark ("), or number sign (#)
- any combination of 86 characters that are lower- or uppercase letters, digits, the forward slash (/), or plus sign (+)

- two equal signs (=)

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression CEP_Regex_AzureStorageAccountKeyGeneric finds content that matches the pattern.

```
<!--Azure Storage Account Key (Generic)-->
<Entity id="7ff41bd0-5419-4523-91d6-383b3a37f084" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
        <IdMatch idRef="CEP_Regex_AzureStorageAccountKeyGeneric" />
  </Pattern>
</Entity>
```

Belgium driver's license number

Format

10 digits without spaces and delimiters

Pattern

10 digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_belgium_eu_driver's_license_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_driver's_license_number` or `Keywords_belgium_eu_driver's_license_number` is found.

```
      <!-- Belgium Driver's License Number -->
      <Entity id="d89fd329-9324-433c-b687-2c37bd5166f3" patternsProximity="300" recommendedConfidence="75">
        <Pattern confidenceLevel="75">
          <IdMatch idRef="Regex_belgium_eu_driver's_license_number" />
          <Any minMatches="1">
            <Match idRef="Keywords_eu_driver's_license_number" />
            <Match idRef="Keywords_belgium_eu_driver's_license_number" />
          </Any>
        </Pattern>
      </Entity>
```

Keywords

Keywords_eu_driver's_license_number

- driverlic
- driverlics
- driverlicense
- driverlicenses
- driverlicence

- driverlicences
- driver lic
- driver lics
- driver license
- driver licenses
- driver licence
- driver licences
- driverslic
- driverslics
- driverslicence
- driverslicences
- driverslicense
- driverslicenses
- drivers lic
- drivers lics
- drivers license
- drivers licenses
- drivers licence
- drivers licences
- driver'lic
- driver'lics
- driver'license
- driver'licenses
- driver'licence
- driver'licences
- driver' lic
- driver' lics
- driver' license
- driver' licenses
- driver' licence
- driver' licences
- driver'slic
- driver'slics
- driver'slicense
- driver'slicenses
- driver'slicence
- driver'slicences
- driver's lic
- driver's lics
- driver's license
- driver's licenses
- driver's licence
- driver's licences
- dl#
- dls#
- driverlic#

- driverlics#
- driverlicense#
- driverlicenses#
- driverlicence#
- driverlicences#
- driver lic#
- driver lics#
- driver license#
- driver licenses#
- driver licences#
- driverslic#
- driverslics#
- driverslicense#
- driverslicenses#
- driverslicence#
- driverslicences#
- drivers lic#
- drivers lics#
- drivers license#
- drivers licenses#
- drivers licence#
- drivers licences#
- driver'lic#
- driver'lics#
- driver'license#
- driver'licenses#
- driver'licence#
- driver'licences#
- driver' lic#
- driver' lics#
- driver' license#
- driver' licenses#
- driver' licence#
- driver' licences#
- driver'slic#
- driver'slics#
- driver'slicense#
- driver'slicenses#
- driver'slicence#
- driver'slicences#
- driver's lic#
- driver's lics#
- driver's license#
- driver's licenses#
- driver's licence#
- driver's licences#

- driving licence
- driving license
- dlno#
- driv lic
- driv licen
- driv license
- driv licenses
- driv licence
- driv licences
- driver licen
- drivers licen
- driver's licen
- driving lic
- driving licen
- driving licenses
- driving licence
- driving licences
- driving permit
- dl no
- dlno
- dl number

Keywords_belgium_eu_driver's_license_number

- rijbewijs
- rijbewijsnummer
- führerschein
- führerscheinnummer
- füehrerscheinnummer
- fuhrerschein
- fuehrerschein
- fuhrerscheinnummer
- fuehrerscheinnummer
- permis de conduire
- numéro permis conduire

Belgium national number

Format

11 digits plus optional delimiters

Pattern

11 digits plus delimiters:

- six digits and two optional periods in the format YY.MM.DD for date of birth
- An optional delimiter from dot, dash, space
- three sequential digits (odd for males, even for females)
- An optional delimiter from dot, dash, space
- two check digits

Checksum

Yes

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_belgium_national_number finds content that matches the pattern.
- A keyword from Keyword_belgium_national_number is found.
- The checksum passes.

A DLP policy has low confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_belgium_national_number finds content that matches the pattern.
- The checksum passes.

```
<!-- Belgium National Number -->
       <Entity id="fb969c9e-0fd1-4b18-8091-a2123c5e6a54" patternsProximity="300" recommendedConfidence="75">
        <Pattern confidenceLevel="75">
          <IdMatch idRef="Func_belgium_national_number" />
          <Match idRef="Keyword_belgium_national_number" />
        </Pattern>
        <Pattern confidenceLevel="65">
          <IdMatch idRef="Func_belgium_national_number" />
        </Pattern>
      </Entity>
```

Keywords

Keyword_belgium_national_number

- belasting aantal
- bnn#
- bnn
- carte d’identité
- identifiant national
- identifiantnational#
- identificatie
- identification
- identifikation
- identifikationsnummer
- identifizierung
- identité
- identiteit
- identiteitskaart
- identity
- inscription
- national number
- national register
- nationalnumber#
- nationalnumber
- nif#
- nif

- numéro d'assuré
- numéro de registre national
- numéro de sécurité
- numéro d'identification
- numéro d'immatriculation
- numéro national
- numéronational#
- personal id number
- personalausweis
- personalidnumber#
- registratie
- registration
- registrationsnumme
- registrierung
- social security number
- ssn#
- ssn
- steuernummer
- tax id
- tax identification no
- tax identification number
- tax no#
- tax no
- tax number
- tax registration number
- taxid#
- taxidno#
- taxidnumber#
- taxno#
- taxnumber#
- taxnumber
- tin id
- tin no
- tin#

Belgium passport number

Format

two letters followed by six digits with no spaces or delimiters

Pattern

two letters and followed by six digits

Checksum

not applicable

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_belgium_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_belgium_eu_passport_number` is found.
- The regular expression `Regex_eu_passport_date2` finds date in the format DD MM YY or a keyword from `Keywords_eu_passport_date` or `Keywords_belgium_eu_passport_number` is found

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_belgium_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_belgium_eu_passport_number` is found.

```
<!-- Belgium Passport Number -->
<Entity id="d7b1315b-21ca-4774-a32a-596010ff78fd" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Regex_belgium_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_belgium_eu_passport_number" />
    </Any>
    <Any minMatches="1">
      <Match idRef="Regex_eu_passport_date2" />
      <Match idRef="Keywords_eu_passport_date" />
      <Match idRef="Keywords_belgium_eu_passport_date" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_belgium_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_belgium_eu_passport_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_passport_number

- passport#
- passport #
- passportid
- passports
- passportno
- passport no
- passportnumber
- passport number
- passportnumbers
- passport numbers

Keywords_belgium_eu_passport_number

- numéro passeport
- paspoort nr
- paspoort-nr
- paspoortnummer
- paspoortnummers
- Passeport carte

- Passeport livre
- Pass-Nr
- Passnummer
- reisepass kein

Keywords_eu_passport_date

- date of issue
- date of expiry

Belgium value added tax number

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

12-character alphanumeric pattern

Pattern

12-character alphanumeric pattern:

- a letter B or b
- a letter E or e
- a digit 0
- a digit from 1 to 9
- an optional dot or Hyphen or space
- four digits
- an optional dot or Hyphen or space
- four digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_belgium_value_added_tax_number finds content that matches the pattern.
- A keyword from Keywords_belgium_value_added_tax_number is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_belgium_value_added_tax_number finds content that matches the pattern.

```
      <!-- Belgium Value Added Tax Number -->
      <Entity id="85b5b3c3-f2de-4ae8-ac46-fd3cb38bf9ed" patternsProximity="300" recommendedConfidence="85">
        <Pattern confidenceLevel="85">
          <IdMatch idRef="Func_belgium_value_added_tax_number" />
          <Match idRef="Keywords_belgium_value_added_tax_number" />
        </Pattern>
        <Pattern confidenceLevel="75">
          <IdMatch idRef="Func_belgium_value_added_tax_number" />
        </Pattern>
      </Entity>
    </Version>
```

Keywords

Keyword_belgium_value_added_tax_number

- nº tva
- vat number
- vat no
- numéro t.v.a
- umsatzsteuer-identifikationsnummer
- umsatzsteuernummer
- btw
- btw#
- vat#

Brazil CPF number

Format

11 digits that include a check digit and can be formatted or unformatted

Pattern

Formatted:

- three digits
- a period
- three digits
- a period
- three digits
- a hyphen
- two digits that are check digits

Unformatted:

- 11 digits where the last two digits are check digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_brazil_cpf finds content that matches the pattern.
- A keyword from Keyword_brazil_cpf is found.
- The checksum passes.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_brazil_cpf finds content that matches the pattern.
- The checksum passes.

```
<!-- Brazil CPF Number -->
<Entity id="78e09124-f2c3-4656-b32a-c1a132cd2711" recommendedConfidence="85" patternsProximity="300">
  <Pattern confidenceLevel="85">
     <IdMatch idRef="Func_brazil_cpf"/>
     <Match idRef="Keyword_brazil_cpf"/>
  </Pattern>
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Func_brazil_cpf"/>
  </Pattern>
</Entity>
```

Keywords

Keyword_brazil_cpf

- CPF
- Identification
- Registration
- Revenue
- Cadastro de Pessoas Físicas
- Imposto
- Identificação
- Inscrição
- Receita

Brazil legal entity number (CNPJ)

Format

14 digits that include a registration number, branch number, and check digits, plus delimiters

Pattern

14 digits, plus delimiters:

- two digits
- a period
- three digits
- a period
- three digits (these first eight digits are the registration number)
- a forward slash
- four-digit branch number
- a hyphen
- two digits that are check digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_brazil_cnpj finds content that matches the pattern.
- A keyword from Keyword_brazil_cnpj is found.
- The checksum passes.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_brazil_cnpj finds content that matches the pattern.
- The checksum passes.

```
<!-- Brazil Legal Entity Number (CNPJ) -->
<Entity id="9b58b5cd-5e90-4df6-b34f-1ebcc88ceae4" recommendedConfidence="85" patternsProximity="300">
   <Pattern confidenceLevel="85">
     <IdMatch idRef="Func_brazil_cnpj"/>
     <Match idRef="Keyword_brazil_cnpj"/>
  </Pattern>
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Func_brazil_cnpj"/>
  </Pattern>
</Entity>
```

Keywords

Keyword_brazil_cnpj

- CNPJ
- CNPJ/MF
- CNPJ-MF
- National Registry of Legal Entities
- Taxpayers Registry
- Legal entity
- Legal entities
- Registration Status
- Business
- Company
- CNPJ
- Cadastro Nacional da Pessoa Jurídica
- Cadastro Geral de Contribuintes
- CGC
- Pessoa jurídica
- Pessoas jurídicas
- Situação cadastral
- Inscrição
- Empresa

Brazil national identification card (RG)

Format

Registro Geral (old format): Nine digits

Registro de Identidade (RIC) (new format): 11 digits

Pattern

Registro Geral (old format):

- two digits
- a period
- three digits
- a period
- three digits
- a hyphen
- one digit that is a check digit

Registro de Identidade (RIC) (new format):

- 10 digits
- a hyphen
- one digit that is a check digit

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_brazil_rg finds content that matches the pattern.
- A keyword from Keyword_brazil_rg is found.
- The checksum passes.

```
<!-- Brazil National ID Card (RG) -->
<Entity id="486de900-db70-41b3-a886-abdf25af119c" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_brazil_rg" />
    <Match idRef="Keyword_brazil_rg" />
  </Pattern>
</Entity>
```

Keywords

Keyword_brazil_rg

- Cédula de identidade
- identity card
- national id
- número de rregistro
- registro de lidentidade
- registro geral
- RG (this keyword is case-sensitive)
- RIC (this keyword is case-sensitive)

Bulgaria driver's license number

Format

nine digits without spaces and delimiters

Pattern

nine digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_bulgaria_eu_driver's_license_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_driver's_license_number` or `Keywords_bulgaria_eu_driver's_license_number` is found.

```
<!-- Bulgaria Driver's License Number -->
<Entity id="66d39258-94c2-43b2-804b-aa312258e54b" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_bulgaria_eu_driver's_license_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_driver's_license_number" />
      <Match idRef="Keywords_bulgaria_eu_driver's_license_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_driver's_license_number

- driverlic
- driverlics
- driverlicense
- driverlicenses
- driverlicence
- driverlicences
- driver lic
- driver lics
- driver license
- driver licenses
- driver licence
- driver licences
- driverslic
- driverslics
- driverslicence
- driverslicences
- driverslicense
- driverslicenses
- drivers lic
- drivers lics
- drivers license
- drivers licenses
- drivers licence
- drivers licences
- driver'lic
- driver'lics
- driver'license
- driver'licenses

- driver'licence
- driver'licences
- driver' lic
- driver' lics
- driver' license
- driver' licenses
- driver' licence
- driver' licences
- driver'slic
- driver'slics
- driver'slicense
- driver'slicenses
- driver'slicence
- driver'slicences
- driver's lic
- driver's lics
- driver's license
- driver's licenses
- driver's licence
- driver's licences
- dl#
- dls#
- driverlic#
- driverlics#
- driverlicense#
- driverlicenses#
- driverlicence#
- driverlicences#
- driver lic#
- driver lics#
- driver license#
- driver licenses#
- driver licences#
- driverslic#
- driverslics#
- driverslicense#
- driverslicenses#
- driverslicence#
- driverslicences#
- drivers lic#
- drivers lics#
- drivers license#
- drivers licenses#
- drivers licence#
- drivers licences#
- driver'lic#

- driver'lics#
- driver'license#
- driver'licenses#
- driver'licence#
- driver'licences#
- driver' lic#
- driver' lics#
- driver' license#
- driver' licenses#
- driver' licence#
- driver' licences#
- driver'slic#
- driver'slics#
- driver'slicense#
- driver'slicenses#
- driver'slicence#
- driver'slicences#
- driver's lic#
- driver's lics#
- driver's license#
- driver's licenses#
- driver's licence#
- driver's licences#
- driving licence
- driving license
- dlno#
- driv lic
- driv licen
- driv license
- driv licenses
- driv licence
- driv licences
- driver licen
- drivers licen
- driver's licen
- driving lic
- driving licen
- driving licenses
- driving licence
- driving licences
- driving permit
- dl no
- dlno
- dl number

Keywords_bulgaria_eu_driver's_license_number

- свидетелство за управление на мпс

- свидетелство за управление на моторно превозно средство
- сумпс
- шофьорска книжка
- шофьорски книжки

Bulgaria uniform civil number

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

10 digits without spaces and delimiters

Pattern

10 digits without spaces and delimiters

- six digits that correspond to the birth date (YYMMDD)
- two digits that correspond to the birth order
- one digit that corresponds to gender: An even digit for male and an odd digit for female
- one check digit

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_bulgaria_eu_national_id_card` finds content that matches the pattern.
- A keyword from `Keywords_bulgaria_eu_national_id_card` is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_bulgaria_eu_national_id_card` finds content that matches the pattern.

```
<!-- Bulgaria Uniform Civil Number -->
<Entity id="100d58b1-0a35-4fb1-aa89-e4a86fb53fcc" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_bulgaria_eu_national_id_card" />
    <Match idRef="Keywords_bulgaria_eu_national_id_card" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_bulgaria_eu_national_id_card" />
    <Any minMatches="0" maxMatches="0">
      <Match idRef="Keywords_bulgaria_eu_telephone_number" />
      <Match idRef="Keywords_bulgaria_eu_mobile_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_bulgaria_eu_national_id_card

- bnn#
- bnn
- bucn#
- bucn
- edinen grazhdanski nomer
- egn#
- egn
- identification number
- national id
- national number
- nationalnumber#
- nationalnumber
- personal id
- personal no
- personal number
- personalidnumber#
- social security number
- ssn#
- ssn
- uniform civil id
- uniform civil no
- uniform civil number
- uniformcivilno#
- uniformcivilno
- uniformcivilnumber#
- uniformcivilnumber
- unique citizenship number
- егн#
- егн
- единен граждански номер
- идентификационен номер
- личен номер
- лична идентификация
- лично не
- национален номер
- номер на гражданството
- униформ id
- униформ граждански id
- униформ граждански не
- униформ граждански номер
- униформгражданскиid#
- униформгражданскине.#

Bulgaria passport number

Format

nine digits without spaces and delimiters

Pattern

nine digits

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_bulgaria_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_bulgaria_eu_passport_number` is found.
- The regular expression `Regex_eu_passport_date1` finds date in the format DD.MM.YYYY or a keyword from `Keywords_eu_passport_date` is found

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_bulgaria_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_bulgaria_eu_passport_number` is found.

```
<!-- Bulgaria Passport Number -->
<Entity id="f7172b82-c588-4216-845e-4e54e397f29a" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Regex_bulgaria_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_bulgaria_eu_passport_number" />
    </Any>
    <Any minMatches="1">
      <Match idRef="Regex_eu_passport_date1" />
      <Match idRef="Keywords_eu_passport_date" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_bulgaria_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_bulgaria_eu_passport_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_passport_number

- passport#
- passport #
- passportid
- passports
- passportno
- passport no
- passportnumber
- passport number

- passportnumbers
- passport numbers

Keywords_bulgaria_eu_passport_number

- номер на паспорта
- номер на паспорт
- паспорт №

Keywords_eu_passport_date

- date of issue
- date of expiry

Canada bank account number

Format

7 or 12 digits

Pattern

A Canada Bank Account Number is 7 or 12 digits.

A Canada bank account transit number is:

- five digits
- a hyphen
- three digits OR
- a zero "0"
- eight digits

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_canada_bank_account_number finds content that matches the pattern.
- A keyword from Keyword_canada_bank_account_number is found.
- The regular expression Regex_canada_bank_account_transit_number finds content that matches the pattern.

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_canada_bank_account_number finds content that matches the pattern.
- A keyword from Keyword_canada_bank_account_number is found.

```
<!-- Canada Bank Account Number -->
<Entity id="552e814c-cb50-4d94-bbaa-bb1d1ffb34de" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
        <IdMatch idRef="Regex_canada_bank_account_number" />
        <Match idRef="Keyword_canada_bank_account_number" />
        <Match idRef="Regex_canada_bank_account_transit_number" />
   </Pattern>
   <Pattern confidenceLevel="75">
        <IdMatch idRef="Regex_canada_bank_account_number" />
        <Match idRef="Keyword_canada_bank_account_number" />
   </Pattern>
</Entity>
```

Keywords

Keyword_canada_bank_account_number

- canada savings bonds
- canada revenue agency
- canadian financial institution
- direct deposit form
- canadian citizen
- legal representative
- notary public
- commissioner for oaths
- child care benefit
- universal child care
- canada child tax benefit
- income tax benefit
- harmonized sales tax
- social insurance number
- income tax refund
- child tax benefit
- territorial payments
- institution number
- deposit request
- banking information
- direct deposit

Canada driver's license number

Format

Varies by province

Pattern

Various patterns covering Alberta, British Columbia, Manitoba, New Brunswick, Newfoundland/Labrador, Nova Scotia, Ontario, Prince Edward Island, Quebec, and Saskatchewan

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_[province_name]_drivers_license_number finds content that matches the pattern.
- A keyword from Keyword_[province_name]_drivers_license_name is found.
- A keyword from Keyword_canada_drivers_license is found.

```
<!-- Canada Driver's License Number -->
    <Entity id="37186abb-8e48-4800-ad3c-e3d1610b3db0" patternsProximity="300" recommendedConfidence="75">
      <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_alberta_drivers_license_number" />
        <Match idRef="Keyword_alberta_drivers_license_name" />
        <Match idRef="Keyword_canada_drivers_license" />
      </Pattern>
      <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_british_columbia_drivers_license_number" />
        <Match idRef="Keyword_british_columbia_drivers_license_name" />
        <Match idRef="Keyword_canada_drivers_license" />
      </Pattern>
      <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_manitoba_drivers_license_number" />
        <Match idRef="Keyword_manitoba_drivers_license_name" />
        <Match idRef="Keyword_canada_drivers_license" />
      </Pattern>
      <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_new_brunswick_drivers_license_number" />
        <Match idRef="Keyword_new_brunswick_drivers_license_name" />
        <Match idRef="Keyword_canada_drivers_license" />
      </Pattern>
      <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_newfoundland_labrador_drivers_license_number" />
        <Match idRef="Keyword_newfoundland_labrador_drivers_license_name" />
        <Match idRef="Keyword_canada_drivers_license" />
      </Pattern>
      <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_nova_scotia_drivers_license_number" />
        <Match idRef="Keyword_nova_scotia_drivers_license_name" />
        <Match idRef="Keyword_canada_drivers_license" />
      </Pattern>
      <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_ontario_drivers_license_number" />
        <Match idRef="Keyword_ontario_drivers_license_name" />
        <Match idRef="Keyword_canada_drivers_license" />
      </Pattern>
      <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_prince_edward_island_drivers_license_number" />
        <Match idRef="Keyword_prince_edward_island_drivers_license_name" />
        <Match idRef="Keyword_canada_drivers_license" />
      </Pattern>
      <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_quebec_drivers_license_number" />
        <Match idRef="Keyword_quebec_drivers_license_name" />
        <Match idRef="Keyword_canada_drivers_license" />
      </Pattern>
      <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_saskatchewan_drivers_license_number" />
        <Match idRef="Keyword_saskatchewan_drivers_license_name" />
        <Match idRef="Keyword_canada_drivers_license" />
      </Pattern>
    </Entity>
```

Keywords

Keyword_[province_name]_drivers_license_name

- The province abbreviation, for example AB
- The province name, for example Alberta

Keyword_canada_drivers_license

- DL
- DLS
- CDL
- CDLS
- DriverLic
- DriverLics
- DriverLicense
- DriverLicenses
- DriverLicence
- DriverLicences
- Driver Lic
- Driver Lics
- Driver License
- Driver Licenses
- Driver Licence
- Driver Licences
- DriversLic
- DriversLics
- DriversLicence
- DriversLicences
- DriversLicense
- DriversLicenses
- Drivers Lic
- Drivers Lics
- Drivers License
- Drivers Licenses
- Drivers Licence
- Drivers Licences
- Driver'Lic
- Driver'Lics
- Driver'License
- Driver'Licenses
- Driver'Licence
- Driver'Licences
- Driver' Lic
- Driver' Lics
- Driver' License
- Driver' Licenses
- Driver' Licence
- Driver' Licences
- Driver'sLic
- Driver'sLics
- Driver'sLicense
- Driver'sLicenses
- Driver'sLicence
- Driver'sLicences

- Driver's Lic
- Driver's Lics
- Driver's License
- Driver's Licenses
- Driver's Licence
- Driver's Licences
- Permis de Conduire
- id
- ids
- idcard number
- idcard numbers
- idcard #
- idcard #s
- idcard card
- idcard cards
- idcard
- identification number
- identification numbers
- identification #
- identification #s
- identification card
- identification cards
- identification
- DL#
- DLS#
- CDL#
- CDLS#
- DriverLic#
- DriverLics#
- DriverLicense#
- DriverLicenses#
- DriverLicence#
- DriverLicences#
- Driver Lic#
- Driver Lics#
- Driver License#
- Driver Licenses#
- Driver License#
- Driver Licences#
- DriversLic#
- DriversLics#
- DriversLicense#
- DriversLicenses#
- DriversLicence#
- DriversLicences#
- Drivers Lic#

- Drivers Lics#
- Drivers License#
- Drivers Licenses#
- Drivers Licence#
- Drivers Licences#
- Driver'Lic#
- Driver'Lics#
- Driver'License#
- Driver'Licenses#
- Driver'Licence#
- Driver'Licences#
- Driver' Lic#
- Driver' Lics#
- Driver' License#
- Driver' Licenses#
- Driver' Licence#
- Driver' Licences#
- Driver'sLic#
- Driver'sLics#
- Driver'sLicense#
- Driver'sLicenses#
- Driver'sLicence#
- Driver'sLicences#
- Driver's Lic#
- Driver's Lics#
- Driver's License#
- Driver's Licenses#
- Driver's Licence#
- Driver's Licences#
- Permis de Conduire#
- id#
- ids#
- idcard card#
- idcard cards#
- idcard#
- identification card#
- identification cards#
- identification#

Canada health service number

Format

10 digits

Pattern

10 digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_canada_health_service_number finds content that matches the pattern.
- A keyword from Keyword_canada_health_service_number is found.

```
<!-- Canada Health Service Number -->
<Entity id="59c0bf39-7fab-482c-af25-00faa4384c94" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
        <IdMatch idRef="Regex_canada_health_service_number" />
        <Any minMatches="1">
          <Match idRef="Keyword_canada_health_service_number" />
        </Any>
  </Pattern>
</Entity>
```

Keywords

Keyword_canada_health_service_number

- personal health number
- patient information
- health services
- speciality services
- automobile accident
- patient hospital
- psychiatrist
- workers compensation
- disability

Canada passport number

Format

two uppercase letters followed by six digits

Pattern

two uppercase letters followed by six digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_canada_passport_number finds content that matches the pattern.
- A keyword from Keyword_canada_passport_number or Keyword_passport is found.

```
<!-- Canada Passport Number -->
<Entity id="14d0db8b-498a-43ed-9fca-f6097ae687eb" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
        <IdMatch idRef="Regex_canada_passport_number" />
        <Any minMatches="1">
          <Match idRef="Keyword_canada_passport_number" />
          <Match idRef="Keyword_passport" />
        </Any>
  </Pattern>
</Entity>
```

Keywords

Keyword_canada_passport_number

- canadian citizenship
- canadian passport
- passport application
- passport photos
- certified translator
- canadian citizens
- processing times
- renewal application

Keyword_passport

- Passport Number
- Passport No
- Passport #
- Passport#
- PassportID
- Passportno
- passportnumber
- パスポート
- パスポート番号
- パスポートのNum
- パスポート#
- Numéro de passeport
- Passeport n °
- Passeport Non
- Passeport #
- Passeport#
- PasseportNon
- Passeportn °

Canada personal health identification number (PHIN)

Format

nine digits

Pattern

nine digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_canada_phin finds content that matches the pattern.
- At least two keywords from Keyword_canada_phin or Keyword_canada_provinces are found.

```
<!-- Canada PHIN -->
<Entity id="722e12ac-c89a-4ec8-a1b7-fea3469f89db" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
        <IdMatch idRef="Regex_canada_phin" />
        <Any minMatches="2">
          <Match idRef="Keyword_canada_phin" />
          <Match idRef="Keyword_canada_provinces" />
        </Any>
  </Pattern>
</Entity>
```

Keywords

Keyword_canada_phin

- social insurance number
- health information act
- income tax information
- manitoba health
- health registration
- prescription purchases
- benefit eligibility
- personal health
- power of attorney
- registration number
- personal health number
- practitioner referral
- wellness professional
- patient referral
- health and wellness

Keyword_canada_provinces

- Nunavut
- Quebec
- Northwest Territories
- Ontario
- British Columbia
- Alberta
- Saskatchewan
- Manitoba
- Yukon
- Newfoundland and Labrador
- New Brunswick
- Nova Scotia

- Prince Edward Island
- Canada

Canada social insurance number

Format

nine digits with optional hyphens or spaces

Pattern

Formatted:

- three digits
- a hyphen or space
- three digits
- a hyphen or space
- three digits

Unformatted: nine digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_canadian_sin finds content that matches the pattern.
- At least two of any combination of the following:
  - A keyword from Keyword_sin is found.
  - A keyword from Keyword_sin_collaborative is found.
  - The function Func_eu_date finds a date in the right date format.
- The checksum passes.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_unformatted_canadian_sin finds content that matches the pattern.
- A keyword from Keyword_sin is found.
- The checksum passes.

```
<!-- Canada Social Insurance Number -->
<Entity id="a2f29c85-ecb8-4514-a610-364790c0773e" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
        <IdMatch idRef="Func_canadian_sin" />
        <Any minMatches="2">
          <Match idRef="Keyword_sin" />
          <Match idRef="Keyword_sin_collaborative" />
          <Match idRef="Func_eu_date" />
        </Any>
  </Pattern>
  <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_unformatted_canadian_sin" />
        <Match idRef="Keyword_sin" />
  </Pattern>
</Entity>
```

Keywords

Keyword_sin

- sin
- social insurance
- numero d'assurance sociale
- sins
- ssn
- ssns
- social security
- numero d'assurance social
- national identification number
- national id
- sin#
- soc ins
- social ins

Keyword_sin_collaborative

- driver's license
- drivers license
- driver's licence
- drivers licence
- DOB
- Birthdate
- Birthday
- Date of Birth

Chile identity card number

Format

seven to eight digits plus delimiters a check digit or letter

Pattern

seven to eight digits plus delimiters:

- one to two digits
- an optional period
- three digits
- an optional period
- three digits
- a dash
- one digit or letter (not case-sensitive) which is a check digit

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_chile_id_card finds content that matches the pattern.
- A keyword from Keyword_chile_id_card is found.

- The checksum passes.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_chile_id_card finds content that matches the pattern.
- The checksum passes.

```
<!-- Chile Identity Card Number -->
<Entity id="4e979794-49a0-407e-a0b9-2c536937b925" recommendedConfidence="85" patternsProximity="300">
  <Pattern confidenceLevel="85">
     <IdMatch idRef="Func_chile_id_card"/>
     <Match idRef="Keyword_chile_id_card"/>
  </Pattern>
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Func_chile_id_card"/>
  </Pattern>
</Entity>
```

Keywords

Keyword_chile_id_card

- cédula de identidad
- identificación
- national identification
- national identification number
- national id
- número de identificación nacional
- rol único nacional
- rol único tributario
- RUN
- RUT
- tarjeta de identificación
- Rol Unico Nacional
- Rol Unico Tributario
- RUN#
- RUT#
- nationaluniqueroleID#
- nacional identidad
- número identificación
- identidad número
- numero identificacion
- identidad numero
- Chilean identity no.
- Chilean identity number
- Chilean identity #
- Unique Tax Registry
- Unique Tributary Role
- Unique Tax Role
- Unique Tributary Number
- Unique National Number
- Unique National Role

- National unique role
- Chile identity no.
- Chile identity number
- Chile identity #

China resident identity card (PRC) number

Format

18 digits

Pattern

18 digits:

- six digits that are an address code
- eight digits in the form YYYYMMDD, which are the date of birth
- three digits that are an order code
- one digit that is a check digit

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_china_resident_id finds content that matches the pattern.
- A keyword from Keyword_china_resident_id is found.
- The checksum passes.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_china_resident_id finds content that matches the pattern.
- The checksum passes.

```
<!-- China Resident Identity Card (PRC) Number -->
<Entity id="c92daa86-2d16-4871-901f-816b3f554fc1" recommendedConfidence="85" patternsProximity="300">
  <Pattern confidenceLevel="85">
     <IdMatch idRef="Func_china_resident_id"/>
     <Match idRef="Keyword_china_resident_id"/>
  </Pattern>
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Func_china_resident_id"/>
  </Pattern>
</Entity>
```

Keywords

Keyword_china_resident_id

- Resident Identity Card
- PRC
- National Identification Card
- 身份证
- 居民 身份证
- 居民身份证

- 鉴定
- 身分證
- 居民 身份證
- 鑑定

Credit card number

Format

14 to 16 digits that can be formatted or unformatted (ddddddddddddddd) and that must pass the Luhn test.

Pattern

Complex and robust pattern that detects cards from all major brands worldwide, including Visa, MasterCard, Discover Card, JCB, American Express, gift cards, and diner cards.

Checksum

Yes, the Luhn checksum

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_credit_card finds content that matches the pattern.
- One of the following is true:
  - A keyword from Keyword_cc_verification is found.
  - A keyword from Keyword_cc_name is found.
  - The function Func_expiration_date finds a date in the right date format.
- The checksum passes.

A DLP policy has low confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_credit_card finds content that matches the pattern.
- The checksum passes.

```
<!-- Credit Card Number -->
<Entity id="50842eb7-edc8-4019-85dd-5a5c1f2bb085" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
        <IdMatch idRef="Func_credit_card" />
        <Any minMatches="1">
          <Match idRef="Keyword_cc_verification" />
          <Match idRef="Keyword_cc_name" />
          <Match idRef="Func_expiration_date" />
        </Any>
  </Pattern>
  <Pattern confidenceLevel="65">
        <IdMatch idRef="Func_credit_card" />
  </Pattern>
</Entity>
```

Keywords

Keyword_cc_verification

- card verification
- card identification number
- cvn
- cid

- cvc2
- cvv2
- pin block
- security code
- security number
- security no
- issue number
- issue no
- cryptogramme
- numéro de sécurité
- numero de securite
- kreditkartenprüfnummer
- kreditkartenprufnummer
- prüfziffer
- prufziffer
- sicherheits Kode
- sicherheitscode
- sicherheitsnummer
- verfalldatum
- codice di verifica
- cod. sicurezza
- cod sicurezza
- n autorizzazione
- código
- codigo
- cod. seg
- cod seg
- código de segurança
- codigo de seguranca
- codigo de segurança
- código de seguranca
- cód. segurança
- cod. seguranca
- cod. segurança
- cód. seguranca
- cód segurança
- cod seguranca
- cod segurança
- cód seguranca
- número de verificação
- numero de verificacao
- ablauf
- gültig bis
- gültigkeitsdatum
- gultig bis
- gultigkeitsdatum

- scadenza
- data scad
- fecha de expiracion
- fecha de venc
- vencimiento
- válido hasta
- valido hasta
- vto
- data de expiração
- data de expiracao
- data em que expira
- validade
- valor
- vencimento
- transaction
- transaction number
- reference number
- セキュリティコード
- セキュリティ コード
- セキュリティナンバー
- セキュリティ ナンバー
- セキュリティ番号

Keyword_cc_name

- amex
- american express
- americanexpress
- americano espresso
- Visa
- mastercard
- master card
- mc
- mastercards
- master cards
- diner's Club
- diners club
- dinersclub
- discover
- discover card
- discovercard
- discover cards
- JCB
- BrandSmart
- japanese card bureau
- carte blanche
- carteblanche
- credit card

- cc#
- cc#:
- expiration date
- exp date
- expiry date
- date d’expiration
- date d'exp
- date expiration
- bank card
- bankcard
- card number
- card num
- cardnumber
- cardnumbers
- card numbers
- creditcard
- credit cards
- creditcards
- ccn
- card holder
- cardholder
- card holders
- cardholders
- check card
- checkcard
- check cards
- checkcards
- debit card
- debitcard
- debit cards
- debitcards
- atm card
- atmcard
- atm cards
- atmcards
- enroute
- en route
- card type
- Cardmember Acct
- cardmember account
- Cardno
- Corporate Card
- Corporate cards
- Type of card
- card account number
- card member account

- Cardmember Acct.
- card no.
- card no
- card number
- carte bancaire
- carte de crédit
- carte de credit
- numéro de carte
- numero de carte
- nº de la carte
- nº de carte
- kreditkarte
- karte
- karteninhaber
- karteninhabers
- kreditkarteninhaber
- kreditkarteninstitut
- kreditkartentyp
- eigentümername
- kartennr
- kartennummer
- kreditkartennummer
- kreditkarten-nummer
- carta di credito
- carta credito
- n. carta
- n carta
- nr. carta
- nr carta
- numero carta
- numero della carta
- numero di carta
- tarjeta credito
- tarjeta de credito
- tarjeta crédito
- tarjeta de crédito
- tarjeta de atm
- tarjeta atm
- tarjeta debito
- tarjeta de debito
- tarjeta débito
- tarjeta de débito
- nº de tarjeta
- no. de tarjeta
- no de tarjeta
- numero de tarjeta

- número de tarjeta
- tarjeta no
- tarjetahabiente
- cartão de crédito
- cartão de credito
- cartao de crédito
- cartao de credito
- cartão de débito
- cartao de débito
- cartão de debito
- cartao de debito
- débito automático
- debito automatico
- número do cartão
- numero do cartão
- número do cartao
- numero do cartao
- número de cartão
- numero de cartão
- número de cartao
- numero de cartao
- nº do cartão
- nº do cartao
- nº. do cartão
- no do cartão
- no do cartao
- no. do cartão
- no. do cartao
- クレジットカード番号
- クレジットカードナンバー
- クレジットカード#
- クレジットカード
- クレジット
- クレカ
- カード番号
- カードナンバー
- カード#
- アメックス
- アメリカンエクスプレス
- アメリカン エクスプレス
- Visaカード
- Visa カード
- マスターカード
- マスター カード
- マスター
- ダイナースクラブ

- ダイナース クラブ
- ダイナース
- 有効期限
- 期限
- キャッシュカード
- キャッシュ カード
- カード名義人
- カードの名義人
- カードの名義
- デビット カード
- デビットカード

Croatia driver's license number

Format

eight digits without spaces and delimiters

Pattern

eight digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_croatia_eu_driver's_license_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_driver's_license_number` or `Keywords_croatia_eu_driver's_license_number` is found.

```
<!-- Croatia Driver's License Number -->
<Entity id="005b3ef1-47dd-4e68-bb02-c6db484d00f2" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_croatia_eu_driver's_license_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_driver's_license_number" />
      <Match idRef="Keywords_croatia_eu_driver's_license_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_driver's_license_number

- driverlic
- driverlics
- driverlicense
- driverlicenses
- driverlicence
- driverlicences
- driver lic
- driver lics

- driver license
- driver licenses
- driver licence
- driver licences
- driverslic
- driverslics
- driverslicence
- driverslicences
- driverslicense
- driverslicenses
- drivers lic
- drivers lics
- drivers license
- drivers licenses
- drivers licence
- drivers licences
- driver'lic
- driver'lics
- driver'license
- driver'licenses
- driver'licence
- driver'licences
- driver' lic
- driver' lics
- driver' license
- driver' licenses
- driver' licence
- driver' licences
- driver'slic
- driver'slics
- driver'slicense
- driver'slicenses
- driver'slicence
- driver'slicences
- driver's lic
- driver's lics
- driver's license
- driver's licenses
- driver's licence
- driver's licences
- dl#
- dls#
- driverlic#
- driverlics#
- driverlicense#
- driverlicenses#

- driverlicence#
- driverlicences#
- driver lic#
- driver lics#
- driver license#
- driver licenses#
- driver licences#
- driverslic#
- driverslics#
- driverslicense#
- driverslicenses#
- driverslicence#
- driverslicences#
- drivers lic#
- drivers lics#
- drivers license#
- drivers licenses#
- drivers licence#
- drivers licences#
- driver'lic#
- driver'lics#
- driver'license#
- driver'licenses#
- driver'licence#
- driver'licences#
- driver' lic#
- driver' lics#
- driver' license#
- driver' licenses#
- driver' licence#
- driver' licences#
- driver'slic#
- driver'slics#
- driver'slicense#
- driver'slicenses#
- driver'slicence#
- driver'slicences#
- driver's lic#
- driver's lics#
- driver's license#
- driver's licenses#
- driver's licence#
- driver's licences#
- driving licence
- driving license
- dlno#

- driv lic
- driv licen
- driv license
- driv licenses
- driv licence
- driv licences
- driver licen
- drivers licen
- driver's licen
- driving lic
- driving licen
- driving licenses
- driving licence
- driving licences
- driving permit
- dl no
- dlno
- dl number

Keywords_croatia_eu_driver's_license_number

- vozačka dozvola
- vozačke dozvole

Croatia identity card number

This sensitive information type entity is included in the EU National Identification Number sensitive information type. It's available as a stand-alone sensitive information type entity.

Format

nine digits

Pattern

nine consecutive digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_croatia_id_card finds content that matches the pattern.
- A keyword from Keyword_croatia_id_card is found.

```
<!--Croatia Identity Card Number-->
<Entity id="ff12f884-c20a-4189-b185-34c8e7258d47" recommendedConfidence="75" patternsProximity="300">
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Func_croatia_id_card"/>
     <Match idRef="Keyword_croatia_id_card"/>
  </Pattern>
</Entity>
```

Keywords

Keyword_croatia_id_card

- majstorski broj građana
- master citizen number
- nacionalni identifikacijski broj
- national identification number
- oib#
- oib
- osobna iskaznica
- osobni id
- osobni identifikacijski broj
- personal identification number
- porezni broj
- porezni identifikacijski broj
- tax id
- tax identification no
- tax identification number
- tax no#
- tax no
- tax number
- tax registration number
- taxid#
- taxidno#
- taxidnumber#
- taxno#
- taxnumber#
- taxnumber
- tin id
- tin no
- tin#

Croatia passport number

Format

nine digits without spaces and delimiters

Pattern

nine digits

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_croatia_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_croatia_eu_passport_number` is found.
- The regular expression `Regex_eu_passport_date1` finds date in the format DD.MM.YYYY or a keyword from `Keywords_eu_passport_date` is found

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_croatia_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_croatia_eu_passport_number` is found.

```
<!-- Croatia Passport Number -->
<Entity id="7d7a729d-32d8-4204-8d01-d5e6a6c25581" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Regex_croatia_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_croatia_eu_passport_number" />
    </Any>
    <Any minMatches="1">
      <Match idRef="Regex_eu_passport_date1" />
      <Match idRef="Keywords_eu_passport_date" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_croatia_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_croatia_eu_passport_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_passport_number_common

- passport#
- passport #
- passportid
- passports
- passportno
- passport no
- passportnumber
- passport number
- passportnumbers
- passport numbers

Keywords_croatia_eu_passport_number

- broj putovnice
- br. Putovnice
- br putovnice

Croatia personal identification (OIB) number

Format

11 digits

Pattern

11 digits:

- 10 digits
- final digit is a check digit

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_croatia_oib_number finds content that matches the pattern.
- A keyword from Keywords_croatia_eu_tax_file_number is found.
- The checksum passes.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_croatia_oib_number finds content that matches the pattern.
- The checksum passes.

```
<!-- Croatia Personal Identification (OIB) Number -->
<Entity id="31983b6d-db95-4eb2-a630-b44bd091968d" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_croatia_oib_number" />
    <Match idRef="Keywords_croatia_eu_tax_file_number" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_croatia_oib_number" />
  </Pattern>
</Entity>
```

Keywords

Keyword_croatia_oib_number

- majstorski broj građana
- master citizen number
- nacionalni identifikacijski broj
- national identification number
- oib#
- oib
- osobna iskaznica
- osobni id
- osobni identifikacijski broj
- personal identification number
- porezni broj
- porezni identifikacijski broj
- tax id
- tax identification no
- tax identification number
- tax no#
- tax no
- tax number
- tax registration number
- taxid#
- taxidno#
- taxidnumber#

- taxno#
- taxnumber#
- taxnumber
- tin id
- tin no
- tin#

Croatia social security number or equivalent identification

This sensitive information type entity is only available in the EU Social Security Number or Equivalent ID sensitive information type.

Format

11 digits without spaces and delimiters

Pattern

11 digits:

- 10 digits
- one check digit

Checksum

Yes

Definition

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_croatia_eu_ssn_or_equivalent` finds content that matches the pattern.
- A keyword from `Keywords_croatia_eu_ssn_or_equivalent` is found.

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_croatia_eu_ssn_or_equivalent` finds content that matches the pattern.

```
 <!-- EU SSN or Equivalent Number -->
<Entity id="d24e32a4-c0bb-4ba8-899d-6303b95742d9" patternsProximity="300" recommendedConfidence="75">
        <Pattern confidenceLevel="85">
          <IdMatch idRef="Func_croatia_eu_ssn_or_equivalent" />
          <Match idRef="Keywords_croatia_eu_ssn_or_equivalent" />
        </Pattern>
       <Pattern confidenceLevel="75">
          <IdMatch idRef="Func_croatia_eu_ssn_or_equivalent" />
        </Pattern>
</Entity>
```

Keywords

Keywords_croatia_eu_ssn_or_equivalent

- personal identification number
- master citizen number
- national identification number
- social security number
- nationalnumber#
- ssn#

- ssn
- nationalnumber
- bnn#
- bnn
- personal id number
- personalidnumber#
- oib
- osobni identifikacijski broj

Cyprus drivers license number

Format

12 digits without spaces and delimiters

Pattern

12 digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_cyprus_eu_driver's_license_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_driver's_license_number` or `Keywords_cyprus_eu_driver's_license_number` is found.

```
<!-- Cyprus Driver's License Number -->
<Entity id="356fa104-f9ac-4aff-a0e4-2e6e65ea06c4" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_cyprus_eu_driver's_license_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_driver's_license_number" />
      <Match idRef="Keywords_cyprus_eu_driver's_license_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_driver's_license_number

- driverlic
- driverlics
- driverlicense
- driverlicenses
- driverlicence
- driverlicences
- driver lic
- driver lics
- driver license
- driver licenses
- driver licence

- driver licences
- driverslic
- driverslics
- driverslicence
- driverslicences
- driverslicense
- driverslicenses
- drivers lic
- drivers lics
- drivers license
- drivers licenses
- drivers licence
- drivers licences
- driver'lic
- driver'lics
- driver'license
- driver'licenses
- driver'licence
- driver'licences
- driver' lic
- driver' lics
- driver' license
- driver' licenses
- driver' licence
- driver' licences
- driver'slic
- driver'slics
- driver'slicense
- driver'slicenses
- driver'slicence
- driver'slicences
- driver's lic
- driver's lics
- driver's license
- driver's licenses
- driver's licence
- driver's licences
- dl#
- dls#
- driverlic#
- driverlics#
- driverlicense#
- driverlicenses#
- driverlicence#
- driverlicences#
- driver lic#

- driver lics#
- driver license#
- driver licenses#
- driver licences#
- driverslic#
- driverslics#
- driverslicense#
- driverslicenses#
- driverslicence#
- driverslicences#
- drivers lic#
- drivers lics#
- drivers license#
- drivers licenses#
- drivers licence#
- drivers licences#
- driver'lic#
- driver'lics#
- driver'license#
- driver'licenses#
- driver'licence#
- driver'licences#
- driver' lic#
- driver' lics#
- driver' license#
- driver' licenses#
- driver' licence#
- driver' licences#
- driver'slic#
- driver'slics#
- driver'slicense#
- driver'slicenses#
- driver'slicence#
- driver'slicences#
- driver's lic#
- driver's lics#
- driver's license#
- driver's licenses#
- driver's licence#
- driver's licences#
- driving licence
- driving license
- dlno#
- driv lic
- driv licen
- driv license

- driv licenses
- driv licence
- driv licences
- driver licen
- drivers licen
- driver's licen
- driving lic
- driving licen
- driving licenses
- driving licence
- driving licences
- driving permit
- dl no
- dlno
- dl number

Keywords_cyprus_eu_driver's_license_number

- άδεια οδήγησης
- αριθμό άδειας οδήγησης
- άδειες οδήγησης

Cyprus identity card

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

10 digits without spaces and delimiters

Pattern

10 digits

Checksum

not applicable

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_cyprus_eu_national_id_card` finds content that matches the pattern.
- A keyword from `Keywords_cyprus_eu_national_id_card` is found.

```
<!-- Cyprus Identity Card -->
<Entity id="3ba8afe5-7a6c-4929-8247-0001b6878438" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_cyprus_eu_national_id_card" />
    <Match idRef="Keywords_cyprus_eu_national_id_card" />
  </Pattern>
</Entity>
```

Keywords

Keywords_cyprus_eu_national_id_card

- id card number
- identity card number
- kimlik karti
- national identification number
- personal id number
- ταυτοτητασ

Cyprus passport number

Format

one letter followed by 6-8 digits with no spaces or delimiters

Pattern

one letter followed by six to eight digits

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_cyprus_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_cyprus_eu_passport_number` is found.
- The regular expression `Regex_cyprus_eu_passport_date` finds date in the format DD/MM/YYYY or a keyword from `Keywords_cyprus_eu_passport_date` is found

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_cyprus_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_cyprus_eu_passport_number` is found.

```
<!-- Cyprus Passport Number -->
<Entity id="9193e2e8-7f8c-43c1-a274-ac40d651936f" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Regex_cyprus_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_cyprus_eu_passport_number" />
    </Any>
    <Any minMatches="1">
      <Match idRef="Regex_cyprus_eu_passport_date" />
      <Match idRef="Keywords_cyprus_eu_passport_date" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_cyprus_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_cyprus_eu_passport_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_passport_number_common

- passport#
- passport #
- passportid
- passports
- passportno
- passport no
- passportnumber
- passport number
- passportnumbers
- passport numbers

Keywords_cyprus_eu_passport_number

- αριθμό διαβατηρίου
- pasaportu
- Αριθμός Διαβατηρίου
- κυπριακό διαβατήριο
- διαβατήριο#
- διαβατήριο
- αριθμός διαβατηρίου
- Pasaport Kimliği
- pasaport numarası
- Pasaport no.
- Αρ. Διαβατηρίου

Keywords_cyprus_eu_passport_date

- expires on
- issued on

Cyprus tax identification number

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

eight digits and one letter in the specified pattern

Pattern

eight digits and one letter:

- a "0" or "9"
- seven digits
- one letter (not case-sensitive)

Checksum

not applicable

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_cyprus_eu_tax_file_number` finds content that matches the pattern.
- A keyword from `Keywords_cyprus_eu_tax_file_number` is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_cyprus_eu_tax_file_number` finds content that matches the pattern.

```
<!-- Cyprus Tax Identification Number -->
<Entity id="40e64bd9-55f3-4a09-9bd6-1db18dced9dd" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_cyprus_eu_tax_file_number" />
    <Match idRef="Keywords_cyprus_eu_tax_file_number" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_cyprus_eu_tax_file_number" />
  </Pattern>
</Entity>
```

Keywords

Keywords_cyprus_eu_tax_file_number

- tax id
- tax identification code
- tax identification no
- tax identification number
- tax no#
- tax no
- tax number
- tax registration number
- taxid#

- taxidno#
- taxidnumber#
- taxno#
- taxnumber#
- taxnumber
- tic#
- tic
- tin id
- tin no
- tin#
- vergi kimlik kodu
- vergi kimlik numarası
- αριθμός φορολογικού μητρώου
- κωδικός φορολογικού μητρώου
- φορολογική ταυτότητα
- φορολογικού κωδικού

Czech driver's license number

Format

two letters followed by six digits

Pattern

eight letters and digits:

- letter 'E' (not case-sensitive)
- a letter
- a space (optional)
- six digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_czech_republic_eu_driver's_license_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_driver's_license_number` or `Keywords_czech_republic_eu_driver's_license_number` is found.

```
<Entity id="86b40d3b-d8ea-4c36-aab0-ef9416a6769c" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_czech_republic_eu_driver's_license_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_driver's_license_number" />
      <Match idRef="Keywords_czech_republic_eu_driver's_license_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_driver's_license_number

- driverlic
- driverlics
- driverlicense
- driverlicenses
- driverlicence
- driverlicences
- driver lic
- driver lics
- driver license
- driver licenses
- driver licence
- driver licences
- driverslic
- driverslics
- driverslicence
- driverslicences
- driverslicense
- driverslicenses
- drivers lic
- drivers lics
- drivers license
- drivers licenses
- drivers licence
- drivers licences
- driver'lic
- driver'lics
- driver'license
- driver'licenses
- driver'licence
- driver'licences
- driver' lic
- driver' lics
- driver' license
- driver' licenses
- driver' licence
- driver' licences
- driver'slic
- driver'slics
- driver'slicense
- driver'slicenses
- driver'slicence
- driver'slicences
- driver's lic
- driver's lics
- driver's license

- driver's licenses
- driver's licence
- driver's licences
- dl#
- dls#
- driverlic#
- driverlics#
- driverlicense#
- driverlicenses#
- driverlicence#
- driverlicences#
- driver lic#
- driver lics#
- driver license#
- driver licenses#
- driver licences#
- driverslic#
- driverslics#
- driverslicense#
- driverslicenses#
- driverslicence#
- driverslicences#
- drivers lic#
- drivers lics#
- drivers license#
- drivers licenses#
- drivers licence#
- drivers licences#
- driver'lic#
- driver'lics#
- driver'license#
- driver'licenses#
- driver'licence#
- driver'licences#
- driver' lic#
- driver' lics#
- driver' license#
- driver' licenses#
- driver' licence#
- driver' licences#
- driver'slic#
- driver'slics#
- driver'slicense#
- driver'slicenses#
- driver'slicence#
- driver'slicences#

- driver's lic#
- driver's lics#
- driver's license#
- driver's licenses#
- driver's licence#
- driver's licences#
- driving licence
- driving license
- dlno#
- driv lic
- driv licen
- driv license
- driv licenses
- driv licence
- driv licences
- driver licen
- drivers licen
- driver's licen
- driving lic
- driving licen
- driving licenses
- driving licence
- driving licences
- driving permit
- dl no
- dlno
- dl number

Keywords_czech_republic_eu_driver's_license_number

- řidičský prúkaz
- řidičské průkazy
- číslo řidičského průkazu
- čísla řidičských průkazů

Czech passport number

Format

eight digits without spaces or delimiters

Pattern

eight digits without spaces or delimiters

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_czech_republic_eu_passport_number` finds content that matches the pattern.

- A keyword from `Keywords_eu_passport_number` or `Keywords_czech_republic_eu_passport_number` is found.
- The regular expression `Regex_eu_passport_date1` finds date in the format DD.MM.YYYY or a keyword from `Keywords_eu_passport_date` is found

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_czech_republic_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_czech_republic_eu_passport_number` is found.

```
<!-- Czech Republic Passport Number -->
<Entity id="7bcd8ce8-5e92-4bbe-bc92-fa669f0369fa" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Regex_czech_republic_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_czech_republic_eu_passport_number" />
    </Any>
    <Any minMatches="1">
      <Match idRef="Regex_eu_passport_date1" />
      <Match idRef="Keywords_eu_passport_date" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_czech_republic_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_czech_republic_eu_passport_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_passport_number_common

- passport#
- passport #
- passportid
- passports
- passportno
- passport no
- passportnumber
- passport number
- passportnumbers
- passport numbers

Keywords_czech_republic_eu_passport_number

- cestovní pas
- číslo pasu
- cestovní pasu
- passeport no
- čísla pasu

Keywords_eu_passport_date

- date of issue
- date of expiry

Czech personal identity number

Format

nine digits with optional forward slash (old format) 10 digits with optional forward slash (new format)

Pattern

nine digits (old format):

- six digits that represent date of birth
- an optional forward slash
- three digits

10 digits (new format):

- six digits that represent date of birth
- an optional forward slash
- four digits where last digit is a check digit

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_czech_id_card finds content that matches the pattern.
- A keyword from Keyword_czech_id_card is found.
- The checksum passes.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_czech_id_card_new_format finds content that matches the pattern.
- The checksum passes.

```
<!-- Czech Personal Identity Number -->
      <!-- Czech Personal Identity Number -->
      <Entity id="60c0725a-4eb6-455b-9dda-05d8a7396497" patternsProximity="300" recommendedConfidence="85">
        <Pattern confidenceLevel="85">
          <IdMatch idRef="Func_czech_id_card" />
          <Match idRef="Keyword_czech_id_card" />
        </Pattern>
        <Version minEngineVersion="15.20.3000.000">
          <Pattern confidenceLevel="75">
            <IdMatch idRef="Func_czech_id_card_new_format" />
          </Pattern>
        </Version>
      </Entity>
```

Keywords

Keyword_czech_id_card

- birth number
- czech republic id
- czechidno#
- daňové číslo
- identifikační číslo

- identity no
- identity number
- identityno#
- identityno
- insurance number
- national identification number
- nationalnumber#
- national number
- osobní číslo
- personalidnumber#
- personal id number
- personal identification number
- personal number
- pid#
- pid
- pojištění číslo
- rč
- rodne cislo
- rodné číslo
- ssn
- ssn#
- social security number
- tax id
- tax identification no
- tax identification number
- tax no#
- tax no
- tax number
- tax registration number
- taxid#
- taxidno#
- taxidnumber#
- taxno#
- taxnumber#
- taxnumber
- tin id
- tin no
- tin#
- unique identification number

Czech social security number or equivalent identification

This sensitive information type entity is only available in the EU Social Security Number or Equivalent ID sensitive information type.

Format

10 digits and a backslash in the specified pattern

Pattern

10 digits and a backslash:

- six digits that correspond to the birth date (YYMMDD):
- a backslash
- three digits that correspond to a serial number that separates persons born on the same date
- one check digit

Checksum

Yes

Definition

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_czech_republic_eu_ssn_or_equivalent` finds content that matches the pattern.
- A keyword from `Keywords_czech_republic_eu_ssn_or_equivalent` is found.

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_czech_republic_eu_ssn_or_equivalent` finds content that matches the pattern.

```
 <!-- EU SSN or Equivalent Number -->
<Entity id="d24e32a4-c0bb-4ba8-899d-6303b95742d9" patternsProximity="300" recommendedConfidence="75">
        <Pattern confidenceLevel="85">
          <IdMatch idRef="Func_czech_republic_eu_ssn_or_equivalent" />
          <Match idRef="Keywords_czech_republic_eu_ssn_or_equivalent" />
        </Pattern>
       <Pattern confidenceLevel="75">
          <IdMatch idRef="Func_czech_republic_eu_ssn_or_equivalent" />
        </Pattern>
</Entity>
```

Keywords

Keywords_czech_republic_eu_ssn_or_equivalent

- birth number
- national identification number
- personal identification number
- social security number
- nationalnumber#
- ssn#
- ssn
- national number
- personal id number
- personalidnumber#
- rč
- rodné číslo
- rodne cislo

Denmark driver's license number

Format

eight digits without spaces and delimiters

Pattern

eight digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_denmark_eu_driver's_license_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_driver's_license_number` or `Keywords_denmark_eu_driver's_license_number` is found.

```
<!-- Denmark Driver's License Number -->
<Entity id="98a95812-6203-451a-a220-d39870ebef0e" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_denmark_eu_driver's_license_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_driver's_license_number" />
      <Match idRef="Keywords_denmark_eu_driver's_license_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_driver's_license_number

- driverlic
- driverlics
- driverlicense
- driverlicenses
- driverlicence
- driverlicences
- driver lic
- driver lics
- driver license
- driver licenses
- driver licence
- driver licences
- driverslic
- driverslics
- driverslicence
- driverslicences
- driverslicense
- driverslicenses
- drivers lic
- drivers lics
- drivers license
- drivers licenses
- drivers licence

- drivers licences
- driver'lic
- driver'lics
- driver'license
- driver'licenses
- driver'licence
- driver'licences
- driver' lic
- driver' lics
- driver' license
- driver' licenses
- driver' licence
- driver' licences
- driver'slic
- driver'slics
- driver'slicense
- driver'slicenses
- driver'slicence
- driver'slicences
- driver's lic
- driver's lics
- driver's license
- driver's licenses
- driver's licence
- driver's licences
- dl#
- dls#
- driverlic#
- driverlics#
- driverlicense#
- driverlicenses#
- driverlicence#
- driverlicences#
- driver lic#
- driver lics#
- driver license#
- driver licenses#
- driver licences#
- driverslic#
- driverslics#
- driverslicense#
- driverslicenses#
- driverslicence#
- driverslicences#
- drivers lic#
- drivers lics#

- drivers license#
- drivers licenses#
- drivers licence#
- drivers licences#
- driver'lic#
- driver'lics#
- driver'license#
- driver'licenses#
- driver'licence#
- driver'licences#
- driver' lic#
- driver' lics#
- driver' license#
- driver' licenses#
- driver' licence#
- driver' licences#
- driver'slic#
- driver'slics#
- driver'slicense#
- driver'slicenses#
- driver'slicence#
- driver'slicences#
- driver's lic#
- driver's lics#
- driver's license#
- driver's licenses#
- driver's licence#
- driver's licences#
- driving licence
- driving license
- dlno#
- driv lic
- driv licen
- driv license
- driv licenses
- driv licence
- driv licences
- driver licen
- drivers licen
- driver's licen
- driving lic
- driving licen
- driving licenses
- driving licence
- driving licences
- driving permit

- dl no
- dlno
- dl number

Keywords_denmark_eu_driver's_license_number

- kørekort
- kørekortnummer

Denmark passport number

Format

nine digits without spaces and delimiters

Pattern

nine digits

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_denmark_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_denmark_eu_passport_number` is found.
- The regular expression `Regex_eu_passport_date2` finds date in the format DD MM YY or a keyword from `Keywords_eu_passport_date` is found

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_denmark_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_denmark_eu_passport_number` is found.

```
<!-- Denmark Passport Number -->
<Entity id="25e8c47e-e6fe-4884-a211-74898f8c0196" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Regex_denmark_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_denmark_eu_passport_number" />
    </Any>
    <Any minMatches="1">
      <Match idRef="Regex_eu_passport_date2" />
      <Match idRef="Keywords_eu_passport_date" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_denmark_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_denmark_eu_passport_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_passport_number_common

- passport#
- passport #
- passportid
- passports
- passportno
- passport no
- passportnumber
- passport number
- passportnumbers
- passport numbers

Keywords_denmark_eu_passport_number

- pasnummer
- Passeport n°
- pasnumre

Keywords_eu_passport_date

- date of issue
- date of expiry

Denmark personal identification number

Format

10 digits containing a hyphen

Pattern

10 digits:

- six digits in the format DDMMYY, which are the date of birth
- a hyphen
- four digits where the final digit is a check digit

Checksum

Yes

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Func_denmark_eu_tax_file_number finds content that matches the pattern.
- A keyword from Keyword_denmark_id is found.
- The checksum passes.

A DLP policy has low confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Func_denmark_eu_tax_file_number finds content that matches the pattern.
- The checksum passes.

```
<!-- Denmark Personal Identification Number -->
   <!-- Denmark Personal Identification Number -->
      <Entity id="6c4f2fef-56e1-4c00-8093-88d7a01cf460" patternsProximity="300" recommendedConfidence="75">
        <Pattern confidenceLevel="75">
          <IdMatch idRef="Func_denmark_eu_tax_file_number" />
          <Match idRef="Keyword_denmark_id" />
        </Pattern>
        <Pattern confidenceLevel="65">
          <IdMatch idRef="Func_denmark_eu_tax_file_number" />
        </Pattern>
      </Entity>
```

Keywords

Keyword_denmark_id

- centrale personregister
- civilt registreringssystem
- cpr
- cpr#
- gesundheitskarte nummer
- gesundheitsversicherungkarte nummer
- health card
- health insurance card number
- health insurance number
- identification number
- identifikationsnummer
- identifikationsnummer#
- identity number
- krankenkassennummer
- nationalid#
- nationalnumber#
- national number
- personalidnumber#
- personalidentityno#
- personal id number
- personnummer
- personnummer#
- reisekrankenversicherungskartenummer
- rejsesygesikringskort
- ssn
- ssn#
- skat id
- skat kode
- skat nummer
- skattenummer
- social security number
- sundhedsforsikringskort
- sundhedsforsikringsnummer
- sundhedskort
- sundhedskortnummer

- sygesikring
- sygesikringkortnummer
- tax code
- travel health insurance card
- uniqueidentityno#
- tax number
- tax registration number
- tax id
- tax identification number
- taxid#
- taxnumber#
- tax no
- taxno#
- taxnumber
- tax identification no
- tin#
- taxidno#
- taxidnumber#
- tax no#
- tin id
- tin no
- cpr.nr
- cprnr
- cprnummer
- personnr
- personregister
- sygesikringsbevis
- sygesikringsbevisnr
- sygesikringsbevisnummer
- sygesikringskort
- sygesikringskortnr
- sygesikringskortnummer
- sygesikringsnr
- sygesikringsnummer

Denmark social security number or equivalent identification

This sensitive information type entity is only available the EU Social Security Number or Equivalent ID sensitive information type.

Format

10 digits and a hyphen in the specified pattern

Pattern

10 digits and a hyphen:

- six digits that correspond to the birth date (DDMMYY)
- a hyphen
- four digits that correspond to a sequence number

Checksum

Yes

Definition

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_denmark_eu_ssn_or_equivalent` finds content that matches the pattern.
- A keyword from `Keywords_denmark_eu_ssn_or_equivalent` is found.

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_denmark_eu_ssn_or_equivalent` finds content that matches the pattern.

```
 <!-- EU SSN or Equivalent Number -->
<Entity id="d24e32a4-c0bb-4ba8-899d-6303b95742d9" patternsProximity="300" recommendedConfidence="75">
        <Pattern confidenceLevel="85">
          <IdMatch idRef="Func_denmark_eu_ssn_or_equivalent" />
          <Match idRef="Keywords_denmark_eu_ssn_or_equivalent" />
        </Pattern>
       <Pattern confidenceLevel="75">
          <IdMatch idRef="Func_denmark_eu_ssn_or_equivalent" />
        </Pattern>
</Entity>
```

Keywords

Keywords_denmark_eu_ssn_or_equivalent

- personal identification number
- national identification number
- social security number
- nationalnumber#
- ssn#
- ssn
- national number
- personal id number
- personalidnumber#
- cpr-nummer
- personnummer

Drug Enforcement Agency (DEA) number

Format

two letters followed by seven digits

Pattern

Pattern must include all of the following:

- one letter (not case-sensitive) from this set of possible letters: abcdefghjklmnprstux, which is a registrant code
- one letter (not case-sensitive), which is the first letter of the registrant's last name or digit '9'
- seven digits, the last of which is the check digit

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_dea_number finds content that matches the pattern.
- A keyword from `Keyword_dea_number` is found
- The checksum passes.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_dea_number finds content that matches the pattern.
- The checksum passes.

```
<!-- DEA Number -->
<Entity id="9a5445ad-406e-43eb-8bd7-cac17ab6d0e4" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_dea_number" />
  </Pattern>
  <Version minEngineVersion="15.20.1207.000" maxEngineVersion="15.20.3134.000">
    <Pattern confidenceLevel="85">
      <IdMatch idRef="Func_dea_number" />
    </Pattern>
  </Version>
  <Version minEngineVersion="15.20.3135.000">
    <Pattern confidenceLevel="85">
      <IdMatch idRef="Func_dea_number" />
      <Match idRef="Keyword_dea_number" />
    </Pattern>
  </Version>
</Entity>
```

Keywords

Keyword_dea_number

- dea
- dea#
- drug enforcement administration
- drug enforcement agency

Estonia driver's license number

Format

two letters followed by six digits

Pattern

two letters and six digits:

- the letters "ET" (not case-sensitive)
- six digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of

300 characters:

- The regular expression `Regex_estonia_eu_driver's_license_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_driver's_license_number` or `Keywords_estonia_eu_driver's_license_number` is found.

```
<!-- Estonia Driver's License Number -->
<Entity id="51da8171-da70-4cc1-9d65-055a59ca4f83" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_estonia_eu_driver's_license_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_driver's_license_number" />
      <Match idRef="Keywords_estonia_eu_driver's_license_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_driver's_license_number

- driverlic
- driverlics
- driverlicense
- driverlicenses
- driverlicence
- driverlicences
- driver lic
- driver lics
- driver license
- driver licenses
- driver licence
- driver licences
- driverslic
- driverslics
- driverslicence
- driverslicences
- driverslicense
- driverslicenses
- drivers lic
- drivers lics
- drivers license
- drivers licenses
- drivers licence
- drivers licences
- driver'lic
- driver'lics
- driver'license
- driver'licenses
- driver'licence
- driver'licences
- driver' lic

- driver' lics
- driver' license
- driver' licenses
- driver' licence
- driver' licences
- driver'slic
- driver'slics
- driver'slicense
- driver'slicenses
- driver'slicence
- driver'slicences
- driver's lic
- driver's lics
- driver's license
- driver's licenses
- driver's licence
- driver's licences
- dl#
- dls#
- driverlic#
- driverlics#
- driverlicense#
- driverlicenses#
- driverlicence#
- driverlicences#
- driver lic#
- driver lics#
- driver license#
- driver licenses#
- driver licences#
- driverslic#
- driverslics#
- driverslicense#
- driverslicenses#
- driverslicence#
- driverslicences#
- drivers lic#
- drivers lics#
- drivers license#
- drivers licenses#
- drivers licence#
- drivers licences#
- driver'lic#
- driver'lics#
- driver'license#
- driver'licenses#

- driver'licence#
- driver'licences#
- driver' lic#
- driver' lics#
- driver' license#
- driver' licenses#
- driver' licence#
- driver' licences#
- driver'slic#
- driver'slics#
- driver'slicense#
- driver'slicenses#
- driver'slicence#
- driver'slicences#
- driver's lic#
- driver's lics#
- driver's license#
- driver's licenses#
- driver's licence#
- driver's licences#
- driving licence
- driving license
- dlno#
- driv lic
- driv licen
- driv license
- driv licenses
- driv licence
- driv licences
- driver licen
- drivers licen
- driver's licen
- driving lic
- driving licen
- driving licenses
- driving licence
- driving licences
- driving permit
- dl no
- dlno
- dl number

Keywords_estonia_eu_driver's_license_number

-- permis de conduire

- juhilubade numbrid
- juhiloa number
- juhiluba

Estonia Personal Identification Code

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

11 digits without spaces and delimiters

Pattern

11 digits:

- one digit that corresponds to sex and century of birth (odd number male, even number female; 1-2: 19th century; 3-4: 20th century; 5-6: 21st century)
- six digits that correspond to date of birth (YYMMDD)
- three digits that correspond to a serial number separating persons born on the same date
- one check digit

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_estonia_eu_national_id_card` finds content that matches the pattern.
- A keyword from `Keywords_estonia_eu_national_id_card` is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_estonia_eu_national_id_card` finds content that matches the pattern.

```
<!-- Estonia Personal Identification Code -->
<Entity id="bfb26de6-dad5-4d48-ab72-4789cdd0654c" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_estonia_eu_national_id_card" />
    <Match idRef="Keywords_estonia_eu_national_id_card" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_estonia_eu_national_id_card" />
    <Any minMatches="0" maxMatches="0">
      <Match idRef="Keywords_estonia_eu_telephone_number" />
      <Match idRef="Keywords_estonia_eu_mobile_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_estonia_eu_national_id_card

- id-kaart
- ik

- isikukood#
- isikukood
- maksu id
- maksukohustuslase identifitseerimisnumber
- maksunumber
- national identification number
- national number
- personal code
- personal id number
- personal identification code
- personal identification number
- personalidnumber#
- tax id
- tax identification no
- tax identification number
- tax no#
- tax no
- tax number
- tax registration number
- taxid#
- taxidno#
- taxidnumber#
- taxno#
- taxnumber#
- taxnumber
- tin id
- tin no
- tin#

Estonia passport number

Format

one letter followed by seven digits with no spaces or delimiters

Pattern

one letter followed by seven digits

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_estonia_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_estonia_eu_passport_number` is found.
- The regular expression `Regex_eu_passport_date1` finds date in the format DD.MM.YYYY or a keyword from `Keywords_eu_passport_date` is found

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of

300 characters:

- The regular expression `Regex_estonia_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_estonia_eu_passport_number` is found.

```
<!-- Estonia Passport Number -->
<Entity id="61f7073a-509e-425b-a754-bc01bb5d5b8c" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Regex_estonia_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_estonia_eu_passport_number" />
    </Any>
    <Any minMatches="1">
      <Match idRef="Regex_eu_passport_date1" />
      <Match idRef="Keywords_eu_passport_date" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_estonia_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_estonia_eu_passport_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_passport_number_common

- passport#
- passport #
- passportid
- passports
- passportno
- passport no
- passportnumber
- passport number
- passportnumbers
- passport numbers

Keywords_estonia_eu_passport_number

eesti kodaniku pass passi number passinumbrid document number document no dokumendi nr

Keywords_eu_passport_date

- date of issue
- date of expiry

EU debit card number

Format

16 digits

Pattern

Complex and robust pattern

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_eu_debit_card finds content that matches the pattern.
- At least one of the following is true:
  - A keyword from Keyword_eu_debit_card is found.
  - A keyword from Keyword_card_terms_dict is found.
  - A keyword from Keyword_card_security_terms_dict is found.
  - A keyword from Keyword_card_expiration_terms_dict is found.
  - The function Func_expiration_date finds a date in the right date format.
- The checksum passes.

```
<!-- EU Debit Card Number -->
<Entity id="0e9b3178-9678-47dd-a509-37222ca96b42" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_eu_debit_card" />
    <Any minMatches="1">
      <Match idRef="Keyword_eu_debit_card" />
      <Match idRef="Keyword_card_terms_dict" />
      <Match idRef="Keyword_card_security_terms_dict" />
      <Match idRef="Keyword_card_expiration_terms_dict" />
      <Match idRef="Func_expiration_date" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keyword_eu_debit_card

- account number
- card number
- card no.
- security number
- cc#

Keyword_card_terms_dict

- acct nbr
- acct num
- acct no
- american express
- americanexpress
- americano espresso
- amex
- atm card
- atm cards
- atm kaart
- atmcard
- atmcards
- atmkaart
- atmkaarten
- bancontact
- bank card

- bankkaart
- card holder
- card holders
- card num
- card number
- card numbers
- card type
- cardano numerico
- cardholder
- cardholders
- cardnumber
- cardnumbers
- carta bianca
- carta credito
- carta di credito
- cartao de credito
- cartao de crédito
- cartao de debito
- cartao de débito
- carte bancaire
- carte blanche
- carte bleue
- carte de credit
- carte de crédit
- carte di credito
- carteblanche
- cartão de credito
- cartão de crédito
- cartão de debito
- cartão de débito
- cb
- ccn
- check card
- check cards
- checkcard
- checkcards
- chequekaart
- cirrus
- cirrus-edc-maestro
- controlekaart
- controlekaarten
- credit card
- credit cards
- creditcard
- creditcards
- debetkaart

- debetkaarten
- debit card
- debit cards
- debitcard
- debitcards
- debito automatico
- diners club
- dinersclub
- discover
- discover card
- discover cards
- discovercard
- discovercards
- débito automático
- edc
- eigentümername
- european debit card
- hoofdkaart
- hoofdkaarten
- in viaggio
- japanese card bureau
- japanse kaartdienst
- jcb
- kaart
- kaart num
- kaartaantal
- kaartaantallen
- kaarthouder
- kaarthouders
- karte
- karteninhaber
- karteninhabers
- kartennr
- kartennummer
- kreditkarte
- kreditkarten-nummer
- kreditkarteninhaber
- kreditkarteninstitut
- kreditkartennummer
- kreditkartentyp
- maestro
- master card
- master cards
- mastercard
- mastercards
- mc

- mister cash
- n carta
- carta
- no de tarjeta
- no do cartao
- no do cartão
- no. de tarjeta
- no. do cartao
- no. do cartão
- nr carta
- nr. carta
- numeri di scheda
- numero carta
- numero de cartao
- numero de carte
- numero de cartão
- numero de tarjeta
- numero della carta
- numero di carta
- numero di scheda
- numero do cartao
- numero do cartão
- numéro de carte
- nº carta
- nº de carte
- nº de la carte
- nº de tarjeta
- nº do cartao
- nº do cartão
- nº. do cartão
- número de cartao
- número de cartão
- número de tarjeta
- número do cartao
- scheda dell'assegno
- scheda dell'atmosfera
- scheda dell'atmosfera
- scheda della banca
- scheda di controllo
- scheda di debito
- scheda matrice
- schede dell'atmosfera
- schede di controllo
- schede di debito
- schede matrici
- scoprono la scheda

- scoprono le schede
- solo
- supporti di scheda
- supporto di scheda
- switch
- tarjeta atm
- tarjeta credito
- tarjeta de atm
- tarjeta de credito
- tarjeta de debito
- tarjeta debito
- tarjeta no
- tarjetahabiente
- tipo della scheda
- ufficio giapponese della
- scheda
- v pay
- v-pay
- visa
- visa plus
- visa electron
- visto
- visum
- vpay

Keyword_card_security_terms_dict

- card identification number
- card verification
- cardi la verifica
- cid
- cod seg
- cod seguranca
- cod segurança
- cod sicurezza
- cod. seg
- cod. seguranca
- cod. segurança
- cod. sicurezza
- codice di sicurezza
- codice di verifica
- codigo
- codigo de seguranca
- codigo de segurança
- crittogramma
- cryptogram
- cryptogramme
- cv2

- cvc
- cvc2
- cvn
- cvv
- cvv2
- cód seguranca
- cód segurança
- cód. seguranca
- cód. segurança
- código
- código de seguranca
- código de segurança
- de kaart controle
- geeft nr uit
- issue no
- issue number
- kaartidentificatienummer
- kreditkartenprufnummer
- kreditkartenprüfnummer
- kwestieaantal
- no. dell'edizione
- no. di sicurezza
- numero de securite
- numero de verificacao
- numero dell'edizione
- numero di identificazione della
- scheda
- numero di sicurezza
- numero van veiligheid
- numéro de sécurité
- nº autorizzazione
- número de verificação
- perno il blocco
- pin block
- prufziffer
- prüfziffer
- security code
- security no
- security number
- sicherheits kode
- sicherheitscode
- sicherheitsnummer
- speldblok
- veiligheid nr
- veiligheidsaantal
- veiligheidscode

- veiligheidsnummer
- verfalldatum

Keyword_card_expiration_terms_dict

- ablauf
- data de expiracao
- data de expiração
- data del exp
- data di exp
- data di scadenza
- data em que expira
- data scad
- data scadenza
- date de validité
- datum afloop
- datum van exp
- de afloop
- espira
- espira
- exp date
- exp datum
- expiration
- expire
- expires
- expiry
- fecha de expiracion
- fecha de venc
- gultig bis
- gultigkeitsdatum
- gültig bis
- gültigkeitsdatum
- la scadenza
- scadenza
- valable
- validade
- valido hasta
- valor
- venc
- vencimento
- vencimiento
- verloopt
- vervaldag
- vervaldatum
- vto
- válido hasta

EU driver's license number

These entities are in the EU Driver's License Number and are sensitive information types.

- Austria
- Belgium
- Bulgaria
- Croatia
- Cyprus
- Czech
- Denmark
- Estonia
- Finland
- France
- Germany
- Greece
- Hungary
- Ireland
- Italy
- Latvia
- Lithuania
- Luxemburg
- Malta
- Netherlands
- Poland
- Portugal
- Romania
- Slovakia
- Slovenia
- Spain
- Sweden
- U.K.

EU national identification number

These entities are in the EU National Identification Number and are sensitive information types.

- Austria
- Belgium
- Bulgaria
- Croatia
- Cyprus
- Czech
- Denmark
- Estonia
- Finland
- France
- Germany
- Greece
- Hungary

- Ireland
- Italy
- Latvia
- Lithuania
- Luxemburg
- Malta
- Netherlands
- Poland
- Portugal
- Romania
- Slovakia
- Slovenia
- Spain
- U.K.

EU passport number

These entities are in the EU passport number and are sensitive information types. These entities are in the EU passport number bundle.

- Austria
- Belgium
- Bulgaria
- Croatia
- Cyprus
- Czech
- Denmark
- Estonia
- Finland
- France
- Germany
- Greece
- Hungary
- Ireland
- Italy
- Latvia
- Lithuania
- Luxemburg
- Malta
- Netherlands
- Poland
- Portugal
- Romania
- Slovakia
- Slovenia
- Spain
- Sweden
- U.K.

EU social security number or equivalent identification

These entities that are in the EU Social Security Number or equivalent identification and are sensitive information types.

- Austria
- Belgium
- Croatia
- Czech
- Denmark
- Finland
- France
- Germany
- Greece
- Hungary
- Portugal
- Spain
- Sweden

EU Tax identification number

These entities are in the EU Tax identification number sensitive information type.

- Austria
- Belgium
- Bulgaria
- Croatia
- Cyprus
- Czech
- Denmark
- Estonia
- Finland
- France
- Germany
- Greece
- Hungary
- Ireland
- Italy
- Latvia
- Lithuania
- Luxemburg
- Malta
- Netherlands
- Poland
- Portugal
- Romania
- Slovakia
- Slovenia
- Spain

- Sweden
- U.K.

Finland driver's license number

Format

10 digits containing a hyphen

Pattern

10 digits containing a hyphen:

- six digits
- a hyphen
- three digits
- a digit or letter

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_finland_eu_driver's_license_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_driver's_license_number` or `Keywords_finland_eu_driver's_license_number` is found.

```
<!-- Finland Driver's License Number -->
<Entity id="bb3b27a3-79bd-4ac4-81a7-f9fca3c7d1a7" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_finland_eu_driver's_license_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_driver's_license_number" />
      <Match idRef="Keywords_finland_eu_driver's_license_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_driver's_license_number

- driverlic
- driverlics
- driverlicense
- driverlicenses
- driverlicence
- driverlicences
- driver lic
- driver lics
- driver license
- driver licenses
- driver licence
- driver licences
- driverslic

- driverslics
- driverslicence
- driverslicences
- driverslicense
- driverslicenses
- drivers lic
- drivers lics
- drivers license
- drivers licenses
- drivers licence
- drivers licences
- driver'lic
- driver'lics
- driver'license
- driver'licenses
- driver'licence
- driver'licences
- driver' lic
- driver' lics
- driver' license
- driver' licenses
- driver' licence
- driver' licences
- driver'slic
- driver'slics
- driver'slicense
- driver'slicenses
- driver'slicence
- driver'slicences
- driver's lic
- driver's lics
- driver's license
- driver's licenses
- driver's licence
- driver's licences
- dl#
- dls#
- driverlic#
- driverlics#
- driverlicense#
- driverlicenses#
- driverlicence#
- driverlicences#
- driver lic#
- driver lics#
- driver license#

- driver licenses#
- driver licences#
- driverslic#
- driverslics#
- driverslicense#
- driverslicenses#
- driverslicence#
- driverslicences#
- drivers lic#
- drivers lics#
- drivers license#
- drivers licenses#
- drivers licence#
- drivers licences#
- driver'lic#
- driver'lics#
- driver'license#
- driver'licenses#
- driver'licence#
- driver'licences#
- driver' lic#
- driver' lics#
- driver' license#
- driver' licenses#
- driver' licence#
- driver' licences#
- driver'slic#
- driver'slics#
- driver'slicense#
- driver'slicenses#
- driver'slicence#
- driver'slicences#
- driver's lic#
- driver's lics#
- driver's license#
- driver's licenses#
- driver's licence#
- driver's licences#
- driving licence
- driving license
- dlno#
- driv lic
- driv licen
- driv license
- driv licenses
- driv licence

- driv licences
- driver licen
- drivers licen
- driver's licen
- driving lic
- driving licen
- driving licenses
- driving licence
- driving licences
- driving permit
- dl no
- dlno
- dl number

Keywords_finland_eu_driver's_license_number

- ajokortti
- permis de conduire
- ajokortin numero
- kuljettaja lic.
- körkort
- körkortnummer
- förare lic.
- ajokortit
- ajokortin numerot

Finland european health insurance number

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

20-digit number

Pattern

20-digit number:

- 10 digits - 8024680246
- an optional space or hyphen
- 10 digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regex Regex_Finland_European_Health_Insurance_Number finds content that matches the pattern.
- A keyword from Keyword_Finland_European_Health_Insurance_Number is found.

```
<!-- Finland European Health Insurance Number -->
<Entity id="60f75aed-81bf-4625-89b0-0846b9248ee7" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_Finland_European_Health_Insurance_Number"/>
    <Match idRef="Keyword_Finland_European_Health_Insurance_Number"/>
  </Pattern>
</Entity>
```

Keywords

Keyword_finland_european_health_insurance_number

- ehic#
- ehic
- finlandehicnumber#
- finska sjukförsäkringskort
- health card
- health insurance card
- health insurance number
- hälsokort
- sairaanhoitokortin
- sairausvakuutuskortti
- sairausvakuutusnumero
- sjukförsäkring nummer
- sjukförsäkringskort
- suomen sairausvakuutuskortti
- terveyskortti

Finland national ID

Format

six digits plus a character indicating a century plus three digits plus a check digit

Pattern

Pattern must include all of the following:

- six digits in the format DDMMYY, which are a date of birth
- century marker (either '-', '+' or 'a')
- three-digit personal identification number
- a digit or letter (case insensitive) which is a check digit

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- the function Func_finnish_national_id finds content that matches the pattern
- a keyword from Keyword_finnish_national_id is found
- the checksum passes

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- the function Func_finnish_national_id finds content that matches the pattern
- the checksum passes

```
<!-- Finnish National ID-->
<Entity id="338FD995-4CB5-4F87-AD35-79BD1DD926C1" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_finnish_national_id" />
    <Match idRef="Keyword_finnish_national_id" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_finnish_national_id" />
  </Pattern>
</Entity>
```

Keywords

- ainutlaatuinen henkilökohtainen tunnus
- henkilökohtainen tunnus
- henkilötunnus
- henkilötunnusnumero#
- henkilötunnusnumero
- hetu
- id no
- id number
- identification number
- identiteetti numero
- identity number
- idnumber
- kansallinen henkilötunnus
- kansallisen henkilökortin
- national id card
- national id no.
- personal id
- personal identity code
- personalidnumber#
- personbeteckning
- personnummer
- social security number
- sosiaaliturvatunnus
- tax id
- tax identification no
- tax identification number
- tax no#
- tax no
- tax number
- tax registration number
- taxid#
- taxidno#

- taxidnumber#
- taxno#
- taxnumber#
- taxnumber
- tin id
- tin no
- tin#
- tunnistenumero
- tunnus numero
- tunnusluku
- tunnusnumero
- verokortti
- veronumero
- verotunniste
- verotunnus

Finland passport number

This sensitive information type entity is available in the EU Passport Number sensitive information type and is available as a stand-alone sensitive information type entity.

Format

combination of nine letters and digits

Pattern

combination of nine letters and digits:

- two letters (not case-sensitive)
- seven digits

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_finland_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keyword_finland_passport_number` is found.
- The regular expression `Regex_eu_passport_date1` finds date in the format DD.MM.YYYY or a keyword from `Keywords_eu_passport_date` is found

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_finland_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keyword_finland_passport_number` is found.

```
<!-- Finland Passport Number -->
<Entity id="d1685ac3-1d3a-40f8-8198-32ef5669c7a5" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Regex_finland_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keyword_finland_passport_number" />
    </Any>
    <Any minMatches="1">
      <Match idRef="Regex_eu_passport_date1" />
      <Match idRef="Keywords_eu_passport_date" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_finland_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keyword_finland_passport_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_passport_number

- passport#
- passport #
- passportid
- passports
- passportno
- passport no
- passportnumber
- passport number
- passportnumbers
- passport numbers

Keyword_finland_passport_number

- suomalainen passi
- passin numero
- passin numero.#
- passin numero#
- passin numero.
- passi#
- passi number

Keywords_eu_passport_date

- date of issue
- date of expiry

France driver's license number

This sensitive information type entity is available in the EU Driver's License Number sensitive information type and is available as a stand-alone sensitive information type entity.

Format

12 digits

Pattern

12 digits with validation to discount similar patterns such as French telephone numbers

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- the function Func_french_drivers_license finds content that matches the pattern.
- a keyword from Keyword_french_drivers_license is found.

```
<!-- France Driver's License Number -->
<Entity id="18e55a36-a01b-4b0f-943d-dc10282a1824" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_french_drivers_license" />
    <Match idRef="Keyword_french_drivers_license" />
  </Pattern>
</Entity>
```

Keywords

Keyword_french_drivers_license

- driverlic
- driverlics
- driverlicense
- driverlicenses
- driverlicence
- driverlicences
- driver lic
- driver lics
- driver license
- driver licenses
- driver licence
- driver licences
- driverslic
- driverslics
- driverslicence
- driverslicences
- driverslicense
- driverslicenses
- drivers lic
- drivers lics
- drivers license
- drivers licenses
- drivers licence
- drivers licences
- driver'lic
- driver'lics
- driver'license

- driver'licenses
- driver'licence
- driver'licences
- driver' lic
- driver' lics
- driver' license
- driver' licenses
- driver' licence
- driver' licences
- driver'slic
- driver'slics
- driver'slicense
- driver'slicenses
- driver'slicence
- driver'slicences
- driver's lic
- driver's lics
- driver's license
- driver's licenses
- driver's licence
- driver's licences
- dl#
- dls#
- driverlic#
- driverlics#
- driverlicense#
- driverlicenses#
- driverlicence#
- driverlicences#
- driver lic#
- driver lics#
- driver license#
- driver licenses#
- driver licences#
- driverslic#
- driverslics#
- driverslicense#
- driverslicenses#
- driverslicence#
- driverslicences#
- drivers lic#
- drivers lics#
- drivers license#
- drivers licenses#
- drivers licence#
- drivers licences#

- driver'lic#
- driver'lics#
- driver'license#
- driver'licenses#
- driver'licence#
- driver'licences#
- driver' lic#
- driver' lics#
- driver' license#
- driver' licenses#
- driver' licence#
- driver' licences#
- driver'slic#
- driver'slics#
- driver'slicense#
- driver'slicenses#
- driver'slicence#
- driver'slicences#
- driver's lic#
- driver's lics#
- driver's license#
- driver's licenses#
- driver's licence#
- driver's licences#
- driving licence
- driving license
- dlno#
- driv lic
- driv licen
- driv license
- driv licenses
- driv licence
- driv licences
- driver licen
- drivers licen
- driver's licen
- driving lic
- driving licen
- driving licenses
- driving licence
- driving licences
- driving permit
- dl no
- dlno
- dl number
- permis de conduire

- licence number
- license number
- licence numbers
- license numbers
- numéros de licence

France health insurance number

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

21-digit number

Pattern

21-digit number:

- 10 digits
- an optional space
- 10 digits
- an optional space
- a digit

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- the regex Regex_France_Health_Insurance_Number finds content that matches the pattern.
- a keyword from Keyword_France_Health_Insurance_Number is found.

```
<!-- France Health Insurance Number -->
<Entity id="9bc2069e-76df-4ff9-ac02-2f519469e236" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_France_Health_Insurance_Number"/>
    <Match idRef="Keyword_France_Health_Insurance_Number"/>
  </Pattern>
</Entity>
```

Keywords

Keyword_France_health_insurance_number

- insurance card
- carte vitale
- carte d'assuré social

France national id card (CNI)

Format

12 digits

Pattern

12 digits

Checksum

No

Definition

A DLP policy has low confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_france_cni finds content that matches the pattern.
- A keyword from Keywords_france_eu_national_id_card is found.

```
<!-- France CNI -->
<Entity id="f741ac74-1bc0-4665-b69b-f0c7f927c0c4" patternsProximity="300" recommendedConfidence="65">
  <Pattern confidenceLevel="65">
    <IdMatch idRef="Regex_france_cni" />
    <Match idRef="Keywords_france_eu_national_id_card" />
  </Pattern>
</Entity>
```

Keywords

Keywords_france_eu_national_id_card

- card number
- carte nationale d'identité
- carte nationale d'idenite no
- cni#
- cni
- compte bancaire
- national identification number
- national identity
- nationalidno#
- numéro d'assurance maladie
- numéro de carte vitale

France passport number

This sensitive information type entity is available in the EU Passport Number sensitive information type. It's available as a stand-alone sensitive information type entity.

Format

nine digits and letters

Pattern

nine digits and letters:

- two digits
- two letters (not case-sensitive)
- five digits

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_fr_passport` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_france_eu_passport_number` is found.
- The regular expression `Regex_eu_passport_date3` finds date in the format DD MM YYYY or a keyword from `Keywords_eu_passport_date` is found

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_fr_passport` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_france_eu_passport_number` is found.

```
<!-- France Passport Number -->
<Entity id="3008b884-8c8c-4cd8-a289-99f34fc7ff5d" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
      <IdMatch idRef="Func_fr_passport" />
      <Any minMatches="1">
        <Match idRef="Keywords_eu_passport_number" />
        <Match idRef="Keywords_france_eu_passport_number" />
      </Any>
      <Any minMatches="1">
        <Match idRef="Regex_eu_passport_date3" />
        <Match idRef="Keywords_eu_passport_date" />
      </Any>
    </Pattern>
    <Pattern confidenceLevel="75">
      <IdMatch idRef="Func_fr_passport" />
      <Any minMatches="1">
        <Match idRef="Keywords_eu_passport_number" />
        <Match idRef="Keywords_france_eu_passport_number" />
      </Any>
    </Pattern>
</Entity>
```

Keywords

Keywords_eu_passport_number

- passport#
- passport #
- passportid
- passports
- passportno
- passport no
- passportnumber
- passport number
- passportnumbers
- passport numbers

Keywords_france_eu_passport_number

- numéro de passeport
- passeport n °
- passeport non

- passeport #
- passeport#
- passeportnon
- passeportn °
- passeport français
- passeport livre
- passeport carte
- numéro passeport
- passeport n°
- n° du passeport
- n° passeport

Keywords_eu_passport_date

- date of issue
- date of expiry

France social security number (INSEE) or equivalent identification

Format

15 digits

Pattern

Must match one of two patterns:

- 13 digits followed by a space followed by two digits
  or
- 15 consecutive digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_french_insee` finds content that matches the pattern.
- A keyword from Keyword_fr_insee is found.
- The checksum passes.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_french_insee or Func_fr_insee finds content that matches the pattern.
- The checksum passes.

```
<!-- France INSEE -->
<Entity id="71f62b97-efe0-4aa1-aa49-e14de253619d" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_french_insee" />
    <Any minMatches="0" maxMatches="0">
      <Match idRef="Keyword_fr_insee" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_french_insee" />
    <Match idRef="Keyword_fr_insee" />
  </Pattern>
</Entity>
```

Keywords

Keyword_fr_insee

- code sécu
- d'identité nationale
- insee
- fssn#
- le numéro d'identification nationale
- le code de la sécurité sociale
- national id
- national identification
- no d'identité
- no. d'identité
- numéro d'assurance
- numéro d'identité
- numero d'identite
- numéro de sécu
- numéro de sécurité sociale
- no d'identite
- no. d'identite
- ssn
- ssn#
- sécurité sociale
- securité sociale
- securite sociale
- socialsecuritynumber
- social security number
- social security code
- social insurance number

France tax identification number

Format

13 digits

Pattern

13 digits

- One digit that must be 0, 1, 2, or 3
- One digit
- A space (optional)
- Two digits
- A space (optional)
- Three digits
- A space (optional)
- Three digits
- A space (optional)
- Three check digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_france_eu_tax_file_number` finds content that matches the pattern.
- A keyword from `Keywords_france_eu_tax_file_number` is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_france_eu_tax_file_number` finds content that matches the pattern.

```
<!-- France Tax Identification Number (numéro SPI.) -->
<Entity id="ed59e77e-171d-442c-9ec1-88e2ebcb5b0a" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_france_eu_tax_file_number" />
    <Match idRef="Keywords_france_eu_tax_file_number" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_france_eu_tax_file_number" />
    <Any minMatches="0" maxMatches="0">
      <Match idRef="Keywords_france_eu_telephone_number" />
      <Match idRef="Keywords_france_eu_mobile_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_france_eu_tax_file_number

- numéro d'identification fiscale
- tax id
- tax identification no
- tax identification number
- tax no#
- tax no
- tax number
- tax registration number
- taxid#
- taxidno#

- taxidnumber#
- taxno#
- taxnumber#
- taxnumber
- tin id
- tin no
- tin#

France value added tax number

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

13 character alphanumeric pattern

Pattern

13 character alphanumeric pattern:

- two letters - FR (case insensitive)
- an optional space or hyphen
- two letters or digits
- an optional space, dot, hyphen, or comma
- three digits
- an optional space, dot, hyphen, or comma
- three digits
- an optional space, dot, hyphen, or comma
- three digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_france_value_added_tax_number finds content that matches the pattern.
- A keyword from Keywords_france_value_added_tax_number is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_france_value_added_tax_number finds content that matches the pattern.

```
<!-- France Value Added Tax Number -->
<Entity id="949121e6-ad9f-4379-8731-710342baea78" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_france_value_added_tax_number" />
    <Match idRef="Keywords_france_value_added_tax_number" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_france_value_added_tax_number" />
  </Pattern>
</Entity>
```

Keywords

Keyword_France_value_added_tax_number

- vat number
- vat no
- vat#
- value added tax
- siren identification no numéro d'identification taxe sur valeur ajoutée
- taxe valeur ajoutée
- taxe sur la valeur ajoutée
- n° tva
- numéro de tva
- numéro d'identification siren

Germany driver's license number

This sensitive information type entity is included in the EU Driver's License Number sensitive information type. It's available as a stand-alone sensitive information type entity.

Format

combination of 11 digits and letters

Pattern

11 digits and letters (not case-sensitive):

- a digit or letter
- two digits
- six digits or letters
- a digit
- a digit or letter

Checksum

Yes

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_german_drivers_license finds content that matches the pattern.
- A keyword from Keyword_german_drivers_license_number is found.
- The checksum passes.

```
<!-- German Driver's License Number -->
<Entity id="91da9335-1edb-45b7-a95f-5fe41a16c63c" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_german_drivers_license" />
    <Match idRef="Keyword_german_drivers_license" />
  </Pattern>
</Entity>
```

Keywords

Keyword_german_drivers_license_number

- ausstellungsdatum
- ausstellungsort
- ausstellende behöde
- ausstellende behorde
- ausstellende behoerde
- führerschein
- fuhrerschein
- fuehrerschein
- führerscheinnummer
- fuhrerscheinnummer
- fuehrerscheinnummer
- führerschein-
- fuhrerschein-
- fuehrerschein-
- führerscheinnummernr
- fuhrerscheinnummernr
- fuehrerscheinnummernr
- führerscheinnummerklasse
- fuhrerscheinnummerklasse
- fuehrerscheinnummerklasse
- nr-führerschein
- nr-fuhrerschein
- nr-fuehrerschein
- no-führerschein
- no-fuhrerschein
- no-fuehrerschein
- n-führerschein
- n-fuhrerschein
- n-fuehrerschein
- permis de conduire
- driverlic
- driverlics
- driverlicense
- driverlicenses
- driverlicence
- driverlicences
- driver lic
- driver lics

- driver license
- driver licenses
- driver licence
- driver licences
- driverslic
- driverslics
- driverslicence
- driverslicences
- driverslicense
- driverslicenses
- drivers lic
- drivers lics
- drivers license
- drivers licenses
- drivers licence
- drivers licences
- driver'lic
- driver'lics
- driver'license
- driver'licenses
- driver'licence
- driver'licences
- driver' lic
- driver' lics
- driver' license
- driver' licenses
- driver' licence
- driver' licences
- driver'slic
- driver'slics
- driver'slicense
- driver'slicenses
- driver'slicence
- driver'slicences
- driver's lic
- driver's lics
- driver's license
- driver's licenses
- driver's licence
- driver's licences
- dl#
- dls#
- driverlic#
- driverlics#
- driverlicense#
- driverlicenses#

- driverlicence#
- driverlicences#
- driver lic#
- driver lics#
- driver license#
- driver licenses#
- driver licences#
- driverslic#
- driverslics#
- driverslicense#
- driverslicenses#
- driverslicence#
- driverslicences#
- drivers lic#
- drivers lics#
- drivers license#
- drivers licenses#
- drivers licence#
- drivers licences#
- driver'lic#
- driver'lics#
- driver'license#
- driver'licenses#
- driver'licence#
- driver'licences#
- driver' lic#
- driver' lics#
- driver' license#
- driver' licenses#
- driver' licence#
- driver' licences#
- driver'slic#
- driver'slics#
- driver'slicense#
- driver'slicenses#
- driver'slicence#
- driver'slicences#
- driver's lic#
- driver's lics#
- driver's license#
- driver's licenses#
- driver's licence#
- driver's licences#
- driving licence
- driving license
- dlno#

- driv lic
- driv licen
- driv license
- driv licenses
- driv licence
- driv licences
- driver licen
- drivers licen
- driver's licen
- driving lic
- driving licen
- driving licenses
- driving licence
- driving licences
- driving permit
- dlno

Germany identity card number

Format

since 1 November 2010: Nine letters and digits

from 1 April 1987 until 31 October 2010: 10 digits

Pattern

since 1 November 2010:

- one letter (not case-sensitive)
- eight digits

from 1 April 1987 until 31 October 2010:

- 10 digits

Checksum

No

Definition

A DLP policy has low confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_germany_id_card finds content that matches the pattern.
- A keyword from Keyword_germany_id_card is found.

```
<!-- Germany Identity Card Number -->
<Entity id="e577372f-c42e-47a0-9d85-bebed1c237d4" recommendedConfidence="65" patternsProximity="300">
  <Pattern confidenceLevel="65">
     <IdMatch idRef="Regex_germany_id_card"/>
     <Match idRef="Keyword_germany_id_card"/>
  </Pattern>
</Entity>
```

Keywords

Keyword_germany_id_card

- ausweis
- gpid
- identification
- identifikation
- identifizierungsnummer
- identity card
- identity number
- id-nummer
- personal id
- personalausweis
- persönliche id nummer
- persönliche identifikationsnummer
- persönliche-id-nummer

Germany passport number

This sensitive information type entity is included in the EU Passport Number sensitive information type and is available as a stand-alone sensitive information type entity.

Format

10 digits or letters

Pattern

Pattern must include all of the following:

- first character is a digit or a letter from this set (C, F, G, H, J, K)
- three digits
- five digits or letters from this set (C, -H, J-N, P, R, T, V-Z)
- a digit

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_german_passport finds content that matches the pattern.
- A keyword from `Keyword_german_passport` or `Keywords_eu_passport_number_common` is found.
- The checksum passes.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_german_passport_data finds content that matches the pattern.
- A keyword from `Keyword_german_passport` or `Keywords_eu_passport_number_common` is found.
- The checksum passes.

```
<!-- German Passport Number -->
<Entity id="2e3da144-d42b-47ed-b123-fbf78604e52c" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_german_passport" />
    <Any minMatches="1">
      <Match idRef="Keyword_german_passport" />
      <Match idRef="Keywords_eu_passport_number_common" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_german_passport_data" />
    <Any minMatches="1">
      <Match idRef="Keyword_german_passport" />
      <Match idRef="Keywords_eu_passport_number_common" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keyword_german_passport

- reisepasse
- reisepassnummer
- No-Reisepass
- Nr-Reisepass
- Reisepass-Nr
- Passnummer
- reisepässe
- passeport no.
- passeport no

Keywords_eu_passport_number_common

- passport#
- passport #
- passportid
- passports
- passportno
- passport no
- passportnumber
- passport number
- passportnumbers
- passport numbers

Germany tax identification number

Format

11 digits without spaces and delimiters

Pattern

11 digits

- Two digits
- An optional space
- Three digits

- An optional space
- Three digits
- An optional space
- Two digits
- one check digit

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_germany_eu_tax_file_number` finds content that matches the pattern.
- A keyword from `Keywords_germany_eu_tax_file_number` is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_germany_eu_tax_file_number` finds content that matches the pattern.

```
<!-- Germany Tax Identification Number -->
<Entity id="43316a89-9880-40cf-b980-04bc7eefcec5" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_germany_eu_tax_file_number" />
    <Match idRef="Keywords_germany_eu_tax_file_number" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_germany_eu_tax_file_number" />
  </Pattern>
</Entity>
```

Keywords

Keywords_germany_eu_tax_file_number

- identifikationsnummer
- steuer id
- steueridentifikationsnummer
- steuernummer
- tax id
- tax identification no
- tax identification number
- tax no#
- tax no
- tax number
- tax registration number
- taxid#
- taxidno#
- taxidnumber#
- taxno#
- taxnumber#
- taxnumber
- tin id

- tin no
- tin#
- zinn#
- zinn
- zinnnummer

Germany value added tax number

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

11 character alphanumeric pattern

Pattern

11-character alphanumeric pattern:

- a letter D or d
- a letter E or e
- an optional space
- three digits
- an optional space or comma
- three digits
- an optional space or comma
- three digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_germany_value_added_tax_number finds content that matches the pattern.
- A keyword from Keywords_germany_value_added_tax_number is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_germany_value_added_tax_number finds content that matches the pattern.

```
<!-- Germany Value Added Tax Number -->
<Entity id="db177eb2-8811-4842-bffc-128c14aa219f" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_germany_value_added_tax_number" />
    <Match idRef="Keywords_germany_value_added_tax_number" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_germany_value_added_tax_number" />
  </Pattern>
</Entity>
```

Keywords

Keyword_germany_value_added_tax_number

- vat number
- vat no
- vat#
- vat# mehrwertsteuer
- mwst
- mehrwertsteuer identifikationsnummer
- mehrwertsteuer nummer

Greece driver's license number

This sensitive information type entity is included in the EU Driver's License Number sensitive information type and is available as a stand-alone sensitive information type entity.

Format

nine digits without spaces and delimiters

Pattern

nine digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_greece_eu_driver's_license_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_driver's_license_number` or `Keywords_greece_eu_driver's_license_number` is found.

```
<!-- Greece Driver's License Number -->
<Entity id="7a2200b5-aacf-4e3c-ab36-136d3e68b7da" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_greece_eu_driver's_license_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_driver's_license_number" />
      <Match idRef="Keywords_greece_eu_driver's_license_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_driver's_license_number

- driverlic
- driverlics
- driverlicense
- driverlicenses
- driverlicence
- driverlicences
- driver lic
- driver lics
- driver license
- driver licenses
- driver licence
- driver licences
- driverslic
- driverslics
- driverslicence
- driverslicences
- driverslicense
- driverslicenses
- drivers lic
- drivers lics
- drivers license
- drivers licenses
- drivers licence
- drivers licences
- driver'lic
- driver'lics
- driver'license
- driver'licenses
- driver'licence
- driver'licences
- driver' lic
- driver' lics
- driver' license
- driver' licenses
- driver' licence
- driver' licences
- driver'slic
- driver'slics
- driver'slicense
- driver'slicenses
- driver'slicence
- driver'slicences
- driver's lic
- driver's lics
- driver's license
- driver's licenses

- driver's licence
- driver's licences
- dl#
- dls#
- driverlic#
- driverlics#
- driverlicense#
- driverlicenses#
- driverlicence#
- driverlicences#
- driver lic#
- driver lics#
- driver license#
- driver licenses#
- driver licences#
- driverslic#
- driverslics#
- driverslicense#
- driverslicenses#
- driverslicence#
- driverslicences#
- drivers lic#
- drivers lics#
- drivers license#
- drivers licenses#
- drivers licence#
- drivers licences#
- driver'lic#
- driver'lics#
- driver'license#
- driver'licenses#
- driver'licence#
- driver'licences#
- driver' lic#
- driver' lics#
- driver' license#
- driver' licenses#
- driver' licence#
- driver' licences#
- driver'slic#
- driver'slics#
- driver'slicense#
- driver'slicenses#
- driver'slicence#
- driver'slicences#
- driver's lic#

- driver's lics#
- driver's license#
- driver's licenses#
- driver's licence#
- driver's licences#
- driving licence
- driving license
- dlno#
- driv lic
- driv licen
- driv license
- driv licenses
- driv licence
- driv licences
- driver licen
- drivers licen
- driver's licen
- driving lic
- driving licen
- driving licenses
- driving licence
- driving licences
- driving permit
- dl no
- dlno
- dl number

Keywords_greece_eu_driver's_license_number

- δεια οδήγησης
- Adeia odigisis
- Άδεια οδήγησης
- Δίπλωμα οδήγησης

Greece national ID card

Format

Combination of 7-8 letters and numbers plus a dash

Pattern

Seven letters and numbers (old format):

- One letter (any letter of the Greek alphabet)
- A dash
- Six digits

Eight letters and numbers (new format):

- Two letters whose uppercase character occurs in both the Greek and Latin alphabets (ABEZHIKMNOPTYX)
- A dash
- Six digits

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_greece_id_card finds content that matches the pattern.
- A keyword from Keyword_greece_id_card is found.

A DLP policy has low confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_greece_id_card finds content that matches the pattern.

```
<!-- Greece National ID Card -->
<Entity id="82568215-1da1-46d3-874a-d2294d81b5ac" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Regex_greece_id_card" />
    <Match idRef="Keyword_greece_id_card" />
  </Pattern>
  <Pattern confidenceLevel="65">
    <IdMatch idRef="Regex_greece_id_card" />
  </Pattern>
</Entity>
```

Keywords

Keyword_greece_id_card

- greek id
- greek national id
- greek personal id card
- greek police id
- identity card
- tautotita
- ταυτότητα
- ταυτότητας

Greece passport number

Format

Two letters followed by seven digits with no spaces or delimiters

Pattern

Two letters followed by seven digits

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_greece_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_greece_eu_passport_number` is found.

- The regular expression `Regex_greece_eu_passport_date` finds date in the format DD MMM YY (Example - 28 Aug 19) or a keyword from `Keywords_greece_eu_passport_date` is found

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_greece_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_greece_eu_passport_number` is found.

```
<!-- Greece Passport Number -->
<Entity id="7e65eb47-cdf9-4f52-8f90-2a27d5ee67e3" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Regex_greece_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_greece_eu_passport_number" />
    </Any>
    <Any minMatches="1">
      <Match idRef="Regex_greece_eu_passport_date" />
      <Match idRef="Keywords_greece_eu_passport_date" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_greece_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_greece_eu_passport_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_passport_number

- passport#
- passport #
- passportid
- passports
- passportno
- passport no
- passportnumber
- passport number
- passportnumbers
- passport numbers

Keywords_greece_eu_passport_number

- αριθμός διαβατηρίου
- αριθμούς διαβατηρίου
- αριθμός διαβατηριο

Greece Social Security Number (AMKA)

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance

- records management
- Microsoft cloud app security

Format

Eleven digits without spaces and delimiters

Pattern

- 6 digits as date of birth YYMMDD
- 4 digits
- a check digit

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_greece_eu_ssn` finds content that matches the pattern.
- A keyword from `Keywords_greece_eu_ssn_or_equivalent` is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_greece_eu_ssn` finds content that matches the pattern.

```
<!-- Greece Social Security Number (AMKA) -->
<Entity id="e39b03f4-50ea-41ae-af7a-a4b9539596ad" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_greece_eu_ssn" />
    <Match idRef="Keywords_greece_eu_ssn_or_equivalent" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_greece_eu_ssn" />
  </Pattern>
</Entity>
```

Keywords

Keywords_greece_eu_ssn_or_equivalent

- ssn
- ssn#
- social security no
- socialsecurityno#
- social security number
- amka
- a.m.k.a.
- Αριθμού Μητρώου Κοινωνικής Ασφάλισης

Greece tax identification number

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance

- records management
- Microsoft cloud app security

Format

Nine digits without spaces and delimiters

Pattern

Nine digits

Checksum

Not applicable

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_greece_eu_tax_file_number` finds content that matches the pattern.
- A keyword from `Keywords_greece_eu_tax_file_number` is found.

```
<!-- Greek Tax Identification Number -->
<Entity id="15a54a5a-53d4-4080-ad43-a2a4fe1d3bf7" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_greece_eu_tax_file_number" />
    <Match idRef="Keywords_greece_eu_tax_file_number" />
  </Pattern>
</Entity>
```

Keywords

Keywords_greece_eu_tax_file_number

- afm#
- afm
- αφμ|αφμ αριθμός
- αφμ
- tax id
- tax identification no
- tax identification number
- tax no#
- tax no
- tax number
- tax registration number
- tax registry no
- tax registry number
- taxid#
- taxidno#
- taxidnumber#
- taxno#
- taxnumber#
- taxnumber
- taxregistryno#
- tin id
- tin no

- tin#
- αριθμός φορολογικού μητρώου
- τον αριθμό φορολογικού μητρώου
- φορολογικού μητρώου νο

Hong Kong identity card (HKID) number

Format

Combination of 8-9 letters and numbers plus optional parentheses around the final character

Pattern

Combination of 8-9 letters:

- 1-2 letters (not case-sensitive)
- Six digits
- The final character (any digit or the letter A), which is the check digit and is optionally enclosed in parentheses.

Checksum

Yes

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_hong_kong_id_card finds content that matches the pattern.
- A keyword from Keyword_hong_kong_id_card is found.
- The checksum passes.

A DLP policy has low confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_hong_kong_id_card finds content that matches the pattern.
- The checksum passes.

```
<!-- Hong Kong Identity Card (HKID) number -->
<Entity id="e63c28a7-ad29-4c17-a41a-3d2a0b70fd9c" recommendedConfidence="75" patternsProximity="300">
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Func_hong_kong_id_card"/>
     <Match idRef="Keyword_hong_kong_id_card"/>
  </Pattern>
  <Pattern confidenceLevel="65">
     <IdMatch idRef="Func_hong_kong_id_card"/>
  </Pattern>
</Entity>
```

Keywords

Keyword_hong_kong_id_card

- hkid
- hong kong identity card
- HKIDC
- id card
- identity card
- hk identity card

- hong kong id
- 香港身份證
- 香港永久性居民身份證
- 身份證
- 身份証
- 身分證
- 身分証
- 香港身份証
- 香港身分證
- 香港身分証
- 香港身份證
- 香港居民身份證
- 香港居民身份証
- 香港居民身分證
- 香港居民身分証
- 香港永久性居民身份証
- 香港永久性居民身分證
- 香港永久性居民身分証
- 香港永久性居民身份證
- 香港非永久性居民身份證
- 香港非永久性居民身份証
- 香港非永久性居民身分證
- 香港非永久性居民身分証
- 香港特別行政區永久性居民身份證
- 香港特別行政區永久性居民身份証
- 香港特別行政區永久性居民身分證
- 香港特別行政區永久性居民身分証
- 香港特別行政區非永久性居民身份證
- 香港特別行政區非永久性居民身份証
- 香港特別行政區非永久性居民身分證
- 香港特別行政區非永久性居民身分証

Hungary driver's license number

Format

Two letters followed by six digits

Pattern

Two letters and six digits:

- Two letters (not case-sensitive)
- Six digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_hungary_eu_driver's_license_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_driver's_license_number` or `Keywords_hungary_eu_driver's_license_number` is found.

```
<Entity id="9d31c46b-6e6b-444c-aeb1-6dd7e604bb24" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_hungary_eu_driver's_license_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_driver's_license_number" />
      <Match idRef="Keywords_hungary_eu_driver's_license_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_driver's_license_number

- driverlic
- driverlics
- driverlicense
- driverlicenses
- driverlicence
- driverlicences
- driver lic
- driver lics
- driver license
- driver licenses
- driver licence
- driver licences
- driverslic
- driverslics
- driverslicence
- driverslicences
- driverslicense
- driverslicenses
- drivers lic
- drivers lics
- drivers license
- drivers licenses
- drivers licence
- drivers licences
- driver'lic
- driver'lics
- driver'license
- driver'licenses
- driver'licence
- driver'licences
- driver' lic
- driver' lics
- driver' license

- driver' licenses
- driver' licence
- driver' licences
- driver'slic
- driver'slics
- driver'slicense
- driver'slicenses
- driver'slicence
- driver'slicences
- driver's lic
- driver's lics
- driver's license
- driver's licenses
- driver's licence
- driver's licences
- dl#
- dls#
- driverlic#
- driverlics#
- driverlicense#
- driverlicenses#
- driverlicence#
- driverlicences#
- driver lic#
- driver lics#
- driver license#
- driver licenses#
- driver licences#
- driverslic#
- driverslics#
- driverslicense#
- driverslicenses#
- driverslicence#
- driverslicences#
- drivers lic#
- drivers lics#
- drivers license#
- drivers licenses#
- drivers licence#
- drivers licences#
- driver'lic#
- driver'lics#
- driver'license#
- driver'licenses#
- driver'licence#
- driver'licences#

- driver' lic#
- driver' lics#
- driver' license#
- driver' licenses#
- driver' licence#
- driver' licences#
- driver'slic#
- driver'slics#
- driver'slicense#
- driver'slicenses#
- driver'slicence#
- driver'slicences#
- driver's lic#
- driver's lics#
- driver's license#
- driver's licenses#
- driver's licence#
- driver's licences#
- driving licence
- driving license
- dlno#
- driv lic
- driv licen
- driv license
- driv licenses
- driv licence
- driv licences
- driver licen
- drivers licen
- driver's licen
- driving lic
- driving licen
- driving licenses
- driving licence
- driving licences
- driving permit
- dl no
- dlno
- dl number

Keywords_hungary_eu_driver's_license_number

- vezetoi engedely
- vezetői engedély
- vezetői engedélyek

Hungary personal identification number

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

11 digits

Pattern

11 digits:

- One digit that corresponds to gender (1-male, 2-female, other numbers are also possible for citizens born before 1900 or citizens with double citizenship)
- Six digits that correspond to birth date (YYMMDD)
- Three digits that correspond to a serial number
- One check digit

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_hungary_eu_national_id_card` finds content that matches the pattern.
- A keyword from `Keywords_hungary_eu_national_id_card` is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_hungary_eu_national_id_card` finds content that matches the pattern.

```
<!-- Hungary Personal Identification Number -->
<Entity id="7b5cc218-7046-47d9-80c9-f325b50896ca" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_hungary_eu_national_id_card" />
    <Match idRef="Keywords_hungary_eu_national_id_card" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_hungary_eu_national_id_card" />
    <Any minMatches="0" maxMatches="0">
      <Match idRef="Keywords_hungary_eu_telephone_number" />
      <Match idRef="Keywords_hungary_eu_mobile_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_hungary_eu_national_id_card

- id number
- identification number
- sz ig
- sz. ig.
- sz.ig.

- személyazonosító igazolvány
- személyi igazolvány

Hungary passport number

Format

Two letters followed by six or seven digits with no spaces or delimiters

Pattern

Two letters followed by six or seven digits

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_hungary_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_hungary_eu_passport_number` is found.
- The regular expression `Regex_hungary_eu_passport_date` finds date in the format DD MMM/MMM YY (Example - 01 MÁR/MAR 12) or a keyword from `Keywords_eu_passport_date` is found

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_hungary_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_hungary_eu_passport_number` is found.

```
<!-- Hungary Passport Number -->
<Entity id="5b483910-9aa7-4c99-9917-f4001464bda7" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Regex_hungary_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_hungary_eu_passport_number" />
    </Any>
    <Any minMatches="1">
      <Match idRef="Regex_hungary_eu_passport_date" />
      <Match idRef="Keywords_eu_passport_date" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_hungary_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_hungary_eu_passport_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_passport_number

- passport#
- passport #
- passportid
- passports

- passportno
- passport no
- passportnumber
- passport number
- passportnumbers
- passport numbers

Keywords_hungary_eu_passport_number

- útlevél száma
- Útlevelek száma
- útlevél szám

Keywords_eu_passport_date

- date of issue
- date of expiry

Hungary social security number (TAJ)

Format

Nine digits without spaces and delimiters

Pattern

Nine digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_hungary_eu_ssn_or_equivalent` finds content that matches the pattern.
- A keyword from `Keywords_hungary_eu_ssn_or_equivalent` is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_hungary_eu_ssn_or_equivalent` finds content that matches the pattern.

```
<!-- Hungarian Social Security Number (TAJ) -->
<Entity id="0de78315-9537-47f5-95ab-b3e77eba3993" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_hungary_eu_ssn_or_equivalent" />
    <Match idRef="Keywords_hungary_eu_ssn_or_equivalent" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_hungary_eu_ssn_or_equivalent" />
  </Pattern>
</Entity>
```

Keywords

Keywords_hungary_eu_ssn_or_equivalent

- hungarian social security number
- social security number
- socialsecuritynumber#

- hssn#
- socialsecuritynno
- hssn
- taj
- taj#
- ssn
- ssn#
- social security no
- áfa
- közösségi adószám
- általános forgalmi adó szám
- hozzáadottérték adó
- áfa szám
- magyar áfa szám

Hungary tax identification number

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

10 digits with no spaces or delimiters

Pattern

10 digits:

- One digit that must be "8"
- Eight digits
- One check digit

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_hungary_eu_tax_file_number` finds content that matches the pattern.
- A keyword from `Keywords_hungary_eu_tax_file_number` is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_hungary_eu_tax_file_number` finds content that matches the pattern.

```
<!-- Hungary Tax Identification Number -->
<Entity id="ede42eb4-59d9-49eb-9603-d7853fbda91d" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_hungary_eu_tax_file_number" />
    <Match idRef="Keywords_hungary_eu_tax_file_number" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_hungary_eu_tax_file_number" />
    <Any minMatches="0" maxMatches="0">
      <Match idRef="Keywords_hungary_eu_telephone_number" />
      <Match idRef="Keywords_hungary_eu_mobile_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_hungary_eu_tax_file_number

- adóazonosító szám
- adóhatóság szám
- adószám
- hungarian tin
- hungatiantin#
- tax authority no
- tax id
- tax identification no
- tax identification number
- tax no#
- tax no
- tax number
- tax registration number
- taxid#
- taxidno#
- taxidnumber#
- taxno#
- taxnumber#
- taxnumber
- tin id
- tin no
- tin#
- vat number

Hungary value added tax number

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

10 character alphanumeric pattern

Pattern

10 character alphanumeric pattern:

- two letters - HU or hu
- optional space
- eight digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_hungarian_value_added_tax_number finds content that matches the pattern.
- A keyword from Keywords_hungarian_value_added_tax_number is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_hungarian_value_added_tax_number finds content that matches the pattern.

```
<!-- Hungarian Value Added Tax Number -->
<Entity id="976349a0-683b-477a-90f8-ff0a220d5592" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_hungarian_value_added_tax_number" />
    <Match idRef="Keywords_hungarian_value_added_tax_number" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_hungarian_value_added_tax_number" />
  </Pattern>
</Entity>
```

Keywords

Keyword_Hungary_value_added_tax_number

- vat
- value added tax number
- vat#
- vatno#
- hungarianvatno#
- tax no.
- value added tax áfa
- közösségi adószám
- általános forgalmi adó szám
- hozzáadottérték adó
- áfa szám

India permanent account number (PAN)

Format

10 letters or digits

Pattern

10 letters or digits:

- Three letters (not case-sensitive)
- A letter in C, P, H, F, A, T, B, L, J, G (not case-sensitive)
- A letter
- Four digits
- A letter that is an alphabetic check digit

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_india_permanent_account_number finds content that matches the pattern.
- A keyword from Keyword_india_permanent_account_number is found.

A DLP policy has low confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_india_permanent_account_number finds content that matches the pattern.

```
<!-- India Permanent Account Number -->
<Entity id="2602bfee-9bb0-47a5-a7a6-2bf3053e2804" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Regex_india_permanent_account_number" />
    <Match idRef="Keyword_india_permanent_account_number" />
  </Pattern>
  <Version minEngineVersion="15.20.3520.000">
    <Pattern confidenceLevel="65">
      <IdMatch idRef="Regex_india_permanent_account_number" />
    </Pattern>
  </Version>
</Entity>
```

Keywords

Keyword_india_permanent_account_number

- Permanent Account Number
- PAN

India unique identification (Aadhaar) number

Format

12 digits containing optional spaces or dashes

Pattern

12 digits:

- A digit which is not 0 or 1
- Three digits
- An optional space or dash
- Four digits
- An optional space or dash
- The final digit, which is the check digit

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_india_aadhaar finds content that matches the pattern.
- A keyword from Keyword_india_aadhar is found.
- The checksum passes.
-

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_india_aadhaar finds content that matches the pattern.
- The checksum passes.

```
<!-- India Unique Identification (Aadhaar) number -->
<Entity id="1ca46b29-76f5-4f46-9383-cfa15e91048f" recommendedConfidence="85" patternsProximity="300">
  <Pattern confidenceLevel="85">
     <IdMatch idRef="Func_india_aadhaar"/>
     <Match idRef="Keyword_india_aadhar"/>
  </Pattern>
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Func_india_aadhaar"/>
  </Pattern>
</Entity>
```

Keywords

Keyword_india_aadhar

- aadhaar
- aadhar
- aadhar#
- uid
- ■■■■
- uidai

Indonesia identity card (KTP) number

Format

16 digits containing optional periods

Pattern

16 digits:

- Two-digit province code
- A period (optional)
- Two-digit regency or city code
- Two-digit subdistrict code
- A period (optional)
- Six digits in the format DDMMYY, which are the date of birth
- A period (optional)

- Four digits

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_indonesia_id_card finds content that matches the pattern.
- A keyword from Keyword_indonesia_id_card is found.

```
<!-- Indonesia Identity Card (KTP) Number -->
<Entity id="da68fdb0-f383-4981-8c86-82689d3b7d55" recommendedConfidence="85" patternsProximity="300">
  <Pattern confidenceLevel="85">
     <IdMatch idRef="Regex_indonesia_id_card"/>
     <Match idRef="Keyword_indonesia_id_card"/>
</Entity>
```

Keywords

Keyword_indonesia_id_card

- KTP
- Kartu Tanda Penduduk
- Nomor Induk Kependudukan

International banking account number (IBAN)

Format

Country code (two letters) plus check digits (two digits) plus bban number (up to 30 characters)

Pattern

Pattern must include all of the following:

- Two-letter country code
- Two check digits (followed by an optional space)
- 1-7 groups of four letters or digits (can be separated by spaces)
- 1-3 letters or digits

The format for each country is slightly different. The IBAN sensitive information type covers these 60 countries:

ad, ae, al, at, az, ba, be, bg, bh, ch, cr, cy, cz, de, dk, do, ee, es, fi, fo, fr, gb, ge, gi, gl, gr, hr, hu, ie, il, is, it, kw, kz, lb, li, lt, lu, lv, mc, md, me, mk, mr, mt, mu, nl, no, pl, pt, ro, rs, sa, se, si, sk, sm, tn, tr, vg

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_iban finds content that matches the pattern.
- The checksum passes.

```
<Entity id="e7dc4711-11b7-4cb0-b88b-2c394a771f0e" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
        <IdMatch idRef="Func_iban" />
  </Pattern>
</Entity>
```

Keywords

None

International classification of diseases (ICD-10-CM)

Format

Dictionary

Pattern

Keyword

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- A keyword from Dictionary_icd_10_updated is found.
- A keyword from Dictionary_icd_10_codes is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- A keyword from Dictionary_icd_10_ updated is found.

```
      <!-- ICD-10 CM -->
      <Entity id="3356946c-6bb7-449b-b253-6ffa419c0ce7" patternsProximity="300" recommendedConfidence="85">
        <Pattern confidenceLevel="85">
          <IdMatch idRef="Dictionary_icd_10_updated" />
          <Match idRef="Dictionary_icd_10_codes" />
        </Pattern>
        <Pattern confidenceLevel="75">
          <IdMatch idRef="Dictionary_icd_10_updated" />
        </Pattern>
```

Keywords

Any term from the Dictionary_icd_10_updated keyword dictionary, which is based on the International Classification of Diseases, Tenth Revision, Clinical Modification (ICD-10-CM). This type looks only for the term, not the insurance codes.

Any term from the Dictionary_icd_10_codes keyword dictionary, which is based on the International Classification of Diseases, Tenth Revision, Clinical Modification (ICD-10-CM). This type looks only for insurance codes, not the description.

International classification of diseases (ICD-9-CM)

Format

Dictionary

Pattern

Keyword

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- A keyword from Dictionary_icd_9_updated is found.
- A keyword from Dictionary_icd_9_codes is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- A keyword from Dictionary_icd_9_updated is found.

```
<Entity id="fa3f9c74-ee07-4c52-b5f2-085d6b2c0ec4" patternsProximity="300" recommendedConfidence="85">
    <Pattern confidenceLevel="85">
      <IdMatch idRef="Dictionary_icd_9_updated" />
      <Match idRef="Dictionary_icd_9_codes" />
    </Pattern>
    <Pattern confidenceLevel="75">
      <IdMatch idRef="Dictionary_icd_9_updated" />
    </Pattern>
  </Entity>
```

Keywords

Any term from the Dictionary_icd_9_updated keyword dictionary, which is based on the International Classification of Diseases,Ninth Revision, Clinical Modification (ICD-9-CM). This type looks only for the term, not the insurance codes.

Any term from the Dictionary_icd_9_codes keyword dictionary, which is based on the International Classification of Diseases,Ninth Revision, Clinical Modification (ICD-9-CM). This type looks only for insurance codes, not the description.

IP address

Format

IPv4:

Complex pattern that accounts for formatted (periods) and unformatted (no periods) versions of the IPv4 addresses

IPv6:

Complex pattern that accounts for formatted IPv6 numbers (which include colons)

Pattern

Checksum

No

Definition

For IPv6, a DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_ipv6_address finds content that matches the pattern.
- No keyword from Keyword_ipaddress is found.

For IPv4, a DLP policy is 95% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_ipv4_address finds content that matches the pattern.
- A keyword from Keyword_ipaddress is found.

For IPv6, a DLP policy is 95% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_ipv6_address finds content that matches the pattern.
- No keyword from Keyword_ipaddress is found.

```
<!-- IP Address -->
<Entity id="1daa4ad5-e2dd-4ca4-a788-54722c09efb2" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Regex_ipv6_address" />
    <Any minMatches="0" maxMatches="0">
      <Match idRef="Keyword_ipaddress" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="95">
    <IdMatch idRef="Regex_ipv4_address" />
    <Any minMatches="1">
      <Match idRef="Keyword_ipaddress" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="95">
    <IdMatch idRef="Regex_ipv6_address" />
    <Any minMatches="1">
      <Match idRef="Keyword_ipaddress" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keyword_ipaddress

- IP (this keyword is case-sensitive)
- ip address
- ip addresses
- internet protocol
- כתובת ה-IP

Ireland driver's license number

Format

Six digits followed by four letters

Pattern

Six digits and four letters:

- Six digits
- Four letters (not case-sensitive)

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_ireland_eu_driver's_license_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_driver's_license_number` or `Keywords_ireland_eu_driver's_license_number` is found.

```
<!-- Ireland Driver's License Number -->
<Entity id="e01bccd9-eb4d-414f-ace1-e9b6a4c4a2ca" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_ireland_eu_driver's_license_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_driver's_license_number" />
      <Match idRef="Keywords_ireland_eu_driver's_license_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_driver's_license_number

- driverlic
- driverlics
- driverlicense
- driverlicenses
- driverlicence
- driverlicences
- driver lic
- driver lics
- driver license
- driver licenses
- driver licence
- driver licences
- driverslic
- driverslics
- driverslicence
- driverslicences
- driverslicense
- driverslicenses
- drivers lic
- drivers lics
- drivers license
- drivers licenses
- drivers licence
- drivers licences
- driver'lic
- driver'lics
- driver'license
- driver'licenses
- driver'licence

- driver'licences
- driver' lic
- driver' lics
- driver' license
- driver' licenses
- driver' licence
- driver' licences
- driver'slic
- driver'slics
- driver'slicense
- driver'slicenses
- driver'slicence
- driver'slicences
- driver's lic
- driver's lics
- driver's license
- driver's licenses
- driver's licence
- driver's licences
- dl#
- dls#
- driverlic#
- driverlics#
- driverlicense#
- driverlicenses#
- driverlicence#
- driverlicences#
- driver lic#
- driver lics#
- driver license#
- driver licenses#
- driver licences#
- driverslic#
- driverslics#
- driverslicense#
- driverslicenses#
- driverslicence#
- driverslicences#
- drivers lic#
- drivers lics#
- drivers license#
- drivers licenses#
- drivers licence#
- drivers licences#
- driver'lic#
- driver'lics#

- driver'license#
- driver'licenses#
- driver'licence#
- driver'licences#
- driver' lic#
- driver' lics#
- driver' license#
- driver' licenses#
- driver' licence#
- driver' licences#
- driver'slic#
- driver'slics#
- driver'slicense#
- driver'slicenses#
- driver'slicence#
- driver'slicences#
- driver's lic#
- driver's lics#
- driver's license#
- driver's licenses#
- driver's licence#
- driver's licences#
- driving licence
- driving license
- dlno#
- driv lic
- driv licen
- driv license
- driv licenses
- driv licence
- driv licences
- driver licen
- drivers licen
- driver's licen
- driving lic
- driving licen
- driving licenses
- driving licence
- driving licences
- driving permit
- dl no
- dlno
- dl number

Keywords_ireland_eu_driver's_license_number

- ceadúnas tiomána
- ceadúnais tiomána

Ireland passport number

Format

Two letters or digits followed by seven digits with no spaces or delimiters

Pattern

Two letters or digits followed by seven digits:

- Two digits or letters (not case-sensitive)
- Seven digits

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_ireland_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_ireland_eu_passport_number` is found.
- The regular expression `Regex_ireland_eu_passport_date` finds date in the format DD MMM/MMM YYYY (Example - 01 BEA/MAY 1988) or a keyword from `Keywords_eu_passport_date` is found

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_ireland_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_ireland_eu_passport_number` is found.

```
<!-- Ireland Passport Number -->
<Entity id="a2130f27-9ee2-4103-84f9-a6b1ee7d0cbf" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Regex_ireland_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_ireland_eu_passport_number" />
    </Any>
    <Any minMatches="1">
      <Match idRef="Regex_ireland_eu_passport_date" />
      <Match idRef="Keywords_eu_passport_date" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_ireland_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_ireland_eu_passport_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_passport_number_common

- passport#
- passport #
- passportid
- passports

- passportno
- passport no
- passportnumber
- passport number
- passportnumbers
- passport numbers

Keywords_ireland_eu_passport_number

- passeport numero
- uimhreacha pasanna
- uimhir pas
- uimhir phas
- uimhreacha pas
- uimhir cárta
- uimhir chárta

Keywords_eu_passport_date

- date of issue
- date of expiry

Ireland personal public service (PPS) number

Format

Old format (until 31 December 2012):

- seven digits followed by 1-2 letters

New format (1 January 2013 and after):

- seven digits followed by two letters

Pattern

Old format (until 31 December 2012):

- seven digits
- one to two letters (not case-sensitive)

New format (1 January 2013 and after):

- seven digits
- a letter (not case-sensitive) which is an alphabetic check digit
- An optional letter in the range A-I, or "W"

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_ireland_pps finds content that matches the pattern.
- A keyword from Keywords_ireland_eu_national_id_card is found.
- The checksum passes.

A DLP policy has low confidence that it's detected this type of sensitive information if, within a proximity of 300

characters:

- The function Func_ireland_pps finds content that matches the pattern.
- The checksum passes.

```
      <!-- Ireland Personal Public Service (PPS) Number -->
      <Entity id="1cdb674d-c19a-4fcf-9f4b-7f56cc87345a" patternsProximity="300" recommendedConfidence="85"
relaxProximity="true">
        <Pattern confidenceLevel="85">
          <IdMatch idRef="Func_ireland_pps" />
          <Match idRef="Keywords_ireland_eu_national_id_card" />
        </Pattern>
        <Pattern confidenceLevel="65">
          <IdMatch idRef="Func_ireland_pps" />
        </Pattern>
      </Entity>
```

Keywords

Keywords_ireland_eu_national_id_card

- client identity service
- identification number
- personal id number
- personal public service number
- personal service no
- phearsanta seirbhíse poiblí
- pps no
- pps number
- pps num
- pps service no
- ppsn
- ppsno#
- ppsno
- psp
- public service no
- publicserviceno#
- publicserviceno
- revenue and social insurance number
- rsi no
- rsi number
- rsin
- seirbhís aitheantais cliant
- uimh
- uimhir aitheantais chánach
- uimhir aitheantais phearsanta
- uimhir phearsanta seirbhíse poiblí
- tax id
- tax identification no
- tax identification number
- tax no#
- tax no

- tax number
- tax registration number
- taxid#
- taxidno#
- taxidnumber#
- taxno#
- taxnumber#
- taxnumber
- tin id
- tin no
- tin#

Israel bank account number

Format

13 digits

Pattern

Formatted:

- two digits
- a dash
- three digits
- a dash
- eight digits

Unformatted:

- 13 consecutive digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_israel_bank_account_number finds content that matches the pattern.
- A keyword from Keyword_israel_bank_account_number is found.

```
<!-- Israel Bank Account Number -->
<Entity id="7d08b2ff-a0b9-437f-957c-aeddbf9b2b25" patternsProximity="300" recommendedConfidence="75">
    <Pattern confidenceLevel="75">
        <IdMatch idRef="Regex_israel_bank_account_number" />
        <Any minMatches="1">
          <Match idRef="Keyword_israel_bank_account_number" />
        </Any>
    </Pattern>
</Entity>
```

Keywords

Keyword_israel_bank_account_number

- Bank Account Number
- Bank Account

- Account Number
- מספר חשבון בנק

Israel national identification number

Format

nine digits

Pattern

nine consecutive digits

Checksum

Yes

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_israeli_national_id_number finds content that matches the pattern.
- A keyword from Keyword_Israel_National_ID is found.
- The checksum passes.

```
<!-- Israel National ID Number -->
<Entity id="e05881f5-1db1-418c-89aa-a3ac5c5277ee" patternsProximity="300" recommendedConfidence="75">
    <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_israeli_national_id_number" />
        <Any minMatches="1">
          <Match idRef="Keyword_Israel_National_ID" />
        </Any>
    </Pattern>
</Entity>
```

Keywords

Keyword_Israel_National_ID

- מספר זהות
- מספר זיה וי
- מספר זיהוי ישר אלי
- זהותישר אלית
- هو ية اسرائيل ية عدد
- هوية إسرائ يلية
- رقم الهوية
- عدد هوية فريدة من نوعها
- idnumber#
- id number
- identity no
- identitynumber#
- identity number
- israeliidentitynumber
- personal id
- unique id

Italy driver's license number

This sensitive information type entity is included in the EU Driver's License Number sensitive information type and is available as a stand-alone sensitive information type entity.

Format

a combination of 10 letters and digits

Pattern

a combination of 10 letters and digits:

- one letter (not case-sensitive)
- the letter "A" or "V" (not case-sensitive)
- seven digits
- one letter (not case-sensitive)

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_italy_drivers_license_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_driver's_license_number` or `Keyword_italy_drivers_license_number` is found.

```
<!-- Italy Driver's license Number -->
<Entity id="97d6244f-9157-41bd-8e0c-9d669a5c4d71" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_italy_drivers_license_number" />
      <Any minMatches="1">
        <Match idRef="Keywords_eu_driver's_license_number" />
        <Match idRef="Keyword_italy_drivers_license_number" />
      </Any>
    </Pattern>
</Entity>
```

Keywords

Keywords_eu_driver's_license_number

- driverlic
- driverlics
- driverlicense
- driverlicenses
- driverlicence
- driverlicences
- driver lic
- driver lics
- driver license
- driver licenses
- driver licence
- driver licences
- driverslic
- driverslics
- driverslicence
- driverslicences

- driverslicense
- driverslicenses
- drivers lic
- drivers lics
- drivers license
- drivers licenses
- drivers licence
- drivers licences
- driver'lic
- driver'lics
- driver'license
- driver'licenses
- driver'licence
- driver'licences
- driver' lic
- driver' lics
- driver' license
- driver' licenses
- driver' licence
- driver' licences
- driver'slic
- driver'slics
- driver'slicense
- driver'slicenses
- driver'slicence
- driver'slicences
- driver's lic
- driver's lics
- driver's license
- driver's licenses
- driver's licence
- driver's licences
- dl#
- dls#
- driverlic#
- driverlics#
- driverlicense#
- driverlicenses#
- driverlicence#
- driverlicences#
- driver lic#
- driver lics#
- driver license#
- driver licenses#
- driver licences#
- driverslic#

- driverslics#
- driverslicense#
- driverslicenses#
- driverslicence#
- driverslicences#
- drivers lic#
- drivers lics#
- drivers license#
- drivers licenses#
- drivers licence#
- drivers licences#
- driver'lic#
- driver'lics#
- driver'license#
- driver'licenses#
- driver'licence#
- driver'licences#
- driver' lic#
- driver' lics#
- driver' license#
- driver' licenses#
- driver' licence#
- driver' licences#
- driver'slic#
- driver'slics#
- driver'slicense#
- driver'slicenses#
- driver'slicence#
- driver'slicences#
- driver's lic#
- driver's lics#
- driver's license#
- driver's licenses#
- driver's licence#
- driver's licences#
- driving licence
- driving license
- dlno#
- driv lic
- driv licen
- driv license
- driv licenses
- driv licence
- driv licences
- driver licen
- drivers licen

- driver's licen
- driving lic
- driving licen
- driving licenses
- driving licence
- driving licences
- driving permit
- dl no
- dlno
- dl number

Keyword_italy_drivers_license_number

- numero di patente
- patente di guida
- patente guida
- patenti di guida
- patenti guida

Italy fiscal code

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

a 16-character combination of letters and digits in the specified pattern

Pattern

A 16-character combination of letters and digits:

- three letters that correspond to the first three consonants in the family name
- three letters that correspond to the first, third, and fourth consonants in the first name
- two digits that correspond to the last digits of the birth year
- one letter that corresponds to the letter for the month of birth—letters are used in alphabetical order, but only the letters A to E, H, L, M, P, R to T are used (so, January is A and October is R)
- two digits that correspond to the day of the month of birth—in order to differentiate between genders, 40 is added to the day of birth for women
- four digits that correspond to the area code specific to the municipality where the person was born (country-wide codes are used for foreign countries)
- one parity digit

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_italy_eu_national_id_card` finds content that matches the pattern.
- A keyword from `Keywords_italy_eu_national_id_card` is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_italy_eu_national_id_card` finds content that matches the pattern.

```
<!-- Italy Fiscal Code -->
<Entity id="4cd79172-8da9-4ff5-9188-98b1e7e2eca6" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_italy_eu_national_id_card" />
    <Match idRef="Keywords_italy_eu_national_id_card" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_italy_eu_national_id_card" />
  </Pattern>
</Entity>
```

Keywords

Keywords_italy_eu_national_id_card

- codice fiscal
- codice fiscale
- codice id personale
- codice personale
- fiscal code
- numero certificato personale
- numero di identificazione fiscale
- numero id personale
- numero personale
- personal certificate number
- personal code
- personal id code
- personal id number
- personalcodeno#
- tax code
- tax id
- tax identification no
- tax identification number
- tax identity number
- tax no#
- tax no
- tax number
- tax registration number
- taxid#
- taxidno#
- taxidnumber#
- taxno#
- taxnumber#
- taxnumber
- tin id

- tin no
- tin#

Italy passport number

Format

two letters or digits followed by seven digits with no spaces or delimiters

Pattern

two letters or digits followed by seven digits:

- two digits or letters (not case-sensitive)
- seven digits

Checksum

not applicable

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_italy_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_italy_eu_passport_number` is found.
- The regular expression `Regex_italy_eu_passport_date` finds date in the format DD MMM/MMM YYYY (Example - 01 GEN/JAN 1988) or a keyword from `Keywords_eu_passport_date` is found

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_italy_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_italy_eu_passport_number` is found.

```
<!-- Italy Passport Number -->
<Entity id="39811019-4750-445f-b26d-4c0e6c431544" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Regex_italy_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_italy_eu_passport_number" />
    </Any>
    <Any minMatches="1">
      <Match idRef="Regex_italy_eu_passport_date" />
      <Match idRef="Keywords_eu_passport_date" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_italy_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_italy_eu_passport_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_passport_number_common

- passport#

- passport #
- passportid
- passports
- passportno
- passport no
- passportnumber
- passport number
- passportnumbers
- passport numbers

Keywords_italy_eu_passport_number

- italiana passaporto
- passaporto italiana
- passaporto numero
- numéro passeport
- numero di passaporto
- numeri del passaporto
- passeport italien

Keywords_eu_passport_date

- date of issue
- date of expiry

Italy value added tax number

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

13 character alphanumeric pattern with optional delimiters

Pattern

13 character alphanumeric pattern with optional delimiters:

- I or i
- T or t
- optional space, dot, hyphen, or comma
- 11 digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_italy_value_added_tax_number finds content that matches the pattern.
- A keyword from Keywords_italy_value_added_tax_number is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_italy_value_added_tax_number finds content that matches the pattern.

```
<!-- Italy Value Added Tax -->
<Entity id="26a8cc07-2283-4a2a-ab1d-4ab643c4c67f" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_italy_value_added_tax_number" />
    <Match idRef="Keywords_italy_value_added_tax_number" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_italy_value_added_tax_number" />
  </Pattern>
</Entity>
```

Keywords

Keyword_italy_value_added_tax_number

- vat number
- vat no
- vat#
- iva
- iva#

Japan bank account number

Format

seven or eight digits

Pattern

bank account number:

- seven or eight digits
- bank account branch code:
- four digits
- a space or dash (optional)
- three digits

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_jp_bank_account finds content that matches the pattern.
- A keyword from Keyword_jp_bank_account is found.
- One of the following is true:
- The function Func_jp_bank_account_branch_code finds content that matches the pattern.
- A keyword from Keyword_jp_bank_branch_code is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_jp_bank_account finds content that matches the pattern.
- A keyword from Keyword_jp_bank_account is found.

```
<!-- Japan Bank Account Number -->
<Entity id="d354f95b-96ee-4b80-80bc-4377312b55bc" patternsProximity="300" recommendedConfidence="75">
  <Version minEngineVersion="15.01.0131.000">
    <Pattern confidenceLevel="85">
          <IdMatch idRef="Func_jp_bank_account" />
          <Match idRef="Keyword_jp_bank_account" />
          <Any minMatches="1">
            <Match idRef="Func_jp_bank_account_branch_code" />
            <Match idRef="Keyword_jp_bank_branch_code" />
          </Any>
      </Pattern>
  </Version>
     <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_jp_bank_account" />
        <Match idRef="Keyword_jp_bank_account" />
    </Pattern>
</Entity>
```

Keywords

Keyword_jp_bank_account

- Checking Account Number
- Checking Account
- Checking Account #
- Checking Acct Number
- Checking Acct #
- Checking Acct No.
- Checking Account No.
- Bank Account Number
- Bank Account
- Bank Account #
- Bank Acct Number
- Bank Acct #
- Bank Acct No.
- Bank Account No.
- Savings Account Number
- Savings Account
- Savings Account #
- Savings Acct Number
- Savings Acct #
- Savings Acct No.
- Savings Account No.
- Debit Account Number
- Debit Account
- Debit Account #
- Debit Acct Number
- Debit Acct #
- Debit Acct No.
- Debit Account No.
- 口座番号

- 銀行口座
- 銀行口座番号
- 総合口座
- 普通預金口座
- 普通口座
- 当座預金口座
- 当座口座
- 預金口座
- 振替口座
- 銀行
- バンク

Keyword_jp_bank_branch_code

- 支店番号
- 支店コード
- 店番号

Japan driver's license number

Format

12 digits

Pattern

12 consecutive digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_jp_drivers_license_number finds content that matches the pattern.
- A keyword from Keyword_jp_drivers_license_number is found.

```
<!-- Japan Driver's License Number -->
<Entity id="c6011143-d087-451c-8313-7f6d4aed2270" patternsProximity="300" recommendedConfidence="75">
    <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_jp_drivers_license_number" />
        <Match idRef ="Keyword_jp_drivers_license_number" />
    </Pattern>
</Entity>
```

Keywords

Keyword_jp_drivers_license_number

- driverlicense
- driverslicense
- driver'slicense
- driverslicenses
- driver'slicenses
- driverlicenses
- dl#

- dls#
- lic#
- lics#
- 運転免許証
- 運転免許
- 免許証
- 免許
- 運転免許証番号
- 運転免許番号
- 免許証番号
- 免許番号
- 運転免許証ナンバー
- 運転免許ナンバー
- 免許証ナンバー
- 運転免許証no
- 運転免許no
- 免許証no
- 免許no
- 運転経歴証明書番号
- 運転経歴証明書
- 運転免許証No.
- 運転免許No.
- 免許証No.
- 免許No.
- 運転免許証#
- 運転免許#
- 免許証#
- 免許#

Japan My Number - Corporate

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

13-digit number

Pattern

13-digit number:

- one digit from one to nine
- 12 digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_japanese_my_number_corporate finds content that matches the pattern.
- A keyword from Keywords_japanese_my_number_corporate is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_japanese_my_number_corporate finds content that matches the pattern.

```
<!-- Japanese My Number – Corporate -->
<Entity id="9e0eaf79-ff20-4ffb-b3e4-e7368d5db6ff" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_japanese_my_number_corporate" />
    <Match idRef="Keywords_japanese_my_number_corporate" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_japanese_my_number_corporate" />
  </Pattern>
</Entity>
```

Keywords

Keyword_japan_my_number_corporate

- corporate number
- マイナンバー
- 共通番号
- マイナンバーカード
- マイナンバーカード番号
- 個人番号カード
- 個人識別番号
- 個人識別ナンバー
- 法人番号
- 指定通知書

Japan My Number - Personal

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

12-digit number

Pattern

12-digit number:

- four digits
- an optional space, dot, or hyphen

- four digits
- an optional space, dot, or hyphen
- four digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_japanese_my_number_personal finds content that matches the pattern.
- A keyword from Keywords_japanese_my_number_personal is found.

A DLP policy has low confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_japanese_my_number_personal finds content that matches the pattern.

```
<!-- Japanese My Number – Personal -->
<Entity id="98da8e66-7299-4ebd-9f82-c871ab37d3ef" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_japanese_my_number_personal" />
    <Match idRef="Keywords_japanese_my_number_personal" />
  </Pattern>
  <Pattern confidenceLevel="65">
    <IdMatch idRef="Func_japanese_my_number_personal" />
  </Pattern>
</Entity>
```

Keywords

Keyword_japan_my_number_personal

- my number
- マイナンバー
- 個人番号
- 共通番号
- マイナンバーカード
- マイナンバーカード番号
- 個人番号カード
- 個人識別番号
- 個人識別ナンバー
- 通知カード

Japan passport number

Format

two letters followed by seven digits

Pattern

two letters (not case-sensitive) followed by seven digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_jp_passport finds content that matches the pattern.
- A keyword from Keyword_jp_passport is found.

```
<!-- Japan Passport Number -->
<Entity id="75177310-1a09-4613-bf6d-833aae3743f8" patternsProximity="300" recommendedConfidence="75">
    <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_jp_passport" />
        <Match idRef="Keyword_jp_passport" />
    </Pattern>
</Entity>
```

Keywords

Keyword_jp_passport

- Passport
- Passport Number
- Passport No.
- Passport #
- パスポート
- パスポート番号
- パスポートナンバー
- パスポート＃
- パスポート#
- パスポートNo.
- 旅券番号
- 旅券番号＃
- 旅券番号♯
- 旅券ナンバー

Japan residence card number

Format

12 letters and digits

Pattern

12 letters and digits:

- two letters (not case-sensitive)
- eight digits
- two letters (not case-sensitive)

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_jp_residence_card_number finds content that matches the pattern.
- A keyword from Keyword_jp_residence_card_number is found.

```
<!--Japan Residence Card Number-->
-<Entity id="ac36fef2-a289-4e2c-bb48-b02366e89fc0" recommendedConfidence="75" patternsProximity="300">
   -<Pattern confidenceLevel="75">
      <IdMatch idRef="Regex_jp_residence_card_number"/>
      <Match idRef="Keyword_jp_residence_card_number"/>
   </Pattern>
</Entity>
```

Keywords

Keyword_jp_residence_card_number

- Residence card number
- Residence card no
- Residence card #
- 在留カード番号
- 在留カード
- 在留番号

Japan resident registration number

Format

11 digits

Pattern

11 consecutive digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_jp_resident_registration_number finds content that matches the pattern.
- A keyword from Keyword_jp_resident_registration_number is found.

```
<!-- Japan Resident Registration Number -->
<Entity id="01c1209b-6389-4faf-a5f8-3f7e13899652" patternsProximity="300" recommendedConfidence="75">
    <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_jp_resident_registration_number" />
        <Match idRef ="Keyword_jp_resident_registration_number" />
    </Pattern>
</Entity>
```

Keywords

Keyword_jp_resident_registration_number

- Resident Registration Number
- Residents Basic Registry Number
- Resident Registration No.
- Resident Register No.
- Residents Basic Registry No.
- Basic Resident Register No.
- 外国人登録証明書番号
- 証明書番号

- 登録番号
- 外国人登録証

Japan social insurance number (SIN)

Format

7-12 digits

Pattern

7-12 digits:

- four digits
- a hyphen (optional)
- six digits OR
- 7-12 consecutive digits

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_jp_sin finds content that matches the pattern.
- A keyword from Keyword_jp_sin is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_jp_sin_pre_1997 finds content that matches the pattern.
- A keyword from Keyword_jp_sin is found.

```
<!-- Japan Social Insurance Number -->
<Entity id="c840e719-0896-45bb-84fd-1ed5c95e45ff" patternsProximity="300" recommendedConfidence="75">
    <Pattern confidenceLevel="85">
        <IdMatch idRef="Func_jp_sin" />
        <Match idRef="Keyword_jp_sin" />
    </Pattern>
    <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_jp_sin_pre_1997" />
        <Match idRef="Keyword_jp_sin" />
    </Pattern>
</Entity>
```

Keywords

Keyword_jp_sin

- Social Insurance No.
- Social Insurance Num
- Social Insurance Number
- 健康保険被保険者番号
- 健保番号
- 基礎年金番号
- 雇用保険被保険者番号
- 雇用保険番号

- 保険証番号
- 社会保険番号
- 社会保険No.
- 社会保険
- 介護保険
- 介護保険被保険者番号
- 健康保険被保険者整理番号
- 雇用保険被保険者整理番号
- 厚生年金
- 厚生年金被保険者整理番号

Latvia driver's license number

Format

three letters followed by six digits

Pattern

three letters and six digits:

- three letters (not case-sensitive)
- six digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_latvia_eu_driver's_license_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_driver's_license_number` or `Keywords_latvia_eu_driver's_license_number` is found.

```
<!-- Latvia Driver's License Number -->
<Entity id="ec996de0-30f2-46b1-b192-4d2ff8805fa7" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_latvia_eu_driver's_license_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_driver's_license_number" />
      <Match idRef="Keywords_latvia_eu_driver's_license_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_driver's_license_number

- driverlic
- driverlics
- driverlicense
- driverlicenses
- driverlicence
- driverlicences
- driver lic

- driver lics
- driver license
- driver licenses
- driver licence
- driver licences
- driverslic
- driverslics
- driverslicence
- driverslicences
- driverslicense
- driverslicenses
- drivers lic
- drivers lics
- drivers license
- drivers licenses
- drivers licence
- drivers licences
- driver'lic
- driver'lics
- driver'license
- driver'licenses
- driver'licence
- driver'licences
- driver' lic
- driver' lics
- driver' license
- driver' licenses
- driver' licence
- driver' licences
- driver'slic
- driver'slics
- driver'slicense
- driver'slicenses
- driver'slicence
- driver'slicences
- driver's lic
- driver's lics
- driver's license
- driver's licenses
- driver's licence
- driver's licences
- dl#
- dls#
- driverlic#
- driverlics#
- driverlicense#

- driverlicenses#
- driverlicence#
- driverlicences#
- driver lic#
- driver lics#
- driver license#
- driver licenses#
- driver licences#
- driverslic#
- driverslics#
- driverslicense#
- driverslicenses#
- driverslicence#
- driverslicences#
- drivers lic#
- drivers lics#
- drivers license#
- drivers licenses#
- drivers licence#
- drivers licences#
- driver'lic#
- driver'lics#
- driver'license#
- driver'licenses#
- driver'licence#
- driver'licences#
- driver' lic#
- driver' lics#
- driver' license#
- driver' licenses#
- driver' licence#
- driver' licences#
- driver'slic#
- driver'slics#
- driver'slicense#
- driver'slicenses#
- driver'slicence#
- driver'slicences#
- driver's lic#
- driver's lics#
- driver's license#
- driver's licenses#
- driver's licence#
- driver's licences#
- driving licence
- driving license

- dlno#
- driv lic
- driv licen
- driv license
- driv licenses
- driv licence
- driv licences
- driver licen
- drivers licen
- driver's licen
- driving lic
- driving licen
- driving licenses
- driving licence
- driving licences
- driving permit
- dl no
- dlno
- dl number

Keywords_latvia_eu_driver's_license_number

- autovadītāja apliecība
- autovadītāja apliecības
- vadītāja apliecība

Latvia personal code

Format

11 digits and an optional hyphen

Pattern

Old format

11 digits and a hyphen:

- six digits that correspond to the birth date (DDMMYY)
- a hyphen
- one digit that corresponds to the century of birth ("0" for 19th century, "1" for 20th century, and "2" for 21st century)
- four digits, randomly generated

New format

11 digits

- Two digits "32"
- Nine digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_latvia_eu_national_id_card` or the regex `Regex_latvia_eu_national_id_card_new_format` finds content that matches the pattern.
- A keyword from `Keywords_latvia_eu_national_id_card` is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_latvia_eu_national_id_card` or the regex `Regex_latvia_eu_national_id_card_new_format` finds content that matches the pattern.

```
<!-- Latvia Personal Code -->
<Entity id="03fcf763-27c2-49ed-9422-2641c6c895c9" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_latvia_eu_national_id_card" />
    <Match idRef="Keywords_latvia_eu_national_id_card" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_latvia_eu_national_id_card" />
    <Any minMatches="0" maxMatches="0">
      <Match idRef="Keywords_latvia_eu_telephone_number" />
      <Match idRef="Keywords_latvia_eu_mobile_number" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Regex_latvia_eu_national_id_card_new_format" />
    <Match idRef="Keywords_latvia_eu_national_id_card" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_latvia_eu_national_id_card_new_format" />
    <Any minMatches="0" maxMatches="0">
      <Match idRef="Keywords_latvia_eu_telephone_number" />
      <Match idRef="Keywords_latvia_eu_mobile_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_latvia_eu_national_id_card

- administrative number
- alvas nē
- birth number
- citizen number
- civil number
- electronic census number
- electronic number
- fiscal code
- healthcare user number
- id#
- id-code
- identification number
- identifikācijas numurs
- id-number
- individual number

- latvija alva
- nacionālais id
- national id
- national identifying number
- national identity number
- national insurance number
- national register number
- nodokļa numurs
- nodokļu id
- nodokļu identifikācija numurs
- personal certificate number
- personal code
- personal id code
- personal id number
- personal identification code
- personal identifier
- personal identity number
- personal number
- personal numeric code
- personalcodeno#
- personas kods
- population identification code
- public service number
- registration number
- revenue number
- social insurance number
- social security number
- state tax code
- tax file number
- tax id
- tax identification no
- tax identification number
- tax no#
- tax no
- tax number
- taxid#
- taxidno#
- taxidnumber#
- taxno#
- taxnumber#
- taxnumber
- tin id
- tin no
- tin#
- voter’s number

Latvia passport number

Format

two letters or digits followed by seven digits with no spaces or delimiters

Pattern

two letters or digits followed by seven digits:

- two digits or letters (not case-sensitive)
- seven digits

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_latvia_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_latvia_eu_passport_number` is found.
- The regular expression `Regex_eu_passport_date1` finds date in the format DD.MM.YYYY or a keyword from `Keywords_eu_passport_date` is found

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_latvia_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_latvia_eu_passport_number` is found.

```
<!-- Latvia Passport Number -->
<Entity id="23ae25ec-cc28-421b-b77a-3054eadf1ede" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Regex_latvia_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_latvia_eu_passport_number" />
    </Any>
    <Any minMatches="1">
      <Match idRef="Regex_eu_passport_date1" />
      <Match idRef="Keywords_eu_passport_date" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_latvia_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_latvia_eu_passport_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_passport_number_common

- passport#
- passport #
- passportid
- passports

- passportno
- passport no
- passportnumber
- passport number
- passportnumbers
- passport numbers

Keywords_latvia_eu_passport_number

- pase numurs
- pase numur
- pases numuri
- pases nr
- passeport no
- n° du Passeport

Keywords_eu_passport_date

- date of issue
- date of expiry

Lithuania driver's license number

Format

eight digits without spaces and delimiters

Pattern

eight digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_lithuania_eu_driver's_license_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_driver's_license_number` or `Keywords_lithuania_eu_driver's_license_number` is found.

```
<!-- Lithuania Driver's License Number -->
<Entity id="86f7628b-e0f4-4dc3-9fbc-e4300e4c7d78" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_lithuania_eu_driver's_license_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_driver's_license_number" />
      <Match idRef="Keywords_lithuania_eu_driver's_license_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_driver's_license_number

- driverlic
- driverlics
- driverlicense

- driverlicenses
- driverlicence
- driverlicences
- driver lic
- driver lics
- driver license
- driver licenses
- driver licence
- driver licences
- driverslic
- driverslics
- driverslicence
- driverslicences
- driverslicense
- driverslicenses
- drivers lic
- drivers lics
- drivers license
- drivers licenses
- drivers licence
- drivers licences
- driver'lic
- driver'lics
- driver'license
- driver'licenses
- driver'licence
- driver'licences
- driver' lic
- driver' lics
- driver' license
- driver' licenses
- driver' licence
- driver' licences
- driver'slic
- driver'slics
- driver'slicense
- driver'slicenses
- driver'slicence
- driver'slicences
- driver's lic
- driver's lics
- driver's license
- driver's licenses
- driver's licence
- driver's licences
- dl#

- dls#
- driverlic#
- driverlics#
- driverlicense#
- driverlicenses#
- driverlicence#
- driverlicences#
- driver lic#
- driver lics#
- driver license#
- driver licenses#
- driver licences#
- driverslic#
- driverslics#
- driverslicense#
- driverslicenses#
- driverslicence#
- driverslicences#
- drivers lic#
- drivers lics#
- drivers license#
- drivers licenses#
- drivers licence#
- drivers licences#
- driver'lic#
- driver'lics#
- driver'license#
- driver'licenses#
- driver'licence#
- driver'licences#
- driver' lic#
- driver' lics#
- driver' license#
- driver' licenses#
- driver' licence#
- driver' licences#
- driver'slic#
- driver'slics#
- driver'slicense#
- driver'slicenses#
- driver'slicence#
- driver'slicences#
- driver's lic#
- driver's lics#
- driver's license#
- driver's licenses#

- driver's licence#
- driver's licences#
- driving licence
- driving license
- dlno#
- driv lic
- driv licen
- driv license
- driv licenses
- driv licence
- driv licences
- driver licen
- drivers licen
- driver's licen
- driving lic
- driving licen
- driving licenses
- driving licence
- driving licences
- driving permit
- dl no
- dlno
- dl number

Keywords_lithuania_eu_driver's_license_number

- vairuotojo pažymėjimas
- vairuotojo pažymėjimo numeris
- vairuotojo pažymėjimo numeriai

Lithuania personal code

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

11 digits without spaces and delimiters

Pattern

11 digits without spaces and delimiters:

- one digit (1-6) that corresponds to the person's gender and century of birth
- six digits that correspond to birth date (YYMMDD)
- three digits that correspond to the serial number of the date of birth
- one check digit

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_lithuania_eu_tax_file_number` finds content that matches the pattern.
- A keyword from `Keywords_lithuania_eu_tax_file_number` is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_lithuania_eu_tax_file_number` finds content that matches the pattern.

```
<!-- Lithuania Personal Code -->
<Entity id="cd6d3786-8ec3-4524-a2cf-1e0095379171" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_lithuania_eu_tax_file_number" />
    <Match idRef="Keywords_lithuania_eu_tax_file_number" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_lithuania_eu_tax_file_number" />
    <Any minMatches="0" maxMatches="0">
      <Match idRef="Keywords_lithuania_eu_telephone_number" />
      <Match idRef="Keywords_lithuania_eu_mobile_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_lithuania_eu_national_id_card

- asmeninis skaitmeninis kodas
- asmens kodas
- citizen service number
- mokesčių id
- mokesčių identifikavimas numeris
- mokesčių identifikavimo numeris
- mokesčių numeris
- national identification number
- personal code
- personal numeric code
- piliečio paslaugos numeris
- tax id
- tax identification no
- tax identification number
- tax no#
- tax no
- tax number
- tax registration number
- taxid#
- taxidno#
- taxidnumber#
- taxno#

- taxnumber#
- taxnumber
- tin id
- tin no
- tin#
- unikalus identifikavimo kodas
- unikalus identifikavimo numeris
- unique identification number
- unique identity number
- uniqueidentityno#

Lithuania passport number

Format

eight digits or letters with no spaces or delimiters

Pattern

eight digits or letters (not case-sensitive)

Checksum

not applicable

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_lithuania_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_lithuania_eu_passport_number` is found.
- The regular expression `Regex_eu_passport_date3` finds date in the format DD MM YYYY or a keyword from `Keywords_eu_passport_date` is found

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_lithuania_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_lithuania_eu_passport_number` is found.

```
<!-- Lithuania Passport Number -->
<Entity id="1b79900f-047b-4c3f-846f-7d73b5534bce" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Regex_lithuania_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_lithuania_eu_passport_number" />
    </Any>
    <Any minMatches="1">
      <Match idRef="Regex_eu_passport_date3" />
      <Match idRef="Keywords_eu_passport_date" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_lithuania_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_lithuania_eu_passport_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_passport_number

- passport#
- passport #
- passportid
- passports
- passportno
- passport no
- passportnumber
- passport number
- passportnumbers
- passport numbers

Keywords_lithuania_eu_passport_number

- paso numeris
- paso numeriai
- paso nr

Keywords_eu_passport_date

- date of issue
- date of expiry

Luxemburg driver's license number

Format

six digits without spaces and delimiters

Pattern

six digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_luxemburg_eu_driver's_license_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_driver's_license_number` or `Keywords_luxemburg_eu_driver's_license_number` is found.

```
<!-- Luxemburg Driver's License Number -->
<Entity id="89daf717-1544-4860-9a2e-fc9166dd8852" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_luxemburg_eu_driver's_license_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_driver's_license_number" />
      <Match idRef="Keywords_luxemburg_eu_driver's_license_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_driver's_license_number

- driverlic
- driverlics
- driverlicense
- driverlicenses
- driverlicence
- driverlicences
- driver lic
- driver lics
- driver license
- driver licenses
- driver licence
- driver licences
- driverslic
- driverslics
- driverslicence
- driverslicences
- driverslicense
- driverslicenses
- drivers lic
- drivers lics
- drivers license
- drivers licenses
- drivers licence
- drivers licences
- driver'lic
- driver'lics
- driver'license
- driver'licenses
- driver'licence
- driver'licences

- driver' lic
- driver' lics
- driver' license
- driver' licenses
- driver' licence
- driver' licences
- driver'slic
- driver'slics
- driver'slicense
- driver'slicenses
- driver'slicence
- driver'slicences
- driver's lic
- driver's lics
- driver's license
- driver's licenses
- driver's licence
- driver's licences
- dl#
- dls#
- driverlic#
- driverlics#
- driverlicense#
- driverlicenses#
- driverlicence#
- driverlicences#
- driver lic#
- driver lics#
- driver license#
- driver licenses#
- driver licences#
- driverslic#
- driverslics#
- driverslicense#
- driverslicenses#
- driverslicence#
- driverslicences#
- drivers lic#
- drivers lics#
- drivers license#
- drivers licenses#
- drivers licence#
- drivers licences#
- driver'lic#
- driver'lics#
- driver'license#

- driver'licenses#
- driver'licence#
- driver'licences#
- driver' lic#
- driver' lics#
- driver' license#
- driver' licenses#
- driver' licence#
- driver' licences#
- driver'slic#
- driver'slics#
- driver'slicense#
- driver'slicenses#
- driver'slicence#
- driver'slicences#
- driver's lic#
- driver's lics#
- driver's license#
- driver's licenses#
- driver's licence#
- driver's licences#
- driving licence
- driving license
- dlno#
- driv lic
- driv licen
- driv license
- driv licenses
- driv licence
- driv licences
- driver licen
- drivers licen
- driver's licen
- driving lic
- driving licen
- driving licenses
- driving licence
- driving licences
- driving permit
- dl no
- dlno
- dl number

Keywords_luxemburg_eu_driver's_license_number

- fahrerlaubnis
- Führerschäin

Luxemburg national identification number (natural persons)

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

13 digits with no spaces or delimiters

Pattern

13 digits:

- 11 digits
- two check digits

Checksum

yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_luxemburg_eu_tax_file_number` finds content that matches the pattern.
- A keyword from `Keywords_luxemburg_eu_national_id_card` is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_luxemburg_eu_tax_file_number` finds content that matches the pattern.

```
<!-- Luxemburg National Identification Number (Natural persons) -->
<Entity id="aaf661ed-29ec-426d-8bf9-880cad298ebb" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_luxemburg_eu_tax_file_number" />
    <Match idRef="Keywords_luxemburg_eu_national_id_card" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_luxemburg_eu_tax_file_number" />
    <Any minMatches="0" maxMatches="0">
      <Match idRef="Keywords_luxemburg_eu_telephone_number" />
      <Match idRef="Keywords_luxemburg_eu_mobile_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_luxemburg_eu_national_id_card

- eindeutige id
- eindeutige id-nummer
- eindeutigeid#
- id personnelle
- idpersonnelle#

- idpersonnelle
- individual code
- individual id
- individual identification
- individual identity
- numéro d'identification personnel
- personal id
- personal identification
- personal identity
- personalidno#
- personalidnumber#
- persönliche identifikationsnummer
- unique id
- unique identity
- uniqueidkey#

Luxemburg passport number

Format

eight digits or letters with no spaces or delimiters

Pattern

eight digits or letters (not case-sensitive)

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_luxemburg_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_luxemburg_eu_passport_number` is found.
- The regular expression `Regex_eu_passport_date3` finds date in the format DD MM YYYY or a keyword from `Keywords_eu_passport_date` is found

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_luxemburg_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_luxemburg_eu_passport_number` is found.

```
<!-- Luxemburg Passport Number -->
<Entity id="81d5c027-bed9-4421-91a0-3b2e55b3eb85" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Regex_luxemburg_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_luxemburg_eu_passport_number" />
    </Any>
    <Any minMatches="1">
      <Match idRef="Regex_eu_passport_date3" />
      <Match idRef="Keywords_eu_passport_date" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_luxemburg_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_luxemburg_eu_passport_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_passport_number

- passport#
- passport #
- passportid
- passports
- passportno
- passport no
- passportnumber
- passport number
- passportnumbers
- passport numbers

Keywords_luxemburg_eu_passport_number

- ausweisnummer
- luxembourg pass
- luxembourg passeport
- luxembourg passport
- no de passeport
- no-reisepass
- nr-reisepass
- numéro de passeport
- pass net
- pass nr
- passnummer
- passeport nombre
- reisepässe
- reisepass-nr
- reisepassnummer

Keywords_eu_passport_date

- date of issue

- date of expiry

Luxemburg national identification number (non-natural persons)

Format

11 digits

Pattern

11 digits

- two digits
- an optional space
- three digits
- an optional space
- three digits
- an optional space
- two digits
- one check digit

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_luxemburg_eu_tax_file_number_non_natural` finds content that matches the pattern.
- A keyword from `Keywords_luxemburg_eu_tax_file_number` is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_luxemburg_eu_tax_file_number_non_natural` finds content that matches the pattern.

```
<!-- Luxemburg National Identification Number (Non-natural persons) -->
<Entity id="84bffa3a-d805-4788-a613-b1e4df3804cf" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_luxemburg_eu_tax_file_number_non_natural" />
    <Match idRef="Keywords_luxemburg_eu_tax_file_number" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_luxemburg_eu_tax_file_number_non_natural" />
    <Any minMatches="0" maxMatches="0">
      <Match idRef="Keywords_luxemburg_eu_telephone_number" />
      <Match idRef="Keywords_luxemburg_eu_mobile_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_luxemburg_eu_tax_file_number

- carte de sécurité sociale
- étain non
- étain#
- identifiant d'impôt

- luxembourg tax identifikatiounsnummer
- numéro d'étain
- numéro d'identification fiscal luxembourgeois
- numéro d'identification fiscale
- social security
- sozialunterstützung
- sozialversécherung
- sozialversicherungsausweis
- steier id
- steier identifikatiounsnummer
- steier nummer
- steuer id
- steueridentifikationsnummer
- steuernummer
- tax id
- tax identification no
- tax identification number
- tax no#
- tax no
- tax number
- tax registration number
- taxid#
- taxidno#
- taxidnumber#
- taxno#
- taxnumber#
- taxnumber
- tin id
- tin no
- tin#
- zinn#
- zinn
- zinnzahl

Malaysia identification card number

Format

12 digits containing optional hyphens

Pattern

12 digits:

- six digits in the format YYMMDD, which are the date of birth
- a dash (optional)
- two-letter place-of-birth code
- a dash (optional)
- three random digits
- one-digit gender code

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_malaysia_id_card_number finds content that matches the pattern.
- A keyword from Keyword_malaysia_id_card_number is found.

```
<!-- Malaysia ID Card Number -->
</Entity>
      <Entity id="7f0e921c-9677-435b-aba2-bb8f1013c749" patternsProximity="300" recommendedConfidence="85">
        <Pattern confidenceLevel="85">
            <IdMatch idRef="Regex_malaysia_id_card_number" />
            <Match idRef="Keyword_malaysia_id_card_number" />
        </Pattern>
</Entity>
```

Keywords

Keyword_malaysia_id_card_number

- digital application card
- i/c
- i/c no
- ic
- ic no
- id card
- identification Card
- identity card
- k/p
- k/p no
- kad akuan diri
- kad aplikasi digital
- kad pengenalan malaysia
- kp
- kp no
- mykad
- mykas
- mykid
- mypr
- mytentera
- malaysia identity card
- malaysian identity card
- nric
- personal identification card

Malta driver's license number

Format

Combination of two characters and six digits in the specified pattern

Pattern

combination of two characters and six digits:

- two characters (digits or letters, not case-sensitive)
- a space (optional)
- three digits
- a space (optional)
- three digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_malta_eu_driver's_license_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_driver's_license_number` or `Keywords_malta_eu_driver's_license_number` is found.

```
<!-- Malta Driver's License Number -->
<Entity id="a3bdaa4a-8371-4735-8fa5-56ee0fb4afc4" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_malta_eu_driver's_license_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_driver's_license_number" />
      <Match idRef="Keywords_malta_eu_driver's_license_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_driver's_license_number

- driverlic
- driverlics
- driverlicense
- driverlicenses
- driverlicence
- driverlicences
- driver lic
- driver lics
- driver license
- driver licenses
- driver licence
- driver licences
- driverslic
- driverslics
- driverslicence
- driverslicences
- driverslicense
- driverslicenses
- drivers lic

- drivers lics
- drivers license
- drivers licenses
- drivers licence
- drivers licences
- driver'lic
- driver'lics
- driver'license
- driver'licenses
- driver'licence
- driver'licences
- driver' lic
- driver' lics
- driver' license
- driver' licenses
- driver' licence
- driver' licences
- driver'slic
- driver'slics
- driver'slicense
- driver'slicenses
- driver'slicence
- driver'slicences
- driver's lic
- driver's lics
- driver's license
- driver's licenses
- driver's licence
- driver's licences
- dl#
- dls#
- driverlic#
- driverlics#
- driverlicense#
- driverlicenses#
- driverlicence#
- driverlicences#
- driver lic#
- driver lics#
- driver license#
- driver licenses#
- driver licences#
- driverslic#
- driverslics#
- driverslicense#
- driverslicenses#

- driverslicence#
- driverslicences#
- drivers lic#
- drivers lics#
- drivers license#
- drivers licenses#
- drivers licence#
- drivers licences#
- driver'lic#
- driver'lics#
- driver'license#
- driver'licenses#
- driver'licence#
- driver'licences#
- driver' lic#
- driver' lics#
- driver' license#
- driver' licenses#
- driver' licence#
- driver' licences#
- driver'slic#
- driver'slics#
- driver'slicense#
- driver'slicenses#
- driver'slicence#
- driver'slicences#
- driver's lic#
- driver's lics#
- driver's license#
- driver's licenses#
- driver's licence#
- driver's licences#
- driving licence
- driving license
- dlno#
- driv lic
- driv licen
- driv license
- driv licenses
- driv licence
- driv licences
- driver licen
- drivers licen
- driver's licen
- driving lic
- driving licen

- driving licenses
- driving licence
- driving licences
- driving permit
- dl no
- dlno
- dl number

Keywords_malta_eu_driver's_license_number

- liċenzja tas-sewqan
- liċenzji tas-sewwieq

Malta identity card number

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

seven digits followed by one letter

Pattern

seven digits followed by one letter:

- seven digits
- one letter in "M, G, A, P, L, H, B, Z" (case insensitive)

Checksum

Not applicable

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_malta_eu_national_id_card` finds content that matches the pattern.
- A keyword from `Keywords_malta_eu_national_id_card` is found.

A DLP policy has low confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_malta_eu_national_id_card` finds content that matches the pattern.

```
<!-- Malta Identity Card Number -->
<Entity id="854b36b3-a388-4ac8-a4ec-677c2b5e4356" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_malta_eu_national_id_card" />
    <Match idRef="Keywords_malta_eu_national_id_card" />
  </Pattern>
  <Pattern confidenceLevel="65">
    <IdMatch idRef="Regex_malta_eu_national_id_card" />
  </Pattern>
</Entity>
```

Keywords

Keywords_malta_eu_national_id_card

- citizen service number
- id tat-taxxa
- identifika numru tal-biljett
- kodiċi numerali personali
- numru ta 'identifikazzjoni personali
- numru ta 'identifikazzjoni tat-taxxa
- numru ta 'identifikazzjoni uniku
- numru ta' identità uniku
- numru tas-servizz taċ-ċittadin
- numru tat-taxxa
- personal numeric code
- unique identification number
- unique identity number
- uniqueidentityno#

Malta passport number

Format

seven digits without spaces or delimiters

Pattern

seven digits

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_malta_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_malta_eu_passport_number` is found.
- A keyword from `Keywords_eu_passport_date` is found

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_malta_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_malta_eu_passport_number` is found.

```
<!-- Malta Passport Number -->
<Entity id="b2b21198-48f9-4d13-b2a5-03969bff0fb8" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Regex_malta_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_malta_eu_passport_number" />
    </Any>
    <Match idRef="Keywords_eu_passport_date" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_malta_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_malta_eu_passport_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_passport_number

- passport#
- passport #
- passportid
- passports
- passportno
- passport no
- passportnumber
- passport number
- passportnumbers
- passport numbers

Keywords_malta_eu_passport_number

- numru tal-passaport
- numri tal-passaport
- Nru tal-passaport

Keywords_eu_passport_date

- date of issue
- date of expiry

Malta tax identification number

Format

For Maltese nationals:

- seven digits and one letter in the specified pattern

Non-Maltese nationals and Maltese entities:

- nine digits

Pattern

Maltese nationals: seven digits and one letter

- seven digits

- one letter (not case-sensitive)

Non-Maltese nationals and Maltese entities: nine digits

- nine digits

Checksum

Not applicable

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regex `Regex_malta_eu_tax_file_number` or `Regex_malta_eu_tax_file_number_non_maltese_national` finds content that matches the pattern.
- A keyword from `Keywords_malta_eu_tax_file_number` is found.

A DLP policy has low confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regex `Regex_malta_eu_tax_file_number` or `Regex_malta_eu_tax_file_number_non_maltese_national` finds content that matches the pattern.

```
<!-- Malta Tax ID Number -->
<Entity id="ec830c63-65f4-45d0-9d8c-910dc8334b20" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_malta_eu_tax_file_number" />
    <Match idRef="Keywords_malta_eu_tax_file_number" />
  </Pattern>
  <Pattern confidenceLevel="65">
    <IdMatch idRef="Regex_malta_eu_tax_file_number" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_malta_eu_tax_file_number_non_maltese_national" />
    <Match idRef="Keywords_malta_eu_tax_file_number" />
  </Pattern>
  <Pattern confidenceLevel="65">
    <IdMatch idRef="Regex_malta_eu_tax_file_number_non_maltese_national" />
  </Pattern>
</Entity>
```

Keywords

Keywords_malta_eu_tax_file_number

- citizen service number
- id tat-taxxa
- identifika numru tal-biljett
- kodiċi numerali personali
- numru ta 'identifikazzjoni personali
- numru ta 'identifikazzjoni tat-taxxa
- numru ta 'identifikazzjoni uniku
- numru ta' identità uniku
- numru tas-servizz taċ-ċittadin
- numru tat-taxxa
- personal numeric code
- tax id
- tax identification no

- tax identification number
- tax no#
- tax no
- tax number
- tax registration number
- taxid#
- taxidno#
- taxidnumber#
- taxno#
- taxnumber#
- taxnumber
- tin id
- tin no
- tin#
- unique identification number
- unique identity number
- uniqueidentityno#

Netherlands citizen's service (BSN) number

Format

eight or nine digits containing optional spaces

Pattern

eight-nine digits:

- three digits
- a space (optional)
- three digits
- a space (optional)
- two-three digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_netherlands_bsn finds content that matches the pattern.
- A keyword from Keyword_netherlands_bsn is found.
- The checksum passes.

```
<!-- Netherlands Citizen's Service (BSN) Number -->
<Entity id="c5f54253-ef7e-44f6-a578-440ed67e946d" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_netherlands_bsn" />
    <Match idRef="Keywords_netherlands_eu_national_id_card" />
  </Pattern>
</Entity>
```

Keywords

Keywords_netherlands_eu_national_id_card

- bsn#
- bsn
- burgerservicenummer
- citizen service number
- person number
- personal number
- personal numeric code
- person-related number
- persoonlijk nummer
- persoonlijke numerieke code
- persoonsgebonden
- persoonsnummer
- sociaal-fiscaal nummer
- social-fiscal number
- sofi
- sofinummer
- uniek identificatienummer
- uniek identiteitsnummer
- unique identification number
- unique identity number
- uniqueidentityno#

Netherlands driver's license number

Format

ten digits without spaces and delimiters

Pattern

ten digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_netherlands_eu_driver's_license_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_driver's_license_number` or `Keywords_netherlands_eu_driver's_license_number` is found.

```
<!-- Netherlands Driver's License Number -->
<Entity id="6247fbea-ab80-4be5-8233-308b7c031401" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_netherlands_eu_driver's_license_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_driver's_license_number" />
      <Match idRef="Keywords_netherlands_eu_driver's_license_number" />
      </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_driver's_license_number

- driverlic
- driverlics
- driverlicense
- driverlicenses
- driverlicence
- driverlicences
- driver lic
- driver lics
- driver license
- driver licenses
- driver licence
- driver licences
- driverslic
- driverslics
- driverslicence
- driverslicences
- driverslicense
- driverslicenses
- drivers lic
- drivers lics
- drivers license
- drivers licenses
- drivers licence
- drivers licences
- driver'lic
- driver'lics
- driver'license
- driver'licenses
- driver'licence
- driver'licences
- driver' lic
- driver' lics
- driver' license
- driver' licenses
- driver' licence
- driver' licences
- driver'slic
- driver'slics
- driver'slicense
- driver'slicenses
- driver'slicence
- driver'slicences
- driver's lic
- driver's lics
- driver's license

- driver's licenses
- driver's licence
- driver's licences
- dl#
- dls#
- driverlic#
- driverlics#
- driverlicense#
- driverlicenses#
- driverlicence#
- driverlicences#
- driver lic#
- driver lics#
- driver license#
- driver licenses#
- driver licences#
- driverslic#
- driverslics#
- driverslicense#
- driverslicenses#
- driverslicence#
- driverslicences#
- drivers lic#
- drivers lics#
- drivers license#
- drivers licenses#
- drivers licence#
- drivers licences#
- driver'lic#
- driver'lics#
- driver'license#
- driver'licenses#
- driver'licence#
- driver'licences#
- driver' lic#
- driver' lics#
- driver' license#
- driver' licenses#
- driver' licence#
- driver' licences#
- driver'slic#
- driver'slics#
- driver'slicense#
- driver'slicenses#
- driver'slicence#
- driver'slicences#

- driver's lic#
- driver's lics#
- driver's license#
- driver's licenses#
- driver's licence#
- driver's licences#
- driving licence
- driving license
- dlno#
- driv lic
- driv licen
- driv license
- driv licenses
- driv licence
- driv licences
- driver licen
- drivers licen
- driver's licen
- driving lic
- driving licen
- driving licenses
- driving licence
- driving licences
- driving permit
- dl no
- dlno
- dl number

Keywords_netherlands_eu_driver's_license_number

- permis de conduire
- rijbewijs
- rijbewijsnummer
- rijbewijzen
- rijbewijs nummer
- rijbewijsnummers

Netherlands passport number

Format

nine letters or digits with no spaces or delimiters

Pattern

nine letters or digits

Checksum

not applicable

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300

characters:

- The regular expression `Regex_netherlands_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_netherlands_eu_passport_number` is found.
- The regular expression `Regex_netherlands_eu_passport_date` finds date in the format DD MMM/MMM YYYY (Example - 26 MAA/MAR 2012)

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_netherlands_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_netherlands_eu_passport_number` is found.

```
<!-- Netherlands Passport Number -->
<Entity id="61786727-bafd-45f6-94d9-888d815e228e" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Regex_netherlands_eu_passport_number" />
    <Match idRef="Regex_netherlands_eu_passport_date" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_netherlands_eu_passport_number" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_netherlands_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_netherlands_eu_passport_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_passport_number

- passport#
- passport #
- passportid
- passports
- passportno
- passport no
- passportnumber
- passport number
- passportnumbers
- passport numbers

Keywords_netherlands_eu_passport_number

- paspoort nummer
- paspoortnummers
- paspoortnummer
- paspoort nr

Netherlands tax identification number

This sensitive information type is only available for use in:

- data loss prevention policies

- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

nine digits without spaces or delimiters

Pattern

nine digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_netherlands_eu_tax_file_number` finds content that matches the pattern.
- A keyword from `Keywords_netherlands_eu_tax_file_number` is found.

A DLP policy has low confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_netherlands_eu_tax_file_number` finds content that matches the pattern.

```
<!-- Netherlands Tax Identification Number -->
<Entity id="01f42a64-eba7-4892-a67b-398237e4ade2" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_netherlands_eu_tax_file_number" />
    <Match idRef="Keywords_netherlands_eu_tax_file_number" />
  </Pattern>
  <Pattern confidenceLevel="65">
    <IdMatch idRef="Func_netherlands_eu_tax_file_number" />
  </Pattern>
</Entity>
```

Keywords

Keywords_netherlands_eu_tax_file_number

- btw nummer
- hollânske tax identification
- hulandes impuesto id number
- hulandes impuesto identification
- identificatienummer belasting
- identificatienummer van belasting
- impuesto identification number
- impuesto number
- nederlands belasting id nummer
- nederlands belasting identificatie
- nederlands belasting identificatienummer
- nederlands belastingnummer
- nederlandse belasting identificatie
- netherlands tax identification

- netherland's tax identification
- netherlands tin
- netherland's tin
- tax id
- tax identification no
- tax identification number
- tax identification tal
- tax no#
- tax no
- tax number
- tax registration number
- tax tal
- taxid#
- taxidno#
- taxidnumber#
- taxno#
- taxnumber#
- taxnumber
- tin id
- tin no
- tin#

Netherlands value added tax number

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

14 character alphanumeric pattern

Pattern

14-character alphanumeric pattern:

- N or n
- L or l
- optional space, dot, or hyphen
- nine digits
- optional space, dot, or hyphen
- B or b
- two digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_netherlands_value_added_tax_number finds content that matches the pattern.
- A keyword from Keywords_netherlands_value_added_tax_number is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_netherlands_value_added_tax_number finds content that matches the pattern.

```
<!-- Netherlands Value Added Tax Number -->
<Entity id="4f320d9b-4972-41ae-b337-88d499bb1ade" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_netherlands_value_added_tax_number" />
    <Match idRef="Keywords_netherlands_value_added_tax_number" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_netherlands_value_added_tax_number" />
  </Pattern>
</Entity>
```

Keywords

Keyword_netherlands_value_added_tax_number

- vat number
- vat no
- vat#
- wearde tafoege tax getal
- btw nûmer
- btw-nummer

New Zealand bank account number

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

14-digit to 16-digit pattern with optional delimiter

Pattern

14-digit to 16-digit pattern with optional delimiter:

- two digits
- an optional hyphen or space
- three to four digits
- an optional hyphen or space
- seven digits
- an optional hyphen or space
- two to three digits

- an options hyphen or space

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_new_zealand_bank_account_number finds content that matches the pattern.
- A keyword from Keywords_new_zealand_bank_account_number is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_new_zealand_bank_account_number finds content that matches the pattern.

```
<!-- New Zealand Bank Account Number -->
<Entity id="1a97fc2b-dd2f-48f1-bc4e-2ddf25813956" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_new_zealand_bank_account_number" />
    <Match idRef="Keywords_new_zFealand_bank_account_number" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_new_zealand_bank_account_number" />
  </Pattern>
</Entity>
```

Keywords

Keyword_new_zealand_bank_account_number

- account number
- bank account
- bank_acct_id
- bank_acct_branch
- bank_acct_nbr

New Zealand driver's license number

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

eight character alphanumeric pattern

Pattern

eight character alphanumeric pattern

- two letters
- six digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_newzealand_driver_license_number finds content that matches the pattern.
- A keyword from Keywords_newzealand_driver_license_number is found.

A DLP policy has low confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_newzealand_driver_license_number finds content that matches the pattern.

```
<!-- New Zealand Driver License Number -->
<Entity id="1924b377-d287-49c9-a737-cfe7a8a2615a" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_newzealand_driver_license_number" />
    <Match idRef="Keywords_newzealand_driver_license_number" />
  </Pattern>
  <Pattern confidenceLevel="65">
    <IdMatch idRef="Func_newzealand_driver_license_number" />
  </Pattern>
</Entity>
```

Keywords

Keyword_new_zealand_drivers_license_number

- driverlicence
- driverlicences
- driver lic
- driver licence
- driver licences
- driverslic
- driverslicence
- driverslicences
- drivers lic
- drivers lics
- drivers licence
- drivers licences
- driver'lic
- driver'lics
- driver'licence
- driver'licences
- driver' lic
- driver' lics
- driver' licence
- driver' licences
- driver'slic
- driver'slics
- driver'slicence
- driver'slicences
- driver's lic

- driver's lics
- driver's licence
- driver's licences
- driverlic#
- driverlics#
- driverlicence#
- driverlicences#
- driver lic#
- driver lics#
- driver licence#
- driver licences#
- driverslic#
- driverslics#
- driverslicence#
- driverslicences#
- drivers lic#
- drivers lics#
- drivers licence#
- drivers licences#
- driver'lic#
- driver'lics#
- driver'licence#
- driver'licences#
- driver' lic#
- driver' lics#
- driver' licence#
- driver' licences#
- driver'slic#
- driver'slics#
- driver'slicence#
- driver'slicences#
- driver's lic#
- driver's lics#
- driver's licence#
- driver's licences#
- international driving permit
- international driving permits
- nz automobile association
- new zealand automobile association

New Zealand inland revenue number

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management

- Microsoft cloud app security

Format

eight or nine digits with optional delimiters

Pattern

eight or nine digits with optional delimiters

- two or three digits
- an optional space or hyphen
- three digits
- an optional space or hyphen
- three digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_new_zealand_inland_revenue_number finds content that matches the pattern.
- A keyword from Keywords_new_zealand_inland_revenue_number is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_new_zealand_inland_revenue_number finds content that matches the pattern.

```
<!-- New Zealand Inland Revenue Number -->
<Entity id="dd0fe2bc-7bcf-455f-bac1-83b1e3eb25fd" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_new_zealand_inland_revenue_number" />
    <Match idRef="Keywords_new_zealand_inland_revenue_number" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_new_zealand_inland_revenue_number" />
  </Pattern>
</Entity>
```

Keywords

Keyword_new_zealand_inland_revenue_number

- ird no.
- ird no#
- nz ird
- new zealand ird
- ird number
- inland revenue number

New Zealand ministry of health number

Format

three letters, a space (optional), and four digits

Pattern

- three letters (not case-sensitive) except 'I' and 'O'
- a space (optional)
- four digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_new_zealand_ministry_of_health_number finds content that matches the pattern.
- A keyword from Keyword_nz_terms is found.
- The checksum passes.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_new_zealand_ministry_of_health_number finds content that matches the pattern.
- The checksum passes.

```
<!-- New Zealand Health Number -->
<Entity id="2b71c1c8-d14e-4430-82dc-fd1ed6bf05c7" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_new_zealand_ministry_of_health_number" />
      <Match idRef="Keyword_nz_terms" />
  </Pattern>
  <Pattern confidenceLevel="75">
      <IdMatch idRef="Func_new_zealand_ministry_of_health_number" />
   </Pattern>
</Entity>
```

Keywords

Keyword_nz_terms

- NHI
- New Zealand
- Health
- treatment
- National Health Index Number
- nhi number
- nhi no.
- NHI#
- National Health Index No.
- National Health Index Id
- National Health Index #

New Zealand social welfare number

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management

- Microsoft cloud app security

Format

nine digits

Pattern

nine digits

- three digits
- an optional hyphen
- three digits
- an optional hyphen
- three digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_newzealand_social_welfare_number finds content that matches the pattern.
- A keyword from Keywords_newzealand_social_welfare_number is found.

A DLP policy has low confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_newzealand_social_welfare_number finds content that matches the pattern.

```
      <!-- Newzealand Social Welfare Number -->
      <Entity id="20f3c48d-4ac1-4cd2-86bd-34ecc1826e9d" patternsProximity="300" recommendedConfidence="85">
        <Pattern confidenceLevel="85">
          <IdMatch idRef="Func_newzealand_social_welfare_number" />
          <Match idRef="Keywords_newzealand_social_welfare_number" />
        </Pattern>
        <Pattern confidenceLevel="65">
          <IdMatch idRef="Func_newzealand_social_welfare_number" />
        </Pattern>
      </Entity>
    </Version>
```

Keywords

Keyword_new_zealand_social_welfare_number

- social welfare #
- social welfare#
- social welfare No.
- social welfare number
- swn#

Norway identification number

Format

11 digits

Pattern

11 digits:

- six digits in the format DDMMYY which are the date of birth
- three-digit individual number
- two check digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_norway_id_number finds content that matches the pattern.
- A keyword from Keyword_norway_id_number is found.
- The checksum passes.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_norway_id_numbe finds content that matches the pattern.
- The checksum passes.

```
<!-- Norway Identification Number -->
<Entity id="d4c8a798-e9f2-4bd3-9652-500d24080fc3" recommendedConfidence="85" patternsProximity="300">
  <Pattern confidenceLevel="85">
     <IdMatch idRef="Func_norway_id_number"/>
     <Match idRef="Keyword_norway_id_number"/>
  </Pattern>
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Func_norway_id_number"/>
  </Pattern>
</Entity>
```

Keywords

Keyword_norway_id_number

- Personal identification number
- Norwegian ID Number
- ID Number
- Identification
- Personnummer
- Fødselsnummer

Philippines unified multi-purpose identification number

Format

12 digits separated by hyphens

Pattern

12 digits:

- four digits
- a hyphen
- seven digits
- a hyphen

- one digit

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_philippines_unified_id finds content that matches the pattern.
- A keyword from Keyword_philippines_id is found.

```
<!-- Philippines Unified Multi-Purpose ID number -->
<Entity id="019b39dd-8c25-4765-91a3-d9c6baf3c3b3" recommendedConfidence="75" patternsProximity="300">
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Regex_philippines_unified_id"/>
     <Match idRef="Keyword_philippines_id"/>
  </Pattern>
</Entity>
```

Keywords

Keyword_philippines_id

- Unified Multi-Purpose ID
- UMID
- Identity Card
- Pinag-isang Multi-Layunin ID

Poland driver's license number

Format

14 digits containing two forward slashes

Pattern

14 digits and two forward slashes:

- five digits
- a forward slash
- two digits
- a forward slash
- seven digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_poland_eu_driver's_license_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_driver's_license_number` or `Keywords_poland_eu_driver's_license_number` is found.

```
<!-- Poland Driver's License Number -->
<Entity id="24d51f99-ee9e-4060-a077-cae58cab1ee4" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_poland_eu_driver's_license_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_driver's_license_number" />
      <Match idRef="Keywords_poland_eu_driver's_license_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_driver's_license_number

- driverlic
- driverlics
- driverlicense
- driverlicenses
- driverlicence
- driverlicences
- driver lic
- driver lics
- driver license
- driver licenses
- driver licence
- driver licences
- driverslic
- driverslics
- driverslicence
- driverslicences
- driverslicense
- driverslicenses
- drivers lic
- drivers lics
- drivers license
- drivers licenses
- drivers licence
- drivers licences
- driver'lic
- driver'lics
- driver'license
- driver'licenses
- driver'licence
- driver'licences
- driver' lic
- driver' lics
- driver' license
- driver' licenses
- driver' licence
- driver' licences

- driver'slic
- driver'slics
- driver'slicense
- driver'slicenses
- driver'slicence
- driver'slicences
- driver's lic
- driver's lics
- driver's license
- driver's licenses
- driver's licence
- driver's licences
- dl#
- dls#
- driverlic#
- driverlics#
- driverlicense#
- driverlicenses#
- driverlicence#
- driverlicences#
- driver lic#
- driver lics#
- driver license#
- driver licenses#
- driver licences#
- driverslic#
- driverslics#
- driverslicense#
- driverslicenses#
- driverslicence#
- driverslicences#
- drivers lic#
- drivers lics#
- drivers license#
- drivers licenses#
- drivers licence#
- drivers licences#
- driver'lic#
- driver'lics#
- driver'license#
- driver'licenses#
- driver'licence#
- driver'licences#
- driver' lic#
- driver' lics#
- driver' license#

- driver' licenses#
- driver' licence#
- driver' licences#
- driver'slic#
- driver'slics#
- driver'slicense#
- driver'slicenses#
- driver'slicence#
- driver'slicences#
- driver's lic#
- driver's lics#
- driver's license#
- driver's licenses#
- driver's licence#
- driver's licences#
- driving licence
- driving license
- dlno#
- driv lic
- driv licen
- driv license
- driv licenses
- driv licence
- driv licences
- driver licen
- drivers licen
- driver's licen
- driving lic
- driving licen
- driving licenses
- driving licence
- driving licences
- driving permit
- dl no
- dlno
- dl number

Keywords_poland_eu_driver's_license_number

- prawo jazdy
- prawa jazdy

Poland identity card

Format

three letters and six digits

Pattern

three letters (not case-sensitive) followed by six digits

Checksum

Yes

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_polish_national_id finds content that matches the pattern.
- A keyword from Keyword_polish_national_id_passport_number is found.
- The checksum passes.

```
<!-- Poland Identity Card-->
<Entity id="25E64989-ED5D-40CA-A939-6C14183BB7BF" patternsProximity="300" recommendedConfidence="85">
      <Pattern confidenceLevel="85">
          <IdMatch idRef="Func_polish_national_id" />
          <Match idRef="Keyword_polish_national_id_passport_number" />
      </Pattern>
</Entity>
```

Keywords

Keyword_poland_national_id_passport_number

- Dowód osobisty
- Numer dowodu osobistego
- Nazwa i numer dowodu osobistego
- Nazwa i nr dowodu osobistego
- Nazwa i nr dowodu tożsamości
- Dowód Tożsamości
- dow. os.

Poland national ID (PESEL)

Format

11 digits

Pattern

- six digits representing date of birth in the format YYMMDD
- four digits
- one check digit

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_pesel_identification_number finds content that matches the pattern.
- A keyword from Keyword_pesel_identification_number is found.
- The checksum passes.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_pesel_identification_number finds content that matches the pattern.

- The checksum passes.

```
<!-- Poland National ID (PESEL) -->
<Entity id="E3AAF206-4297-412F-9E06-BA8487E22456" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_pesel_identification_number" />
    <Match idRef="Keyword_pesel_identification_number" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_pesel_identification_number" />
  </Pattern>
</Entity>
```

Keywords

Keyword_pesel_identification_number

- dowód osobisty
- dowódosobisty
- niepowtarzalny numer
- niepowtarzalnynumer
- nr.-pesel
- nr-pesel
- numer identyfikacyjny
- pesel
- tożsamości narodowej

Poland passport number

This sensitive information type entity is included in the EU Passport Number sensitive information type. It's available as a stand-alone sensitive information type entity.

Format

two letters and seven digits

Pattern

Two letters (not case-sensitive) followed by seven digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_polish_passport_number_v2` finds content that matches the pattern.
- The checksum passes.
- A keyword from `Keywords_eu_passport_number` or `Keyword_polish_national_passport_number` is found.
- A keyword from `Keywords_eu_passport_date` is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_polish_passport_number_v2` finds content that matches the pattern.
- The checksum passes.
- A keyword from `Keywords_eu_passport_number` or `Keyword_polish_national_passport_number` is found.

A DLP policy has low confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_polish_passport_number_v2` finds content that matches the pattern.
- The checksum passes.

```
<!-- Poland Passport Number -->
<Entity id="03937FB5-D2B6-4487-B61F-0F8BFF7C3517" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_polish_passport_number_v2" />
    <Match idRef="Keywords_eu_passport_date" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keyword_polish_national_passport_number" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_polish_passport_number_v2" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keyword_polish_national_passport_number" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="65">
    <IdMatch idRef="Func_polish_passport_number_v2" />
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_passport_number

- passport#
- passport #
- passportid
- passports
- passportno
- passport no
- passportnumber
- passport number
- passportnumbers
- passport numbers

Keyword_polish_national_passport_number

- numer paszportu
- numery paszportów
- numery paszportowe
- nr paszportu
- nr. paszportu
- nr paszportów
- n° passeport
- passeport n°

Keywords_eu_passport_date

- date of issue
- date of expiry

Poland REGON number

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

9-digit or 14-digit number

Pattern

nine digit or 14-digit number:

- nine digits or
- nine digits
- hyphen
- five digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_polish_regon_number finds content that matches the pattern.
- A keyword from Keywords_polish_regon_number is found.

A DLP policy has low confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_polish_regon_number finds content that matches the pattern.

```
<!-- Polish REGON Number  -->
<Entity id="fc87b421-f437-4f8b-b739-29a735ead0d9" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_polish_regon_number" />
    <Match idRef="Keywords_polish_regon_number" />
  </Pattern>
  <Pattern confidenceLevel="65">
    <IdMatch idRef="Func_polish_regon_number" />
  </Pattern>
</Entity>
```

Keywords

Keywords_poland_regon_number

- regon id
- statistical number
- statistical id
- statistical no
- regon number
- regonid#

- regonno#
- company id
- companyid#
- companyidno#
- numer statystyczny
- numeru regon
- numerstatystyczny#
- numeruregon#

Poland tax identification number

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

11 digits with no spaces or delimiters

Pattern

11 digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_poland_eu_tax_file_number` finds content that matches the pattern.
- A keyword from `Keywords_poland_eu_tax_file_number` is found.

```
<!-- Poland Tax Identification Number -->
<Entity id="1ff28b4d-40f2-49e9-b677-9606a88e2bca" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_poland_eu_tax_file_number" />
    <Match idRef="Keywords_poland_eu_tax_file_number" />
  </Pattern>
</Entity>
```

Keywords

Keywords_poland_eu_tax_file_number

- nip#
- nip
- numer identyfikacji podatkowej
- numeridentyfikacjipodatkowej#
- tax id
- tax identification no
- tax identification number

- tax no#
- tax no
- tax number
- tax registration number
- taxid#
- taxidno#
- taxidnumber#
- taxno#
- taxnumber#
- taxnumber
- tin id
- tin no
- tin#
- vat id#
- vat id
- vat no
- vat number
- vatid#
- vatid
- vatno#

Portugal citizen card number

Format

eight digits

Pattern

eight digits

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_portugal_citizen_card finds content that matches the pattern.
- A keyword from Keyword_portugal_citizen_card is found.

```
<!-- Portugal Citizen Card Number -->
<Entity id="91a7ece2-add4-4986-9a15-c84544d81ecd" recommendedConfidence="85" patternsProximity="300">
  <Pattern confidenceLevel="85">
     <IdMatch idRef="Regex_portugal_citizen_card"/>
     <Match idRef="Keyword_portugal_citizen_card"/>
  </Pattern>
</Entity>
```

Keywords

Keyword_portugal_citizen_card

- bilhete de identidade
- cartão de cidadão

- citizen card
- document number
- documento de identificação
- id number
- identification no
- identification number
- identity card no
- identity card number
- national id card
- nic
- número bi de portugal
- número de identificação civil
- número de identificação fiscal
- número do documento
- portugal bi number

Portugal driver's license number

Format

two patterns - two letters followed by 5-8 digits with special characters

Pattern

Pattern 1: Two letters followed by 5/6 with special characters:

- Two letters (not case-sensitive)
- A hyphen
- Five or Six digits
- A space
- One digit

Pattern 2: One letter followed by 6/8 digits with special characters:

- One letter (not case-sensitive)
- A hyphen
- Six or eight digits
- A space
- One digit

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_portugal_eu_driver's_license_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_driver's_license_number` or `Keywords_portugal_eu_driver's_license_number` is found.

```
<!-- Portugal Driver's License Number -->
<Entity id="977f1e5a-2c33-4bcc-b516-95bb275cff23" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_portugal_eu_driver's_license_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_driver's_license_number" />
      <Match idRef="Keywords_portugal_eu_driver's_license_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_driver's_license_number

- driverlic
- driverlics
- driverlicense
- driverlicenses
- driverlicence
- driverlicences
- driver lic
- driver lics
- driver license
- driver licenses
- driver licence
- driver licences
- driverslic
- driverslics
- driverslicence
- driverslicences
- driverslicense
- driverslicenses
- drivers lic
- drivers lics
- drivers license
- drivers licenses
- drivers licence
- drivers licences
- driver'lic
- driver'lics
- driver'license
- driver'licenses
- driver'licence
- driver'licences
- driver' lic
- driver' lics
- driver' license
- driver' licenses
- driver' licence
- driver' licences

- driver'slic
- driver'slics
- driver'slicense
- driver'slicenses
- driver'slicence
- driver'slicences
- driver's lic
- driver's lics
- driver's license
- driver's licenses
- driver's licence
- driver's licences
- dl#
- dls#
- driverlic#
- driverlics#
- driverlicense#
- driverlicenses#
- driverlicence#
- driverlicences#
- driver lic#
- driver lics#
- driver license#
- driver licenses#
- driver licences#
- driverslic#
- driverslics#
- driverslicense#
- driverslicenses#
- driverslicence#
- driverslicences#
- drivers lic#
- drivers lics#
- drivers license#
- drivers licenses#
- drivers licence#
- drivers licences#
- driver'lic#
- driver'lics#
- driver'license#
- driver'licenses#
- driver'licence#
- driver'licences#
- driver' lic#
- driver' lics#
- driver' license#

- driver' licenses#
- driver' licence#
- driver' licences#
- driver'slic#
- driver'slics#
- driver'slicense#
- driver'slicenses#
- driver'slicence#
- driver'slicences#
- driver's lic#
- driver's lics#
- driver's license#
- driver's licenses#
- driver's licence#
- driver's licences#
- driving licence
- driving license
- dlno#
- driv lic
- driv licen
- driv license
- driv licenses
- driv licence
- driv licences
- driver licen
- drivers licen
- driver's licen
- driving lic
- driving licen
- driving licenses
- driving licence
- driving licences
- driving permit
- dl no
- dlno
- dl number

Keywords_portugal_eu_driver's_license_number

- carteira de motorista
- carteira motorista
- carteira de habilitação
- carteira habilitação
- número de licença
- número licença
- permissão de condução
- permissão condução
- Licença condução Portugal

- carta de condução

Portugal passport number

Format

one letter followed by six digits with no spaces or delimiters

Pattern

one letter followed by six digits:

- one letter (not case-sensitive)
- six digits

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_portugal_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_portugal_eu_passport_number` is found.
- The regular expression `Regex_eu_passport_date1` finds date in the format DD.MM.YYYY or a keyword from `Keywords_eu_passport_date` is found

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_portugal_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_portugal_eu_passport_number` is found.

```
<!-- Portugal Passport Number -->
<Entity id="080a52fd-a7bc-431e-b54d-51f08f59db11" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Regex_portugal_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_portugal_eu_passport_number" />
    </Any>
    <Any minMatches="1">
      <Match idRef="Regex_eu_passport_date1" />
      <Match idRef="Keywords_eu_passport_date" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_portugal_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_portugal_eu_passport_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_passport_number

- passport#
- passport #

- passportid
- passports
- passportno
- passport no
- passportnumber
- passport number
- passportnumbers
- passport numbers

Keywords_portugal_eu_passport_number

- número do passaporte
- portuguese passport
- portuguese passeport
- portuguese passaporte
- passaporte nº
- passeport nº
- números de passaporte
- portuguese passports
- número passaporte
- números passaporte

Keywords_eu_passport_date

- date of issue
- date of expiry

Portugal tax identification number

Format

nine digits with optional spaces

Pattern

- three digits
- an optional space
- three digits
- an optional space
- three digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_portugal_eu_tax_file_number` finds content that matches the pattern.
- A keyword from `Keywords_portugal_eu_tax_file_number` is found.

A DLP policy has low confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_portugal_eu_tax_file_number` finds content that matches the pattern.

```
<!-- Portugal Tax Identification Number -->
<Entity id="65372402-3131-4f1e-9983-4439841d1f15" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_portugal_eu_tax_file_number" />
    <Match idRef="Keywords_portugal_eu_tax_file_number" />
  </Pattern>
  <Pattern confidenceLevel="65">
    <IdMatch idRef="Func_portugal_eu_tax_file_number" />
  </Pattern>
</Entity>
```

Keywords

Keywords_portugal_eu_tax_file_number

- cpf#
- cpf
- nif#
- nif
- número de identificação fisca
- numero fiscal
- tax id
- tax identification no
- tax identification number
- tax no#
- tax no
- tax number
- tax registration number
- taxid#
- taxidno#
- taxidnumber#
- taxno#
- taxnumber#
- taxnumber
- tin id
- tin no
- tin#

Romania driver's license number

Format

one character followed by eight digits

Pattern

one character followed by eight digits:

- one letter (not case-sensitive) or digit
- eight digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of

300 characters:

- The regular expression `Regex_romania_eu_driver's_license_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_driver's_license_number` or `Keywords_romania_eu_driver's_license_number` is found.

```
<!-- Romania Driver's License Number -->
<Entity id="b5511ace-2fd8-4ae4-b6fc-c7c6e4689e3c" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_romania_eu_driver's_license_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_driver's_license_number" />
      <Match idRef="Keywords_romania_eu_driver's_license_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_driver's_license_number

- driverlic
- driverlics
- driverlicense
- driverlicenses
- driverlicence
- driverlicences
- driver lic
- driver lics
- driver license
- driver licenses
- driver licence
- driver licences
- driverslic
- driverslics
- driverslicence
- driverslicences
- driverslicense
- driverslicenses
- drivers lic
- drivers lics
- drivers license
- drivers licenses
- drivers licence
- drivers licences
- driver'lic
- driver'lics
- driver'license
- driver'licenses
- driver'licence
- driver'licences
- driver' lic

- driver' lics
- driver' license
- driver' licenses
- driver' licence
- driver' licences
- driver'slic
- driver'slics
- driver'slicense
- driver'slicenses
- driver'slicence
- driver'slicences
- driver's lic
- driver's lics
- driver's license
- driver's licenses
- driver's licence
- driver's licences
- dl#
- dls#
- driverlic#
- driverlics#
- driverlicense#
- driverlicenses#
- driverlicence#
- driverlicences#
- driver lic#
- driver lics#
- driver license#
- driver licenses#
- driver licences#
- driverslic#
- driverslics#
- driverslicense#
- driverslicenses#
- driverslicence#
- driverslicences#
- drivers lic#
- drivers lics#
- drivers license#
- drivers licenses#
- drivers licence#
- drivers licences#
- driver'lic#
- driver'lics#
- driver'license#
- driver'licenses#

- driver'licence#
- driver'licences#
- driver' lic#
- driver' lics#
- driver' license#
- driver' licenses#
- driver' licence#
- driver' licences#
- driver'slic#
- driver'slics#
- driver'slicense#
- driver'slicenses#
- driver'slicence#
- driver'slicences#
- driver's lic#
- driver's lics#
- driver's license#
- driver's licenses#
- driver's licence#
- driver's licences#
- driving licence
- driving license
- dlno#
- driv lic
- driv licen
- driv license
- driv licenses
- driv licence
- driv licences
- driver licen
- drivers licen
- driver's licen
- driving lic
- driving licen
- driving licenses
- driving licence
- driving licences
- driving permit
- dl no
- dlno
- dl number

Keywords_romania_eu_driver's_license_number

- permis de conducere
- permisului de conducere
- permisului conducere
- permisele de conducere

- permisele conducere
- permis conducere

Romania personal numeric code (CNP)

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

13 digits without spaces and delimiters

Pattern

- one digit from 1-9
- six digits representing date of birth (YYMMDD)
- two digits, which can be 01-52 or 99
- four digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_romania_eu_national_id_card` finds content that matches the pattern.
- A keyword from `Keywords_romania_eu_national_id_card` is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_romania_eu_national_id_card` finds content that matches the pattern.

```
<!-- Romania Personal Numerical Code (CNP) -->
<Entity id="eb5fa399-fe28-4c67-8188-d63a616ed89c" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_romania_eu_national_id_card" />
    <Match idRef="Keywords_romania_eu_national_id_card" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_romania_eu_national_id_card" />
  </Pattern>
</Entity>
```

Keywords

Keywords_romania_eu_national_id_card

- cnp#
- cnp
- cod identificare personal
- cod numeric personal
- cod unic identificare

- codnumericpersonal#
- codul fiscal nr.
- identificarea fiscală nr#
- id-ul taxei
- insurance number
- insurancenumber#
- national id#
- national id
- national identification number
- număr identificare personal
- număr identitate
- număr personal unic
- număridentitate#
- număridentitate
- numărpersonalunic#
- numărpersonalunic
- număru de identificare fiscală
- numărul de identificare fiscală
- personal numeric code
- pin#
- pin
- tax file no
- tax file number
- tax id
- tax identification no
- tax identification number
- tax no#
- tax no
- tax number
- tax registration number
- taxid#
- taxidno#
- taxidnumber#
- taxno#
- taxnumber#
- taxnumber
- tin id
- tin no
- tin#
- unique identification number
- unique identity number
- uniqueidentityno#
- uniqueidentityno

Romania passport number

Format

eight or nine digits without spaces and delimiters

Pattern

eight or nine digits

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_romania_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_romania_eu_passport_number` is found.
- The regular expression `Regex_romania_eu_passport_date` finds date in the format DD MMM/MMM YY (Example- 01 FEB/FEB 10) or a keyword from `Keywords_eu_passport_date` is found

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_romania_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_romania_eu_passport_number` is found.

```
<!-- Romania Passport Number -->
<Entity id="5d31b90c-7fe2-4a76-a14b-767b8fd19d6c" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Regex_romania_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_romania_eu_passport_number" />
    </Any>
    <Any minMatches="1">
      <Match idRef="Regex_romania_eu_passport_date" />
      <Match idRef="Keywords_eu_passport_date" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_romania_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_romania_eu_passport_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_passport_number

- passport#
- passport #
- passportid
- passports
- passportno
- passport no
- passportnumber
- passport number
- passportnumbers

- passport numbers

Keywords_romania_eu_passport_number

numărul pașaportului numarul pasaportului numerele pașaportului Pașaport nr

Keywords_eu_passport_date

- date of issue
- date of expiry

Russia passport number domestic

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

10-digit number

Pattern

10-digit number:

- two digits
- an optional space or hyphen
- two digits
- an optional space
- six digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regex Regex_Russian_Passport_Number_Domestic finds content that matches the pattern.
- A keyword from Keyword_Russian_Passport_Number is found.

```
<!-- Russian Passport Number Domestic -->
<Entity id="76ec2f5d-cedb-48e1-8070-1998794af445" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_Russian_Passport_Number_Domestic" />
    <Match idRef="Keyword_Russian_Passport_Number" />
  </Pattern>
</Entity>
```

Keywords

Keyword_russia_passport_number_domestic

- passport number
- passport no
- passport #
- passport id

- passportno#
- passportnumber#
- паспорт нет
- паспорт id
- российской паспорт
- русский номер паспорта
- паспорт#
- паспортid#
- номер паспорта
- номерпаспорта#

Russia passport number international

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

nine-digit number

Pattern

nine-digit number:

- two digits
- an optional space or hyphen
- seven digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regex Regex_Russian_Passport_Number_International finds content that matches the pattern.
- A keyword from Keyword_Russian_Passport_Number is found.

```
<!-- Russian Passport Number International -->
<Entity id="ac5f4878-75e4-4b82-af2d-02e13ea9f411" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_Russian_Passport_Number_International" />
    <Match idRef="Keyword_Russian_Passport_Number" />
  </Pattern>
</Entity>
```

Keywords

Keywords_russia_passport_number_international

- passport number
- passport no

- passport #
- passport id
- passportno#
- passportnumber#
- паспорт нет
- паспорт id
- российской паспорт
- русский номер паспорта
- паспорт#
- паспортid#
- номер паспорта
- номерпаспорта#

Saudi Arabia National ID

Format

10 digits

Pattern

10 consecutive digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_saudi_arabia_national_id finds content that matches the pattern.
- A keyword from Keyword_saudi_arabia_national_id is found.

```
<!-- Saudi Arabia National ID -->
<Entity id="8c5a0ba8-404a-41a3-8871-746aa21ee6c0" patternsProximity="300" recommendedConfidence="75">
    <Pattern confidenceLevel="75">
        <IdMatch idRef="Regex_saudi_arabia_national_id" />
        <Any minMatches="1">
          <Match idRef="Keyword_saudi_arabia_national_id" />
        </Any>
    </Pattern>
</Entity>
```

Keywords

Keyword_saudi_arabia_national_id

- Identification Card
- I card number
- ID number
- الوطنية الهوية بطاقة رقم

Singapore national registration identity card (NRIC) number

Format

nine letters and digits

Pattern

- nine letters and digits:
- the letter "F", "G", "S", or "T" (not case-sensitive)
- seven digits
- an alphabetic check digit

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_singapore_nric finds content that matches the pattern.
- A keyword from Keyword_singapore_nric is found.
- The checksum passes.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_singapore_nric finds content that matches the pattern.
- The checksum passes.

```
<!-- Singapore National Registration Identity Card (NRIC) Number -->
<Entity id="cead390a-dd83-4856-9751-fb6dc98c34da" recommendedConfidence="75" patternsProximity="300">
  <Pattern confidenceLevel="85">
     <IdMatch idRef="Regex_singapore_nric"/>
     <Match idRef="Keyword_singapore_nric"/>
  </Pattern>
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Regex_singapore_nric"/>
  </Pattern>
</Entity>
```

Keywords

Keyword_singapore_nric

- National Registration Identity Card
- Identity Card Number
- NRIC
- IC
- Foreign Identification Number
- FIN
- 身份证
- 身份證

Slovakia driver's license number

Format

one character followed by seven digits

Pattern

one character followed by seven digits

- one letter (not case-sensitive) or digit

- seven digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_slovakia_eu_driver's_license_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_driver's_license_number` or `Keywords_slovakia_eu_driver's_license_number` is found.

```
<!-- Slovakia Driver's License Number -->
<Entity id="14240c22-b6de-4ce5-a90b-137f74252513" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_slovakia_eu_driver's_license_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_driver's_license_number" />
      <Match idRef="Keywords_slovakia_eu_driver's_license_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_driver's_license_number

- driverlic
- driverlics
- driverlicense
- driverlicenses
- driverlicence
- driverlicences
- driver lic
- driver lics
- driver license
- driver licenses
- driver licence
- driver licences
- driverslic
- driverslics
- driverslicence
- driverslicences
- driverslicense
- driverslicenses
- drivers lic
- drivers lics
- drivers license
- drivers licenses
- drivers licence
- drivers licences
- driver'lic

- driver'lics
- driver'license
- driver'licenses
- driver'licence
- driver'licences
- driver' lic
- driver' lics
- driver' license
- driver' licenses
- driver' licence
- driver' licences
- driver'slic
- driver'slics
- driver'slicense
- driver'slicenses
- driver'slicence
- driver'slicences
- driver's lic
- driver's lics
- driver's license
- driver's licenses
- driver's licence
- driver's licences
- dl#
- dls#
- driverlic#
- driverlics#
- driverlicense#
- driverlicenses#
- driverlicence#
- driverlicences#
- driver lic#
- driver lics#
- driver license#
- driver licenses#
- driver licences#
- driverslic#
- driverslics#
- driverslicense#
- driverslicenses#
- driverslicence#
- driverslicences#
- drivers lic#
- drivers lics#
- drivers license#
- drivers licenses#

- drivers licence#
- drivers licences#
- driver'lic#
- driver'lics#
- driver'license#
- driver'licenses#
- driver'licence#
- driver'licences#
- driver' lic#
- driver' lics#
- driver' license#
- driver' licenses#
- driver' licence#
- driver' licences#
- driver'slic#
- driver'slics#
- driver'slicense#
- driver'slicenses#
- driver'slicence#
- driver'slicences#
- driver's lic#
- driver's lics#
- driver's license#
- driver's licenses#
- driver's licence#
- driver's licences#
- driving licence
- driving license
- dlno#
- driv lic
- driv licen
- driv license
- driv licenses
- driv licence
- driv licences
- driver licen
- drivers licen
- driver's licen
- driving lic
- driving licen
- driving licenses
- driving licence
- driving licences
- driving permit
- dl no
- dlno

- dl number

Keywords_slovakia_eu_driver's_license_number

- vodičský preukaz
- vodičské preukazy
- vodičského preukazu
- vodičských preukazov

Slovakia personal number

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

nine or ten digits containing optional backslash

Pattern

- six digits representing date of birth
- optional slash (/)
- three digits
- one optional check digit

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_slovakia_eu_national_id_card` finds content that matches the pattern.
- A keyword from `Keywords_slovakia_eu_national_id_card` is found.

A DLP policy has low confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_slovakia_eu_national_id_card` finds content that matches the pattern.

```
    <!-- Slovakia Personal Number -->
    <Entity id="951c26b7-3b35-4f73-924b-15dd599cb9ab" patternsProximity="300" recommendedConfidence="85">
      <Pattern confidenceLevel="85">
        <IdMatch idRef="Func_slovakia_eu_national_id_card" />
        <Match idRef="Keywords_slovakia_eu_national_id_card" />
      </Pattern>
      <Pattern confidenceLevel="65">
        <IdMatch idRef="Func_slovakia_eu_national_id_card" />
      </Pattern>
    </Entity>
  </Version>
```

Keywords

Keywords_slovakia_eu_national_id_card

- azonosító szám
- birth number
- číslo národnej identifikačnej karty
- číslo občianského preukazu
- daňové číslo
- id number
- identification no
- identification number
- identifikačná karta č
- identifikačné číslo
- identity card no
- identity card number
- národná identifikačná značka č
- national number
- nationalnumber#
- nemzeti személyazonosító igazolvány
- personalidnumber#
- rč
- rodne cislo
- rodné číslo
- social security number
- ssn#
- ssn
- személyi igazolvány szám
- személyi igazolvány száma
- személyigazolvány szám
- tax file no
- tax file number
- tax id
- tax identification no
- tax identification number
- tax no#
- tax no
- tax number
- tax registration number
- taxid#
- taxidno#
- taxidnumber#
- taxno#
- taxnumber#
- taxnumber
- tin id
- tin no
- tin#

Slovakia passport number

Format

one digit or letter followed by seven digits with no spaces or delimiters

Pattern

one digit or letter (not case-sensitive) followed by seven digits

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_slovakia_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_slovakia_eu_passport_number` is found.
- The regular expression `Regex_eu_passport_date1` finds date in the format DD.MM.YYYY or a keyword from `Keywords_eu_passport_date` is found

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_slovakia_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_slovakia_eu_passport_number` is found.

```
<!-- Slovakia Passport Number -->
<Entity id="238e1f08-d80e-4793-af33-9b57918335b7" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Regex_slovakia_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_slovakia_eu_passport_number" />
    </Any>
    <Any minMatches="1">
      <Match idRef="Regex_eu_passport_date1" />
      <Match idRef="Keywords_eu_passport_date" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_slovakia_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_slovakia_eu_passport_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_passport_number

- passport#
- passport #
- passportid
- passports
- passportno
- passport no
- passportnumber
- passport number

- passportnumbers
- passport numbers

Keywords_slovakia_eu_passport_number

- číslo pasu
- čísla pasov
- pas č.
- Passeport n°
- n° Passeport

Keywords_eu_passport_date

- date of issue
- date of expiry

Slovenia driver's license number

Format

nine digits without spaces and delimiters

Pattern

nine digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_slovenia_eu_driver's_license_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_driver's_license_number` or `Keywords_slovenia_eu_driver's_license_number` is found.

```
<!-- Slovenia Driver's License Number -->
<Entity id="d5bc089a-f2ee-433d-a6b1-5c253051d6f2" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_slovenia_eu_driver's_license_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_driver's_license_number" />
      <Match idRef="Keywords_slovenia_eu_driver's_license_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_driver's_license_number

- driverlic
- driverlics
- driverlicense
- driverlicenses
- driverlicence
- driverlicences
- driver lic
- driver lics

- driver license
- driver licenses
- driver licence
- driver licences
- driverslic
- driverslics
- driverslicence
- driverslicences
- driverslicense
- driverslicenses
- drivers lic
- drivers lics
- drivers license
- drivers licenses
- drivers licence
- drivers licences
- driver'lic
- driver'lics
- driver'license
- driver'licenses
- driver'licence
- driver'licences
- driver' lic
- driver' lics
- driver' license
- driver' licenses
- driver' licence
- driver' licences
- driver'slic
- driver'slics
- driver'slicense
- driver'slicenses
- driver'slicence
- driver'slicences
- driver's lic
- driver's lics
- driver's license
- driver's licenses
- driver's licence
- driver's licences
- dl#
- dls#
- driverlic#
- driverlics#
- driverlicense#
- driverlicenses#

- driverlicence#
- driverlicences#
- driver lic#
- driver lics#
- driver license#
- driver licenses#
- driver licences#
- driverslic#
- driverslics#
- driverslicense#
- driverslicenses#
- driverslicence#
- driverslicences#
- drivers lic#
- drivers lics#
- drivers license#
- drivers licenses#
- drivers licence#
- drivers licences#
- driver'lic#
- driver'lics#
- driver'license#
- driver'licenses#
- driver'licence#
- driver'licences#
- driver' lic#
- driver' lics#
- driver' license#
- driver' licenses#
- driver' licence#
- driver' licences#
- driver'slic#
- driver'slics#
- driver'slicense#
- driver'slicenses#
- driver'slicence#
- driver'slicences#
- driver's lic#
- driver's lics#
- driver's license#
- driver's licenses#
- driver's licence#
- driver's licences#
- driving licence
- driving license
- dlno#

- driv lic
- driv licen
- driv license
- driv licenses
- driv licence
- driv licences
- driver licen
- drivers licen
- driver's licen
- driving lic
- driving licen
- driving licenses
- driving licence
- driving licences
- driving permit
- dl no
- dlno
- dl number

Keywords_slovenia_eu_driver's_license_number

- vozniško dovoljenje
- vozniška številka licence
- vozniških dovoljenj
- številka vozniškega dovoljenja
- številke vozniških dovoljenj

Slovenia Unique Master Citizen Number

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

13 digits without spaces or delimiters

Pattern

13 digits in the specified pattern:

- seven digits that correspond to the birth date (DDMMLLL) where "LLL" corresponds to the last three digits of the birth year
- two digits that correspond to the area of birth "50"
- three digits that correspond to a combination of gender and serial number for persons born on the same day (000-499 for male and 500-999 for female)
- one check digit

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_slovenia_eu_national_id_card` finds content that matches the pattern.
- A keyword from `Keywords_slovenia_eu_national_id_card` is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_slovenia_eu_national_id_card` finds content that matches the pattern.

```
<!-- Slovenia Unique Master Citizen Number -->
<Entity id="68948b27-803d-41e4-adf1-13e05eb541bb" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_slovenia_eu_national_id_card" />
    <Match idRef="Keywords_slovenia_eu_national_id_card" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_slovenia_eu_national_id_card" />
  </Pattern>
</Entity>
```

Keywords

Keywords_slovenia_eu_national_id_card

- edinstvena številka glavnega državljana
- emšo
- enotna maticna številka obcana
- id card
- identification number
- identifikacijska številka
- identity card
- nacionalna id
- nacionalni potni list
- national id
- osebna izkaznica
- osebni koda
- osebni ne
- osebni številka
- personal code
- personal number
- personal numeric code
- številka državljana
- unique citizen number
- unique id number
- unique identity number
- unique master citizen number
- unique registration number
- uniqueidentityno #
- uniqueidentityno#

Slovenia passport number

Format

two letters followed by seven digits with no spaces or delimiters

Pattern

two letters followed by seven digits:

- the letter "P"
- one uppercase letter
- seven digits

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_slovenia_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_slovenia_eu_passport_number` is found.
- The regular expression `Regex_eu_passport_date1` finds date in the format DD.MM.YYYY or a keyword from `Keywords_eu_passport_date` is found

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_slovenia_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_slovenia_eu_passport_number` is found.

```
<!-- Slovenia Passport Number -->
<Entity id="235b7976-7bbe-4df5-bb40-08678e749d1a" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Regex_slovenia_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_slovenia_eu_passport_number" />
    </Any>
    <Any minMatches="1">
      <Match idRef="Regex_eu_passport_date1" />
      <Match idRef="Keywords_eu_passport_date" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_slovenia_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_slovenia_eu_passport_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_passport_number

- passport#
- passport #
- passportid

- passports
- passportno
- passport no
- passportnumber
- passport number
- passportnumbers
- passport numbers

Keywords_slovenia_eu_passport_number

- številka potnega lista
- potek veljavnosti
- potni list#
- datum rojstva
- potni list
- številke potnih listov

Keywords_eu_passport_date

- date of issue
- date of expiry

Slovenia tax identification number

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

eight digits with no spaces or delimiters

Pattern

- one digit from 1-9
- six digits
- one check digit

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_slovenia_eu_tax_file_number` finds content that matches the pattern.
- A keyword from `Keywords_slovenia_eu_tax_file_number` is found.

A DLP policy has low confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_slovenia_eu_tax_file_number` finds content that matches the pattern.

```
<!-- Slovenia Tax Identification Number -->
<Entity id="e47b071e-c352-4d70-8241-8c215ad65505" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_slovenia_eu_tax_file_number" />
    <Match idRef="Keywords_slovenia_eu_tax_file_number" />
  </Pattern>
  <Pattern confidenceLevel="65">
    <IdMatch idRef="Func_slovenia_eu_tax_file_number" />
  </Pattern>
</Entity>
```

Keywords

Keywords_slovenia_eu_tax_file_number

- davčna številka
- identifikacijska številka davka
- številka davčne datoteke
- tax file no
- tax file number
- tax id
- tax identification no
- tax identification number
- tax no#
- tax no
- tax number
- tax registration number
- taxid#
- taxidno#
- taxidnumber#
- taxno#
- taxnumber#
- taxnumber
- tin id
- tin no
- tin#

South Africa identification number

Format

13 digits that may contain spaces

Pattern

13 digits:

- six digits in the format YYMMDD, which are the date of birth
- four digits
- a single-digit citizenship indicator
- the digit "8" or "9"
- one digit, which is a checksum digit

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_south_africa_identification_number finds content that matches the pattern.
- A keyword from Keyword_south_africa_identification_number is found.
- The checksum passes.

```
<!-- South Africa Identification Number -->
<Entity id="e2adf7cb-8ea6-4048-a2ed-d89eb65f2780" recommendedConfidence="85" patternsProximity="300">
  <Pattern confidenceLevel="85">
     <IdMatch idRef="Func_south_africa_identification_number"/>
     <Match idRef="Keyword_south_africa_identification_number"/>
  </Pattern>
</Entity>
```

Keywords

Keyword_south_africa_identification_number

- Identity card
- ID
- Identification

South Korea resident registration number

Format

13 digits containing a hyphen

Pattern

13 digits:

- six digits in the format YYMMDD, which are the date of birth
- a hyphen
- one digit determined by the century and gender
- four-digit region-of-birth code
- one digit used to differentiate people for whom the preceding numbers are identical
- a check digit.

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_south_korea_resident_number finds content that matches the pattern.
- A keyword from Keyword_south_korea_resident_number is found.
- The checksum passes.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_south_korea_resident_number finds content that matches the pattern.
- The checksum passes.

```
<!-- South Korea Resident Registration Number -->
<Entity id="5b802e18-ba80-44c4-bc83-bf2ad36ae36a" recommendedConfidence="85" patternsProximity="300">
  <Pattern confidenceLevel="85">
     <IdMatch idRef="Func_south_korea_resident_number"/>
     <Match idRef="Keyword_south_korea_resident_number"/>
  </Pattern>
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Func_south_korea_resident_number"/>
  </Pattern>
</Entity>
```

Keywords

Keyword_south_korea_resident_number

- National ID card
- Citizen's Registration Number
- Jumin deungnok beonho
- RRN
- 주민등록번호

Spain driver's license number

Format

eight digits followed by one character

Pattern

eight digits followed by one character:

- eight digits
- one digit or letter (not case-sensitive)

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_spain_eu_DL_and_NI_number_citizen` or `Func_spain_eu_DL_and_NI_number_foreigner` finds content that matches the pattern.
- A keyword from `Keywords_eu_driver's_license_number` or `Keywords_spain_eu_driver's_license_number` is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_spain_eu_DL_and_NI_number_citizen` or `Func_spain_eu_DL_and_NI_number_foreigner` finds content that matches the pattern.

```
<!-- Spain Driver's License Number -->
<Entity id="d5a82922-b501-4f40-8868-341321146aa2" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_spain_eu_DL_and_NI_number_citizen" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_driver's_license_number" />
      <Match idRef="Keywords_spain_eu_driver's_license_number" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_spain_eu_DL_and_NI_number_citizen" />
  </Pattern>
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_spain_eu_DL_and_NI_number_foreigner" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_driver's_license_number" />
      <Match idRef="Keywords_spain_eu_driver's_license_number" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_spain_eu_DL_and_NI_number_foreigner" />
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_driver's_license_number

- driverlic
- driverlics
- driverlicense
- driverlicenses
- driverlicence
- driverlicences
- driver lic
- driver lics
- driver license
- driver licenses
- driver licence
- driver licences
- driverslic
- driverslics
- driverslicence
- driverslicences
- driverslicense
- driverslicenses
- drivers lic
- drivers lics
- drivers license
- drivers licenses
- drivers licence
- drivers licences
- driver'lic
- driver'lics
- driver'license

- driver'licenses
- driver'licence
- driver'licences
- driver' lic
- driver' lics
- driver' license
- driver' licenses
- driver' licence
- driver' licences
- driver'slic
- driver'slics
- driver'slicense
- driver'slicenses
- driver'slicence
- driver'slicences
- driver's lic
- driver's lics
- driver's license
- driver's licenses
- driver's licence
- driver's licences
- dl#
- dls#
- driverlic#
- driverlics#
- driverlicense#
- driverlicenses#
- driverlicence#
- driverlicences#
- driver lic#
- driver lics#
- driver license#
- driver licenses#
- driver licences#
- driverslic#
- driverslics#
- driverslicense#
- driverslicenses#
- driverslicence#
- driverslicences#
- drivers lic#
- drivers lics#
- drivers license#
- drivers licenses#
- drivers licence#
- drivers licences#

- driver'lic#
- driver'lics#
- driver'license#
- driver'licenses#
- driver'licence#
- driver'licences#
- driver' lic#
- driver' lics#
- driver' license#
- driver' licenses#
- driver' licence#
- driver' licences#
- driver'slic#
- driver'slics#
- driver'slicense#
- driver'slicenses#
- driver'slicence#
- driver'slicences#
- driver's lic#
- driver's lics#
- driver's license#
- driver's licenses#
- driver's licence#
- driver's licences#
- driving licence
- driving license
- dlno#
- driv lic
- driv licen
- driv license
- driv licenses
- driv licence
- driv licences
- driver licen
- drivers licen
- driver's licen
- driving lic
- driving licen
- driving licenses
- driving licence
- driving licences
- driving permit
- dl no
- dlno
- dl number

Keywords_spain_eu_driver's_license_number

- permiso de conducción
- permiso conducción
- licencia de conducir
- licencia conducir
- permiso conducir
- permiso de conducir
- permisos de conducir
- permisos conducir
- carnet conducir
- carnet de conducir
- licencia de manejo
- licencia manejo

Spain DNI

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

eight digits followed by one character

Pattern

seven digits followed by one character

- eight digits
- An optional space or hyphen
- one check letter (not case-sensitive)

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_spain_eu_DL_and_NI_number_citizen` or `Func_spain_eu_DL_and_NI_number_foreigner` finds content that matches the pattern.
- A keyword from `Keywords_spain_eu_national_id_card"` is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_spain_eu_DL_and_NI_number_citizen` or `Func_spain_eu_DL_and_NI_number_foreigner` finds content that matches the pattern.

```
<!-- Spain DNI -->
<Entity id="8e6251b9-47b4-40e8-a42b-0f80876be192" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_spain_eu_DL_and_NI_number_citizen" />
    <Match idRef="Keywords_spain_eu_national_id_card" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_spain_eu_DL_and_NI_number_citizen" />
  </Pattern>
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_spain_eu_DL_and_NI_number_foreigner" />
    <Match idRef="Keywords_spain_eu_national_id_card" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_spain_eu_DL_and_NI_number_foreigner" />
  </Pattern>
</Entity>
```

Keywords

Keywords_spain_eu_national_id_card

- carné de identidad
- dni#
- dni
- dninúmero#
- documento nacional de identidad
- identidad único
- identidadúnico#
- insurance number
- national identification number
- national identity
- nationalid#
- nationalidno#
- nie#
- nie
- nienúmero#
- número de identificación
- número nacional identidad
- personal identification number
- personal identity no
- unique identity number
- uniqueid#

Spain passport number

Format

an eight- or nine-character combination of letters and numbers with no spaces or delimiters

Pattern

an eight- or nine-character combination of letters and numbers:

- two digits or letters
- one digit or letter (optional)
- six digits

Checksum

Not applicable

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_spain_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_spain_eu_passport_number` is found.
- The regular expression `Regex_spain_eu_passport_date` finds date in the format DD-MM-YYYY or a keyword from `Keywords_eu_passport_date` is found

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_spain_eu_passport_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_spain_eu_passport_number` is found.

```
<!-- Spain Passport Number -->
<Entity id="d17a57de-9fa5-4e9f-85d3-85c26d89686e" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Regex_spain_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_spain_eu_passport_number" />
    </Any>
    <Any minMatches="1">
      <Match idRef="Regex_spain_eu_passport_date" />
      <Match idRef="Keywords_eu_passport_date" />
    </Any>
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_spain_eu_passport_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_passport_number" />
      <Match idRef="Keywords_spain_eu_passport_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_passport_number

- passport#
- passport #
- passportid
- passports
- passportno
- passport no
- passportnumber
- passport number
- passportnumbers
- passport numbers

Keywords_spain_eu_passport_number

- libreta pasaporte
- número pasaporte

- españa pasaporte
- números de pasaporte
- número de pasaporte
- números pasaporte
- pasaporte no
- Passeport n°
- n° Passeport
- pasaporte no.
- pasaporte n°
- spain passport

Keywords_eu_passport_date

- date of issue
- date of expiry

Spain social security number (SSN)

Format

11-12 digits

Pattern

11-12 digits:

- two digits
- a forward slash (optional)
- seven to eight digits
- a forward slash (optional)
- two digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_spanish_social_security_number finds content that matches the pattern.
- The checksum passes.
- ○ A keyword from `Keywords_spain_eu_ssn_or_equivalent` is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_spanish_social_security_number finds content that matches the pattern.
- The checksum passes.

```
    <!-- Spain SSN -->
    <Entity id="5df987c0-8eae-4bce-ace7-b316347f3070" patternsProximity="300" recommendedConfidence="85"
relaxProximity="true" >
      <Pattern confidenceLevel="85">
          <IdMatch idRef="Func_spanish_social_security_number" />
          <Match idRef="Keywords_spain_eu_ssn_or_equivalent" />
        </Pattern>
        <Pattern confidenceLevel="75">
          <IdMatch idRef="Func_spanish_social_security_number" />
        </Pattern>
    </Entity>
```

Keywords

Keywords_spain_eu_passport_number

- ssn
- ssn#
- socialsecurityno
- social security no
- social security number
- número de la seguridad social

Spain tax identification number

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

seven or eight digits and one or two letters in the specified pattern

Pattern

Spanish Natural Persons with a Spain National Identity Card:

- eight digits
- one uppercase letter (case-sensitive)

Non-resident Spaniards without a Spain National Identity Card

- one uppercase letter "L" (case-sensitive)
- seven digits
- one uppercase letter (case-sensitive)

Resident Spaniards under the age of 14 years without a Spain National Identity Card:

- one uppercase letter "K" (case-sensitive)
- seven digits
- one uppercase letter (case-sensitive)

Foreigners with a Foreigner's Identification Number

- one uppercase letter that is "X", "Y", or "Z" (case-sensitive)

- seven digits
- one uppercase letter (case-sensitive)

Foreigners without a Foreigner's Identification Number

- one uppercase letter that is "M" (case-sensitive)
- seven digits
- one uppercase letter (case-sensitive)

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_spain_eu_tax_file_number` or `Func_spain_eu_DL_and_NI_number_citizen` finds content that matches the pattern.
- A keyword from `Keywords_spain_eu_tax_file_number` is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_spain_eu_tax_file_number` or `Func_spain_eu_DL_and_NI_number_citizen` finds content that matches the pattern.

```
<!-- Spain Tax Identification Number -->
<Entity id="10f0d113-b0e1-47dc-872a-a4f45b9376a3" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_spain_eu_tax_file_number" />
    <Match idRef="Keywords_spain_eu_tax_file_number" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_spain_eu_tax_file_number" />
  </Pattern>
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_spain_eu_DL_and_NI_number_citizen" />
    <Match idRef="Keywords_spain_eu_tax_file_number" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_spain_eu_DL_and_NI_number_citizen" />
  </Pattern>
</Entity>
```

Keywords

Keywords_spain_eu_tax_file_number

- cif
- cifid#
- cifnúmero#
- número de contribuyente
- número de identificación fiscal
- número de impuesto corporativo
- spanishcifid#
- spanishcifid
- spanishcifno#
- spanishcifno

- tax file no
- tax file number
- tax id
- tax identification no
- tax identification number
- tax no#
- tax no
- tax number
- tax registration number
- taxid#
- taxidno#
- taxidnumber#
- taxno#
- taxnumber#
- taxnumber
- tin id
- tin no
- tin#

SQL Server connection string

Format

The string "User Id", "User ID", "uid", or "UserId" followed by the characters and strings outlined in the pattern below.

Pattern

- the string "User Id", "User ID", "uid", or "UserId"
- any combination of between 1-200 lower- or uppercase letters, digits, symbols, special characters, or spaces
- the string "Password" or "pwd" where "pwd" isn't preceded by a lowercase letter
- an equal sign (=)
- any character that isn't a dollar sign ($), percent symbol (%), greater than symbol (>), at symbol (@), quotation mark ("), semicolon (;), left brace([), or left bracket ({)
- any combination of 7-128 characters that are not a semicolon (;), forward slash (/), or quotation mark (")
- a semicolon (;) or quotation mark (")

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression CEP_Regex_SQLServerConnectionString finds content that matches the pattern.
- A keyword from CEP_GlobalFilter isn't found.
- The regular expression CEP_PasswordPlaceHolder doesn't find content that matches the pattern.
- The regular expression CEP_CommonExampleKeywords doesn't find content that matches the pattern.

```
<!---SQL Server Connection String>
<Entity id="e76b6205-d3cb-46f2-bd63-c90153f2f97d" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
        <IdMatch idRef="CEP_Regex_SQLServerConnectionString" />
        <Any minMatches="0" maxMatches="0">
            <Match idRef="CEP_GlobalFilter" />
            <Match idRef="CEP_PasswordPlaceHolder" />
            <Match idRef="CEP_CommonExampleKeywords" />
        </Any>
    </Pattern>
</Entity>
```

Keywords

CEP_GlobalFilter

- some-password
- somepassword
- secretPassword
- sample

CEP_PasswordPlaceHolder

(Note that technically, this sensitive information type identifies these keywords by using a regular expression, not a keyword list.)

- Password or pwd followed by 0-2 spaces, an equal sign (=), 0-2 spaces, and an asterisk (*) -OR-
- Password or pwd followed by:
  - Equal sign (=)
  - Less than symbol (<)
  - Any combination of 1-200 characters that are upper- or lowercase letters, digits, an asterisk (*), hyphen (-), underline (_), or whitespace character
  - Greater than symbol (>)

CEP_CommonExampleKeywords

(Note that technically, this sensitive information type identifies these keywords by using a regular expression, not a keyword list.)

- contoso
- fabrikam
- northwind
- sandbox
- onebox
- localhost
- 127.0.0.1
- testacs.com
- s-int.net

Sweden driver's license number

Format

ten digits containing a hyphen

Pattern

ten digits containing a hyphen:

- six digits

- a hyphen
- four digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression `Regex_sweden_eu_driver's_license_number` finds content that matches the pattern.
- A keyword from `Keywords_eu_driver's_license_number` or `Keywords_sweden_eu_driver's_license_number` is found.

```
<!-- Sweden Driver's License Number -->
<Entity id="70088720-90dd-47f5-805e-5525f3567391" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_sweden_eu_driver's_license_number" />
    <Any minMatches="1">
      <Match idRef="Keywords_eu_driver's_license_number" />
      <Match idRef="Keywords_sweden_eu_driver's_license_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_driver's_license_number

- driverlic
- driverlics
- driverlicense
- driverlicenses
- driverlicence
- driverlicences
- driver lic
- driver lics
- driver license
- driver licenses
- driver licence
- driver licences
- driverslic
- driverslics
- driverslicence
- driverslicences
- driverslicense
- driverslicenses
- drivers lic
- drivers lics
- drivers license
- drivers licenses
- drivers licence
- drivers licences

- driver'lic
- driver'lics
- driver'license
- driver'licenses
- driver'licence
- driver'licences
- driver' lic
- driver' lics
- driver' license
- driver' licenses
- driver' licence
- driver' licences
- driver'slic
- driver'slics
- driver'slicense
- driver'slicenses
- driver'slicence
- driver'slicences
- driver's lic
- driver's lics
- driver's license
- driver's licenses
- driver's licence
- driver's licences
- dl#
- dls#
- driverlic#
- driverlics#
- driverlicense#
- driverlicenses#
- driverlicence#
- driverlicences#
- driver lic#
- driver lics#
- driver license#
- driver licenses#
- driver licences#
- driverslic#
- driverslics#
- driverslicense#
- driverslicenses#
- driverslicence#
- driverslicences#
- drivers lic#
- drivers lics#
- drivers license#

- drivers licenses#
- drivers licence#
- drivers licences#
- driver'lic#
- driver'lics#
- driver'license#
- driver'licenses#
- driver'licence#
- driver'licences#
- driver' lic#
- driver' lics#
- driver' license#
- driver' licenses#
- driver' licence#
- driver' licences#
- driver'slic#
- driver'slics#
- driver'slicense#
- driver'slicenses#
- driver'slicence#
- driver'slicences#
- driver's lic#
- driver's lics#
- driver's license#
- driver's licenses#
- driver's licence#
- driver's licences#
- driving licence
- driving license
- dlno#
- driv lic
- driv licen
- driv license
- driv licenses
- driv licence
- driv licences
- driver licen
- drivers licen
- driver's licen
- driving lic
- driving licen
- driving licenses
- driving licence
- driving licences
- driving permit
- dl no

- dlno
- dl number

Keywords_sweden_eu_driver's_license_number

- ajokortti
- permis de conducere
- ajokortin numero
- kuljettajat lic.
- drivere lic.
- körkort
- numărul permisului de conducere
- שאָפער דערלויבעניש נומער
- förare lic.
- דריווערס דערלויבעניש
- körkortsnummer

Sweden national ID

Format

10 or 12 digits and an optional delimiter

Pattern

10 or 12 digits and an optional delimiter:

- two digits (optional)
- Six digits in date format YYMMDD
- delimiter of "-" or "+" (optional)
- four digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_swedish_national_identifier` finds content that matches the pattern.
- A keyword from `Keywords_swedish_national_identifier` is found
- The checksum passes.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_swedish_national_identifier` finds content that matches the pattern.
- The checksum passes.

```
<!-- Sweden National ID -->
<Entity id="f69aaf40-79be-4fac-8f05-fd1910d272c8" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_swedish_national_identifier" />
    <Match idRef="Keywords_swedish_national_identifier" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_swedish_national_identifier" />
  </Pattern>
</Entity>
```

Keywords

Keywords_swedish_national_identifier

- id no
- id number
- id#
- identification no
- identification number
- identifikationsnumret#
- identifikationsnumret
- identitetshandling
- identity document
- identity no
- identity number
- id-nummer
- personal id
- personnummer#
- personnummer
- skatteidentifikationsnummer

Sweden passport number

Format

eight digits

Pattern

eight consecutive digits

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- the regular expression Regex_sweden_passport_number finds content that matches the pattern.
- a keyword from `Keywords_eu_passport_number` or `Keyword_sweden_passport` is found.
- the regular expression `Regex_sweden_eu_passport_date` finds a date in the format DD MMM/MMM YY (01 JAN/JAN 12) or a keyword from `Keywords_eu_passport_date` is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- the regular expression Regex_sweden_passport_number finds content that matches the pattern.
- a keyword from `Keywords_eu_passport_number` or `Keyword_sweden_passport` is found.

```
<!-- Sweden Passport Number -->
<Entity id="ba4e7456-55a9-4d89-9140-c33673553526" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
      <IdMatch idRef="Regex_sweden_passport_number" />
      <Any minMatches="1">
        <Match idRef="Keywords_eu_passport_number" />
        <Match idRef="Keyword_sweden_passport" />
      </Any>
      <Any minMatches="1">
        <Match idRef="Regex_sweden_eu_passport_date" />
        <Match idRef="Keywords_eu_passport_date" />
      </Any>
    </Pattern>
    <Pattern confidenceLevel="75">
      <IdMatch idRef="Regex_sweden_passport_number" />
      <Any minMatches="1">
        <Match idRef="Keywords_eu_passport_number" />
        <Match idRef="Keyword_sweden_passport" />
      </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_eu_passport_number

- passport#
- passport #
- passportid
- passports
- passportno
- passport no
- passportnumber
- passport number
- passportnumbers
- passport numbers

Keyword_sweden_passport

- alien registration card
- g3 processing fees
- multiple entry
- Numéro de passeport
- passeport n °
- passeport non
- passeport #
- passeport#
- passeportnon
- passeportn °
- passnummer
- pass nr
- schengen visa
- schengen visas
- single entry

- sverige pass
- visa requirements
- visa processing
- visa type

Keywords_eu_passport_date

- date of issue
- date of expiry

Sweden social security number or equivalent identification

This sensitive information type entity is only available in the EU Social Security Number or Equivalent ID sensitive information type.

Format

12 digits without spaces and delimiters

Pattern

12 digits:

- eight digits that correspond to the birth date (YYYYMMDD)
- three digits that correspond to a serial number where:
  - the last digit in the serial number indicates gender by the assignment of an odd number for male and an even number for female
  - Before 1990, the assignment of a serial number corresponded to the county where the bearer of the number was born. Or (if born before 1947) where they had been living, according to tax records, on January 1, 1947, with a special code (usually 9 as the seventh digit) for immigrants.
- one check digit

Checksum

Yes

Definition

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_sweden_eu_ssn_or_equivalent` finds content that matches the pattern.
- A keyword from `Keywords_sweden_eu_ssn_or_equivalent` is found.

A DLP policy is 75% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_sweden_eu_ssn_or_equivalent` finds content that matches the pattern.

```
 <!-- EU SSN or Equivalent Number -->
<Entity id="d24e32a4-c0bb-4ba8-899d-6303b95742d9" patternsProximity="300" recommendedConfidence="75">
        <Pattern confidenceLevel="85">
          <IdMatch idRef="Func_sweden_eu_ssn_or_equivalent" />
          <Match idRef="Keywords_sweden_eu_ssn_or_equivalent" />
        </Pattern>
       <Pattern confidenceLevel="75">
          <IdMatch idRef="Func_sweden_eu_ssn_or_equivalent" />
        </Pattern>
</Entity>
```

Keywords

Keywords_sweden_eu_ssn_or_equivalent

- personal id number
- identification number
- personal id no
- identity no
- identification no
- personal identification no
- personnummer id
- personligt id-nummer
- unikt id-nummer
- personnummer
- identifikationsnumret
- personnummer#
- identifikationsnumret#

Sweden tax identification number

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

10 digits and a symbol in the specified pattern

Pattern

10 digits and a symbol:

- six digits that correspond to the birth date (YYMMDD)
- a plus sign or minus sign
- three digits that make the identification number unique where:
  - for numbers issued before 1990, the seventh and eighth digit identify the county of birth or foreign-born people
  - the digit in the ninth position indicates gender by either odd for male or even for female
- one check digit

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_sweden_eu_tax_file_number` finds content that matches the pattern.
- A keyword from `Keywords_sweden_eu_tax_file_number` is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_sweden_eu_tax_file_number` finds content that matches the pattern.

```
<!-- Sweden Tax Identification Number -->
<Entity id="139acba0-a5bc-4fbb-876d-f7a493ae8a40" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_sweden_eu_tax_file_number" />
    <Match idRef="Keywords_sweden_eu_tax_file_number" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_sweden_eu_tax_file_number" />
    <Any minMatches="0" maxMatches="0">
      <Match idRef="Keywords_sweden_eu_telephone_number" />
      <Match idRef="Keywords_sweden_eu_mobile_number" />
    </Any>
  </Pattern>
</Entity>
```

Keywords

Keywords_sweden_eu_tax_file_number

- personal id number
- personnummer
- skatt id nummer
- skatt identifikation
- skattebetalarens identifikationsnummer
- sverige tin
- tax file
- tax id
- tax identification no
- tax identification number
- tax no#
- tax no
- tax number
- tax registration number
- taxid#
- taxidno#
- taxidnumber#
- taxno#
- taxnumber#
- taxnumber
- tin id
- tin no
- tin#

SWIFT code

Format

four letters followed by 5-31 letters or digits

Pattern

four letters followed by 5-31 letters or digits:

- four-letter bank code (not case-sensitive)

- an optional space
- 4-28 letters or digits (the Basic Bank Account Number (BBAN))
- an optional space
- one to three letters or digits (remainder of the BBAN)

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_swift finds content that matches the pattern.
- A keyword from Keyword_swift is found.

```
<Entity id="cb2ab58c-9cb8-4c81-baf8-a4e106791df4" patternsProximity="300" recommendedConfidence="75">
<Pattern confidenceLevel="75">
        <IdMatch idRef="Regex_swift" />
        <Match idRef="Keyword_swift" />
    </Pattern>
</Entity>
```

Keywords

Keyword_swift

- international organization for standardization 9362
- iso 9362
- iso9362
- swift#
- swiftcode
- swiftnumber
- swiftroutingnumber
- swift code
- swift number #
- swift routing number
- bic number
- bic code
- bic #
- bic#
- bank identifier code
- Organisation internationale de normalisation 9362
- rapide #
- code SWIFT
- le numéro de swift
- swift numéro d'acheminement
- le numéro BIC
- # BIC
- code identificateur de banque
- SWIFTコード
- SWIFT番号
- BIC番号

- BICコード
- SWIFT コード
- SWIFT 番号
- BIC 番号
- BIC コード
- 金融機関識別コード
- 金融機関コード
- 銀行コード

Switzerland SSN AHV number

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

13-digit number

Pattern

13-digit number:

- three digits - 756
- an optional dot
- four digits
- an optional dot
- four digits
- an optional dot
- two digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_swiss_social_security_number_ahv finds content that matches the pattern.
- A keyword from Keywords_swiss_social_security_number_ahv is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_swiss_social_security_number_ahv finds content that matches the pattern.

```
<!-- Swiss SSN AHV Number -->
<Entity id="277cfa4b-6eaa-4a1b-9492-099dec849971" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_swiss_social_security_number_ahv" />
    <Match idRef="Keywords_swiss_social_security_number_ahv" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_swiss_social_security_number_ahv" />
  </Pattern>
</Entity>
```

Keywords

Keyword_swiss_ssn_AHV_number

- ahv
- ssn
- pid
- insurance number
- personalidno#
- social security number
- personal id number
- personal identification no.
- insuranceno#
- uniqueidno#
- unique identification no.
- avs number
- personal identity no versicherungsnummer
- identifikationsnummer
- einzigartige identität nicht
- sozialversicherungsnummer
- identification personnelle id
- numéro de sécurité sociale

Taiwan national identification number

Format

one letter (in English) followed by nine digits

Pattern

one letter (in English) followed by nine digits:

- one letter (in English, not case-sensitive)
- the digit "1" or "2"
- eight digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_taiwanese_national_id finds content that matches the pattern.

- A keyword from Keyword_taiwanese_national_id is found.
- The checksum passes.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_taiwanese_national_id finds content that matches the pattern.
- The checksum passes.

```
<!-- Taiwanese National ID -->
<Entity id="4C7BFC34-8DD1-421D-8FB7-6C6182C2AF03" patternsProximity="300" recommendedConfidence="85">
      <Pattern confidenceLevel="85">
          <IdMatch idRef="Func_taiwanese_national_id" />
          <Match idRef="Keyword_taiwanese_national_id" />
      </Pattern>
       <Pattern confidenceLevel="75">
         <IdMatch idRef="Func_taiwanese_national_id" />
       </Pattern>
</Entity>
```

Keywords

Keyword_taiwan_national_id

- 身份證字號
- 身份證
- 身份證號碼
- 身份證號
- 身分證字號
- 身分證
- 身分證號碼
- 身份證號
- 身分證統一編號
- 國民身分證統一編號
- 簽名
- 蓋章
- 簽名或蓋章
- 簽章

Taiwan passport number

Format

- biometric passport number: nine digits
- non-biometric passport number: nine digits

Pattern

biometric passport number:

- the character "3"
- eight digits

non-biometric passport number:

- nine digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_taiwan_passport finds content that matches the pattern.
- A keyword from Keyword_taiwan_passport is found.

```
<!-- Taiwan Passport Number -->
<Entity id="e7251cb4-4c2c-41df-963e-924eb3dae04a" recommendedConfidence="75" patternsProximity="300">
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Regex_taiwan_passport"/>
     <Match idRef="Keyword_taiwan_passport"/>
  </Pattern>
</Entity>
```

Keywords

Keyword_taiwan_passport

- ROC passport number
- Passport number
- Passport no
- Passport Num
- Passport #
- 护照
- 中華民國護照
- Zhōnghuá Mínguó hùzhào

Taiwan-resident certificate (ARC/TARC) number

Format

10 letters and digits

Pattern

10 letters and digits:

- two letters (not case-sensitive)
- eight digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_taiwan_resident_certificate finds content that matches the pattern.
- A keyword from Keyword_taiwan_resident_certificate is found.

```
<!-- Taiwan Resident Certificate (ARC/TARC) -->
<Entity id="48269fec-05ea-46ea-b326-f5623a58c6e9" recommendedConfidence="75" patternsProximity="300">
  <Pattern confidenceLevel="75">
     <IdMatch idRef="Regex_taiwan_resident_certificate"/>
     <Match idRef="Keyword_taiwan_resident_certificate"/>
  </Pattern>
</Entity>
```

Keywords

Keyword_taiwan_resident_certificate

- Resident Certificate
- Resident Cert
- Resident Cert.
- Identification card
- Alien Resident Certificate
- ARC
- Taiwan Area Resident Certificate
- TARC
- 居留證
- 外僑居留證
- 台灣地區居留證

Thai population identification code

Format

13 digits

Pattern

13 digits:

- first digit isn't zero or nine
- 12 digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_Thai_Citizen_Id finds content that matches the pattern.
- A keyword from Keyword_Thai_Citizen_Id is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_Thai_Citizen_Id finds content that matches the pattern.

```
<!-- Thai Citizen ID -->
-<Entity id="44ca9e86-ead7-4c5d-884a-e2eaa401515e" recommendedConfidence="75" patternsProximity="300">
   -<Pattern confidenceLevel="85">
      <IdMatch idRef="Func_Thai_Citizen_Id"/>
      <Match idRef="Keyword_Thai_Citizen_Id"/>
   </Pattern>
   -<Pattern confidenceLevel="75">
      <IdMatch idRef="Func_Thai_Citizen_Id"/>
   </Pattern>
</Entity>
```

Keywords

Keyword_thai_citizen_Id

- ID Number
- Identification Number
- **บัตรประชาชน**
- **รหัสบัตรประชาชน**
- **บัตรประชาชน**
- **รหัสบัตรประชาชน**

Turkish national identification number

Format

11 digits

Pattern

11 digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_Turkish_National_Id finds content that matches the pattern.
- A keyword from Keyword_Turkish_National_Id is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_Turkish_National_Id finds content that matches the pattern.

```
<!-- Turkish National Identity -->
-<Entity id="fb621f20-3876-4cfc-acec-8c8e73ca32c7" recommendedConfidence="75" patternsProximity="300">
   -<Pattern confidenceLevel="85">
      <IdMatch idRef="Func_Turkish_National_Id"/>
      <Match idRef="Keyword_Turkish_National_Id"/>
   </Pattern>
   -<Pattern confidenceLevel="75">
      <IdMatch idRef="Func_Turkish_National_Id"/>
   </Pattern>
</Entity>
```

Keywords

Keyword_turkish_national_id

- TC Kimlik No
- TC Kimlik numarası
- Vatandaşlık numarası
- Vatandaşlık no

U.K. driver's license number

Format

Combination of 18 letters and digits in the specified format

Pattern

18 letters and digits:

- Five letters (not case-sensitive) or the digit "9" in place of a letter.
- One digit.
- Five digits in the date format MMDDY for date of birth. The seventh character is incremented by 50 if driver is female; for exampe, 51 to 62 instead of 01 to 12.
- Two letters (not case-sensitive) or the digit "9" in place of a letter.
- Five digits.

Checksum

Yes

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_uk_drivers_license` finds content that matches the pattern.
- A keyword from `Keywords_eu_driver's_license_number` is found.
- The checksum passes.

A DLP policy has low confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_uk_drivers_license` finds content that matches the pattern.
- The checksum passes.

```
    <!-- U.K. Driver's License Number -->
    <Entity id="f93de4be-d94c-40df-a8be-461738047551" patternsProximity="300" recommendedConfidence="75"
relaxProximity="true" >
      <Pattern confidenceLevel="75">
          <IdMatch idRef="Func_uk_drivers_license" />
          <Match idRef="Keywords_eu_driver's_license_number" />
        </Pattern>
        <Pattern confidenceLevel="65">
          <IdMatch idRef="Func_uk_drivers_license" />
        </Pattern>
    </Entity>
```

Keywords

Keywords_eu_driver's_license_number

- driverlic
- driverlics
- driverlicense

- driverlicenses
- driverlicence
- driverlicences
- driver lic
- driver lics
- driver license
- driver licenses
- driver licence
- driver licences
- driverslic
- driverslics
- driverslicence
- driverslicences
- driverslicense
- driverslicenses
- drivers lic
- drivers lics
- drivers license
- drivers licenses
- drivers licence
- drivers licences
- driver'lic
- driver'lics
- driver'license
- driver'licenses
- driver'licence
- driver'licences
- driver' lic
- driver' lics
- driver' license
- driver' licenses
- driver' licence
- driver' licences
- driver'slic
- driver'slics
- driver'slicense
- driver'slicenses
- driver'slicence
- driver'slicences
- driver's lic
- driver's lics
- driver's license
- driver's licenses
- driver's licence
- driver's licences
- dl#

- dls#
- driverlic#
- driverlics#
- driverlicense#
- driverlicenses#
- driverlicence#
- driverlicences#
- driver lic#
- driver lics#
- driver license#
- driver licenses#
- driver licences#
- driverslic#
- driverslics#
- driverslicense#
- driverslicenses#
- driverslicence#
- driverslicences#
- drivers lic#
- drivers lics#
- drivers license#
- drivers licenses#
- drivers licence#
- drivers licences#
- driver'lic#
- driver'lics#
- driver'license#
- driver'licenses#
- driver'licence#
- driver'licences#
- driver' lic#
- driver' lics#
- driver' license#
- driver' licenses#
- driver' licence#
- driver' licences#
- driver'slic#
- driver'slics#
- driver'slicense#
- driver'slicenses#
- driver'slicence#
- driver'slicences#
- driver's lic#
- driver's lics#
- driver's license#
- driver's licenses#

- driver's licence#
- driver's licences#
- driving licence
- driving license
- dlno#
- driv lic
- driv licen
- driv license
- driv licenses
- driv licence
- driv licences
- driver licen
- drivers licen
- driver's licen
- driving lic
- driving licen
- driving licenses
- driving licence
- driving licences
- driving permit
- dl no
- dlno
- dl number

U.K. electoral roll number

Format

two letters followed by 1-4 digits

Pattern

two letters (not case-sensitive) followed by 1-4 numbers

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_uk_electoral finds content that matches the pattern.
- A keyword from Keyword_uk_electoral is found.

```
<!-- U.K. Electoral Number -->
<Entity id="a3eea206-dc0c-4f06-9e22-aa1be3059963" patternsProximity="300" recommendedConfidence="75">
    <Pattern confidenceLevel="75">
        <IdMatch idRef="Regex_uk_electoral" />
        <Any minMatches="1">
          <Match idRef="Keyword_uk_electoral" />
        </Any>
    </Pattern>
</Entity>
```

Keywords

Keyword_uk_electoral

- council nomination
- nomination form
- electoral register
- electoral roll

U.K. national health service number

Format

10-17 digits separated by spaces

Pattern

10-17 digits:

- either 3 or 10 digits
- a space
- three digits
- a space
- four digits

Checksum

Yes

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_uk_nhs_number finds content that matches the pattern.
- One of the following is true:
  - A keyword from Keyword_uk_nhs_number is found.
  - A keyword from Keyword_uk_nhs_number1 is found.
  - A keyword from Keyword_uk_nhs_number_dob is found.
- The checksum passes.

```
<!-- U.K. NHS Number -->
<Entity id="3192014e-2a16-44e9-aa69-4b20375c9a78" patternsProximity="300" recommendedConfidence="85">
    <Pattern confidenceLevel="85">
        <IdMatch idRef="Func_uk_nhs_number" />
        <Any minMatches="1">
          <Match idRef="Keyword_uk_nhs_number" />
          <Match idRef="Keyword_uk_nhs_number1" />
          <Match idRef="Keyword_uk_nhs_number_dob" />
        </Any>
    </Pattern>
</Entity>
```

Keywords

Keyword_uk_nhs_number

- national health service
- nhs
- health services authority
- health authority

Keyword_uk_nhs_number1

- patient id
- patient identification
- patient no
- patient number

Keyword_uk_nhs_number_dob

- GP
- DOB
- D.O.B
- Date of Birth
- Birth Date

U.K. national insurance number (NINO)

This sensitive information type entity is included in the EU National Identification Number sensitive information type. It's available as a stand-alone sensitive information type entity.

Format

seven characters or nine characters separated by spaces or dashes

Pattern

two possible patterns:

- two letters (valid NINOs use only certain characters in this prefix, which this pattern validates; not case-sensitive)
- six digits
- either 'A', 'B', 'C', or 'D' (like the prefix, only certain characters are allowed in the suffix; not case-sensitive)

OR

- two letters
- a space or dash
- two digits
- a space or dash
- two digits
- a space or dash
- two digits
- a space or dash
- either 'A', 'B', 'C', or 'D'

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_uk_nino finds content that matches the pattern.
- A keyword from Keyword_uk_nino is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_uk_nino finds content that matches the pattern.

```
    <!-- U.K. NINO -->
    <Entity id="16c07343-c26f-49d2-a987-3daf717e94cc" patternsProximity="300" recommendedConfidence="75"
relaxProximity="true">
      <Pattern confidenceLevel="85">
        <IdMatch idRef="Func_uk_nino" />
        <Match idRef="Keyword_uk_nino" />
      </Pattern>
      <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_uk_nino" />
      </Pattern>
    </Entity>
```

Keywords

Keyword_uk_nino

- national insurance number
- national insurance contributions
- protection act
- insurance
- social security number
- insurance application
- medical application
- social insurance
- medical attention
- social security
- great britain
- NI Number
- NI No.
- NI #
- NI#
- insurance#
- insurancenumber
- nationalinsurance#
- nationalinsurancenumber

U.K. Unique Taxpayer Reference Number

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

10 digits without spaces and delimiters

Pattern

10 digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_uk_eu_tax_file_number` finds content that matches the pattern.
- A keyword from `Keywords_uk_eu_tax_file_number` is found.

```
<!-- U.K. Unique Taxpayer Reference Number -->
<Entity id="ad4a8116-0db8-439a-b545-6d967642f0ec" patternsProximity="300" recommendedConfidence="85">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_uk_eu_tax_file_number" />
    <Match idRef="Keywords_uk_eu_tax_file_number" />
  </Pattern>
</Entity>
```

Keywords

Keywords_uk_eu_tax_file_number

- tax number
- tax file
- tax id
- tax identification no
- tax identification number
- tax no#
- tax no
- tax registration number
- taxid#
- taxidno#
- taxidnumber#
- taxno#
- taxnumber#
- taxnumber
- tin id
- tin no
- tin#

U.S. bank account number

Format

6-17 digits

Pattern

6-17 consecutive digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regular expression Regex_usa_bank_account_number finds content that matches the pattern.
- A keyword from Keyword_usa_Bank_Account is found.

```
<!-- U.S. Bank Account Number -->
<Entity id="a2ce32a8-f935-4bb6-8e96-2a5157672e2c" patternsProximity="300" recommendedConfidence="75">
    <Pattern confidenceLevel="75">
        <IdMatch idRef="Regex_usa_bank_account_number" />
        <Match idRef="Keyword_usa_Bank_Account" />
    </Pattern>
</Entity>
```

Keywords

Keyword_usa_Bank_Account

- Checking Account Number
- Checking Account
- Checking Account #
- Checking Acct Number
- Checking Acct #
- Checking Acct No.
- Checking Account No.
- Bank Account Number
- Bank Account #
- Bank Acct Number
- Bank Acct #
- Bank Acct No.
- Bank Account No.
- Savings Account Number
- Savings Account.
- Savings Account #
- Savings Acct Number
- Savings Acct #
- Savings Acct No.
- Savings Account No.
- Debit Account Number
- Debit Account
- Debit Account #
- Debit Acct Number
- Debit Acct #
- Debit Acct No.
- Debit Account No.

U.S. driver's license number

Format

Depends on the state

Pattern

depends on the state - for example, New York:

- nine digits formatted like ddd ddd ddd will match.

- nine digits like ddddddddd will not match.

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_new_york_drivers_license_number finds content that matches the pattern.
- A keyword from Keyword_[state_name]_drivers_license_name is found.
- A keyword from Keyword_us_drivers_license is found.

A DLP policy has low confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_new_york_drivers_license_number finds content that matches the pattern.
- A keyword from Keyword_[state_name]_drivers_license_name is found.
- A keyword from Keyword_us_drivers_license_abbreviations is found.
- No keyword from Keyword_us_drivers_license is found.

```
<Entity id="dfeb356f-61cd-459e-bf0f-7c6d28b458c6 patternsProximity="300">
    <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_new_york_drivers_license_number" />
        <Match idRef="Keyword_new_york_drivers_license_name" />
        <Match idRef="Keyword_us_drivers_license" />
    </Pattern>
    <Pattern confidenceLevel="65">
        <IdMatch idRef="Func_new_york_drivers_license_number" />
        <Match idRef="Keyword_new_york_drivers_license_name" />
        <Match idRef="Keyword_us_drivers_license_abbreviations" />
        <Any minMatches="0" maxMatches="0">
          <Match idRef="Keyword_us_drivers_license" />
        </Any>
    </Pattern>
</Entity>
```

Keywords

Keyword_us_drivers_license_abbreviations

- DL
- DLS
- CDL
- CDLS
- ID
- IDs
- DL#
- DLS#
- CDL#
- CDLS#
- ID#
- IDs#
- ID number
- ID numbers
- LIC

- LIC#

Keyword_us_drivers_license

- DriverLic
- DriverLics
- DriverLicense
- DriverLicenses
- Driver Lic
- Driver Lics
- Driver License
- Driver Licenses
- DriversLic
- DriversLics
- DriversLicense
- DriversLicenses
- Drivers Lic
- Drivers Lics
- Drivers License
- Drivers Licenses
- Driver'Lic
- Driver'Lics
- Driver'License
- Driver'Licenses
- Driver' Lic
- Driver' Lics
- Driver' License
- Driver' Licenses
- Driver'sLic
- Driver'sLics
- Driver'sLicense
- Driver'sLicenses
- Driver's Lic
- Driver's Lics
- Driver's License
- Driver's Licenses
- identification number
- identification numbers
- identification #
- id card
- id cards
- identification card
- identification cards
- DriverLic#
- DriverLics#
- DriverLicense#
- DriverLicenses#
- Driver Lic#

- Driver Lics#
- Driver License#
- Driver Licenses#
- DriversLic#
- DriversLics#
- DriversLicense#
- DriversLicenses#
- Drivers Lic#
- Drivers Lics#
- Drivers License#
- Drivers Licenses#
- Driver'Lic#
- Driver'Lics#
- Driver'License#
- Driver'Licenses#
- Driver' Lic#
- Driver' Lics#
- Driver' License#
- Driver' Licenses#
- Driver'sLic#
- Driver'sLics#
- Driver'sLicense#
- Driver'sLicenses#
- Driver's Lic#
- Driver's Lics#
- Driver's License#
- Driver's Licenses#
- id card#
- id cards#
- identification card#
- identification cards#

Keyword_[state_name]_drivers_license_name

- state abbreviation (for example, "NY")
- state name (for example, "New York")

U.S. individual taxpayer identification number (ITIN)

Format

nine digits that start with a "9" and contain a "7" or "8" as the fourth digit, optionally formatted with spaces or dashes

Pattern

formatted:

- the digit "9"
- two digits
- a space or dash
- a "7" or "8"

- a digit
- a space, or dash
- four digits

unformatted:

- the digit "9"
- two digits
- a "7" or "8"
- five digits

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_formatted_itin finds content that matches the pattern.
- A keyword from Keyword_itin is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_unformatted_itin finds content that matches the pattern.
- A keyword from Keyword_itin is found.

A DLP policy has low confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_formatted_itin or Func_unformatted_itin finds content that matches the pattern.

```
<!-- U.S. Individual Taxpayer Identification Number (ITIN) -->
<Entity id="e55e2a32-f92d-4985-a35d-a0b269eb687b" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="85">
    <IdMatch idRef="Func_formatted_itin" />
    <Match idRef="Keyword_itin" />
  </Pattern>
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Func_unformatted_itin" />
    <Match idRef="Keyword_itin" />
  </Pattern>
  <Pattern confidenceLevel="65">
    <IdMatch idRef="Func_formatted_itin" />
  </Pattern>
  <Pattern confidenceLevel="65">
    <IdMatch idRef="Func_unformatted_itin" />
  </Pattern>
</Entity>
```

Keywords

Keyword_itin

- taxpayer
- tax id
- tax identification
- itin
- i.t.i.n.

- ssn
- tin
- social security
- tax payer
- itins
- taxid
- individual taxpayer

U.S. social security number (SSN)

Format

nine digits, which may be in a formatted or unformatted pattern

> **NOTE**
> If issued before mid-2011, an SSN has strong formatting where certain parts of the number must fall within certain ranges to be valid (but there's no checksum).

Pattern

four functions look for SSNs in four different patterns:

- Func_ssn finds SSNs with pre-2011 strong formatting that are formatted with dashes or spaces (ddd-dd-dddd OR ddd dd dddd)
- Func_unformatted_ssn finds SSNs with pre-2011 strong formatting that are unformatted as nine consecutive digits (ddddddddd)
- Func_randomized_formatted_ssn finds post-2011 SSNs that are formatted with dashes or spaces (ddd-dd-dddd OR ddd dd dddd)
- Func_randomized_unformatted_ssn finds post-2011 SSNs that are unformatted as nine consecutive digits (ddddddddd)

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_ssn finds content that matches the pattern.
- A keyword from Keyword_ssn is found.

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_unformatted_ssn finds content that matches the pattern.
- A keyword from Keyword_ssn is found.

A DLP policy has low confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_randomized_formatted_ssn finds content that matches the pattern.
- A keyword from Keyword_ssn is found.

A DLP policy is 55% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_randomized_unformatted_ssn finds content that matches the pattern.
- A keyword from Keyword_ssn is found.

```
<!-- U.S. Social Security Number (SSN) -->
  <Entity id="a44669fe-0d48-453d-a9b1-2cc83f2cba77" patternsProximity="300" recommendedConfidence="75">
      <Pattern confidenceLevel="85">
        <IdMatch idRef="Func_ssn" />
        <Match idRef="Keyword_ssn" />
      </Pattern>
      <Pattern confidenceLevel="75">
        <IdMatch idRef="Func_unformatted_ssn" />
        <Match idRef="Keyword_ssn" />
      </Pattern>
      <Pattern confidenceLevel="65">
        <IdMatch idRef="Func_randomized_formatted_ssn" />
        <Match idRef="Keyword_ssn" />
      </Pattern>
      <Pattern confidenceLevel="55">
        <IdMatch idRef="Func_randomized_unformatted_ssn" />
        <Match idRef="Keyword_ssn" />
      </Pattern>
  </Entity>
```

Keywords

Keyword_ssn

- SSA Number
- social security number
- social security #
- social security#
- social security no
- Social Security#
- Soc Sec
- SSN
- SSNS
- SSN#
- SS#
- SSID

U.S. / U.K. passport number

Format

nine digits

Pattern

nine consecutive digits

Checksum

No

Definition

A DLP policy has high confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_usa_uk_passport finds content that matches the pattern.

- A keyword from `Keywords_eu_passport_number` or `Keywords_uk_eu_passport_number` is found.
- A keyword from `Keywords_eu_passport_date` is found

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function Func_usa_uk_passport finds content that matches the pattern.
- A keyword from `Keywords_eu_passport_number` or `Keywords_uk_eu_passport_number` is found.

```
<!-- U.S. / U.K. Passport Number -->
<Entity id="178ec42a-18b4-47cc-85c7-d62c92fd67f8" patternsProximity="300" recommendedConfidence="75">
   <Pattern confidenceLevel="85">
      <IdMatch idRef="Func_usa_uk_passport" />
      <Match idRef="Keywords_eu_passport_date" />
      <Any minMatches="1">
        <Match idRef="Keywords_eu_passport_number" />
        <Match idRef="Keywords_uk_eu_passport_number" />
      </Any>
    </Pattern>
    <Pattern confidenceLevel="75">
      <IdMatch idRef="Func_usa_uk_passport" />
      <Any minMatches="1">
        <Match idRef="Keywords_eu_passport_number" />
        <Match idRef="Keywords_uk_eu_passport_number" />
      </Any>
    </Pattern>
</Entity>
```

Keywords

Keywords_eu_passport_number

- passport#
- passport #
- passportid
- passports
- passportno
- passport no
- passportnumber
- passport number
- passportnumbers
- passport numbers

Keywords_uk_eu_passport_number

- british passport
- uk passport

Ukraine passport domestic

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

nine digits

Pattern

nine digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regex Regex_Ukraine_Passport_Domestic finds content that matches the pattern.
- A keyword from Keyword_Ukraine_Passport_Domestic is found.

```
<!-- Ukraine Passport Domestic -->
<Entity id="1817a540-221f-4459-9202-3bd78b81d803" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
  <IdMatch idRef="Regex_Ukraine_Passport_Domestic"/>
    <Match idRef="Keyword_Ukraine_Passport_Domestic"/>
  </Pattern>
</Entity>
```

Keywords

Keyword_ukraine_passport_domestic

- ukraine passport
- passport number
- passport no
- паспорт України
- номер паспорта
- персональний

Ukraine passport international

This sensitive information type is only available for use in:

- data loss prevention policies
- communication compliance policies
- information governance
- records management
- Microsoft cloud app security

Format

eight-character alphanumeric pattern

Pattern

eight-character alphanumeric pattern:

- two letters or digits
- six digits

Checksum

No

Definition

A DLP policy has medium confidence that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The regex Regex_Ukraine_Passport_International finds content that matches the pattern.
- A keyword from Keyword_Ukraine_Passport_International is found.

```
<!-- Ukraine Passport International -->
<Entity id="cfbe032d-22e0-4f28-ab68-d66e9641f1e2" patternsProximity="300" recommendedConfidence="75">
  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_Ukraine_Passport_International"/>
    <Match idRef="Keyword_Ukraine_Passport_International"/>
  </Pattern>
</Entity>
```

Keywords

Keyword_ukraine_passport_international

- ukraine passport
- passport number
- passport no
- паспорт України
- номер паспорта

What the DLP functions look for

11/2/2020 • 10 minutes to read • Edit Online

Data loss prevention (DLP) policies can use sensitive information types to identify sensitive items. Credit card number and EU debit card number are examples of sensitive information types. Sensitive information types look for specific patterns. Sensitive information types validate the data by looking at it's format, it's checksums, and looks for relevant keywords or other information. Some of this functionality is performed by internal functions. For example, the Credit Card Number sensitive information type uses a function to look for dates that are formatted like an expiration date. This helps to corroborate that a number is a credit card number.

This article explains what these functions look for, to help you understand how the predefined sensitive information types work. For more information, see Sensitive information type entity definitions

Table of functions

FUNCTION NAME	FUNCTION ACTION	IS A VALIDATOR	
Func_Argentina_Unique_Tax_Key	detects and validates Argentina Unique tax key	no	
Func_aba_routing	detects ABA routing number	yes	
Func_alabama_drivers_license_number	detects Alabama driver's license number	no	
Func_alaska_delaware_oregon_drivers_license_number	detects Alaska, Delaware, Oregon driver's license number	no	
Func_alaska_drivers_license_number	detects Alaska driver's license number	no	
Func_alberta_drivers_license_number	detects Alberta driver's license number	no	
Func_Argentina_Unique_Tax_Key	detects Argentina Unique tax key	no	
Func_arizona_drivers_license_number	detects Arizona driver's license number	no	
Func_arkansas_drivers_license_number	detects Arkansas driver's license number	no	
Func_australian_business_number	detects Australia business number	no	
Func_Australian_Company_Number	detects Australia company number	no	

FUNCTION NAME	FUNCTION ACTION	IS A VALIDATOR	
Func_australian_medical_account_number	detects Australia medical account number	no	
Func_australian_tax_file_number	detects Australia tax file number	yes	
Func_austria_eu_ssn_or_equivalent	detects Austria social security number	no	
Func_austria_eu_tax_file_number	detects Austria tax file number	no	
Func_Austria_Value_Added_Tax	detects Austria Value Added Tax	no	
Func_belgium_national_number	detects Belgium national number	no	
Func_belgium_value_added_tax_number	detects Belgium value added tax number	no	
Func_brazil_cnpj	detects Brazil legal entity number (CNPJ)	yes	
Func_brazil_cpf	detects Brazil CPF	yes	
Func_brazil_rg	detects Brazil RG	no	
Func_british_columbia_drivers_license_number	detects British Columbia driver's license number	no	
Func_bulgaria_eu_national_id_card	detects Bulgaria uniform civil number	no	
Func_california_drivers_license_number	detects California driver's license number	no	
Func_canadian_sin	detects Canada sin	yes	
Func_chile_id_card	detects Chile ID card	no	
Func_china_resident_id	detects China-resident ID	no	
Func_colorado_drivers_license_number	detects Colorado driver's license number	no	
Func_connecticut_drivers_license_number	detects Connecticut driver's license number	no	
Func_credit_card	detects credit card	yes	no
Func_croatia_id_card	detects Croatia ID card	no	

FUNCTION NAME	FUNCTION ACTION	IS A VALIDATOR	
Func_croatia_oib_number	detects Croatia OIB number	no	
Func_cyprus_eu_tax_file_number	detects Cyprus tax file number	no	
Func_czech_id_card	detects Czech ID card	no	
Func_czech_id_card_new_format	detects Czech ID card in new format	no	
Func_dea_number	detects DEA number	yes	
Func_denmark_eu_tax_file_number	detects Denmark personal identification number	no	
Func_district_of_columbia_drivers_license_number	detects District of Columbia driver's license number	no	
Func_estonia_eu_national_id_card	detects Estonia Personal Identification Code	no	
Func_eu_debit_card	detects EU debit card	no	
Func_finnish_national_id	detects Finnish national ID	no	
Func_florida_drivers_license_number	detects Florida driver's license number	no	
Func_florida_maryland_michigan_minnesota_drivers_license_number	detects Florida, Maryland, Michigan, Minnesota driver's license number	no	
Func_formatted_itin	detects formatted US ITIN	yes	
Func_fr_insee	detects France INSEE	no	
Func_fr_passport	detects France passport	no	
Func_france_eu_tax_file_number	detects France tax file number	no	
Func_france_value_added_tax_number	detects France value added tax number	no	
Func_french_drivers_license	detects French driver's license	no	
Func_french_insee	detects French INSEE	no	
Func_georgia_drivers_license_number	detects Georgia driver's license number	no	

FUNCTION NAME	FUNCTION ACTION	IS A VALIDATOR	
Func_german_drivers_license	detects Germany driver's license	no	
Func_german_passport	detects Germany passport	no	
Func_german_passport_data	detects Germany passport	no	
Func_germany_eu_tax_file_number	detects Germany tax file number	no	
Func_germany_value_added_tax_number	detects Germany value added tax number	no	
Func_greece_eu_ssn	detects Greece sin (AMKA)	no	
Func_hawaii_drivers_license_number	detects Hawaii driver's license number	no	
Func_hong_kong_id_card	detects Hong Kong ID card	no	
Func_hungarian_value_added_tax_number	detects Hungary value added tax number	no	
Func_hungary_eu_national_id_card	detects Hungary personal identification number	no	
Func_hungary_eu_ssn_or_equivalent	detects Hungary social security number	no	
Func_hungary_eu_tax_file_number	detects Hungary tax file number	no	
Func_iban	detects IBAN	yes	
Func_idaho_drivers_license_number	detects Idaho driver's license number	no	
Func_illinois_drivers_license_number	detects Illinois driver's license number	no	
Func_india_aadhaar	detects India aadhaar	yes	
Func_indiana_drivers_license_number	detects Indiana driver's license number	no	
Func_iowa_drivers_license_number	detects Iowa driver's license number	no	
Func_ireland_pps	detects Ireland PPS	no	
Func_israeli_national_id_number	detects Israel national ID number	no	

FUNCTION NAME	FUNCTION ACTION	IS A VALIDATOR	
Func_italy_eu_national_id_card	detects Italy fiscal code	no	
Func_italy_value_added_tax_number	detects Italy value added tax number	no	
Func_japanese_my_number_corporate	detects Japan my number corporate	yes	
Func_japanese_my_number_personal	detects Japan my number personal	yes	
Func_jp_bank_account	detects Japan bank account	no	
Func_jp_bank_account_branch_code	detects Japan bank account branch code	no	
Func_jp_drivers_license_number	detects Japan driver's license number	no	
Func_jp_passport	detects Japan passport	no	
Func_jp_resident_registration_number	detects Japan-resident registration number	no	
Func_jp_sin	detects Japan SIN	no	
Func_jp_sin_pre_1997	detects Japan sin pre 1997	no	
Func_kansas_drivers_license_number	detects Kansas driver's license number	no	
Func_kentucky_drivers_license_number	detects Kentucky driver's license number	no	
Func_kentucky_massachusetts_virginia_drivers_license_number	detects Kentucky, Massachusetts, Virginia driver's license number	no	
Func_latvia_eu_national_id_card	detects Latvia personal code	no	
Func_lithuania_eu_tax_file_number	detects Lithuania personal code	no	
Func_louisiana_drivers_license_number	detects Louisiana driver's license number	no	
Func_luxemburg_eu_tax_file_number	detects Luxemburg national identification number (natural persons)	no	

FUNCTION NAME	FUNCTION ACTION	IS A VALIDATOR	
Func_luxemburg_eu_tax_file_number_non_natural	detects Luxemburg national identification number (non-natural persons)	no	
Func_maine_drivers_license_number	detects Maine driver's license number	no	
Func_manitoba_drivers_license_number	detects Manitoba driver's license number	no	
Func_maryland_drivers_license_number	detects Maryland driver's license number	no	
Func_massachusetts_drivers_license_number	detects Massachusetts driver's license number	no	
Func_mexico_population_registry_code	detects Mexico population registry code	no	
Func_michigan_minnesota_drivers_license_number	detects Michigan, Minnesota driver's license number	no	
Func_minnesota_drivers_license_number	detects Minnesota driver's license number	no	
Func_mississippi_oklahoma_drivers_license_number	detects Mississippi, Oklahoma driver's license number	no	
Func_missouri_drivers_license_number	detects Missouri driver's license number	no	
Func_montana_drivers_license_number	detects Montana driver's license number	no	
Func_nebraska_drivers_license_number	detects Nebraska driver's license number	no	
Func_netherlands_bsn	detects Netherlands BSN	no	
Func_netherlands_eu_tax_file_number	detects Netherlands tax file number	no	
Func_netherlands_value_added_tax_number	detects Netherlands value added tax number	no	
Func_nevada_drivers_license_number	detects Nevada driver's license number	no	
Func_new_brunswick_drivers_license_number	detects New Brunswick driver's license number	no	

FUNCTION NAME	FUNCTION ACTION	IS A VALIDATOR	
Func_new_hampshire_drivers_license_number	detects New Hampshire driver's license number	no	
Func_new_jersey_drivers_license_number	detects New Jersey driver's license number	no	
Func_new_mexico_drivers_license_number	detects New Mexico driver's license number	no	
Func_new_york_drivers_license_number	detects New York driver's license number	no	
Func_new_zealand_bank_account_number	detects New Zealand bank account number	no	
Func_new_zealand_inland_revenue_number	detects New Zealand inland revenue number	no	
Func_new_zealand_ministry_of_health_number	detects New Zealand ministry of health number	no	
Func_newfoundland_labrador_drivers_license_number	detects Newfoundland Labrador driver's license number	no	
Func_newzealand_driver_license_number	detects New Zealand driver license number	no	
Func_newzealand_social_welfare_number	detects New Zealand social welfare number	no	
Func_north_carolina_drivers_license_number	detects North Carolina driver's license number	no	
Func_north_dakota_drivers_license_number	detects North Dakota driver's license number	no	
Func_norway_id_number	detects Norway ID number	no	
Func_nova_scotia_drivers_license_number	detects Nova Scotia driver's license number	no	
Func_ohio_drivers_license_number	detects Ohio driver's license number	no	
Func_ontario_drivers_license_number	detects Ontario driver's license number	no	
Func_pennsylvania_drivers_license_number	detects Pennsylvania driver's license number	no	
Func_pesel_identification_number	detects Poland National ID (PESEL)	no	

FUNCTION NAME	FUNCTION ACTION	IS A VALIDATOR	
Func_poland_eu_tax_file_number	detects Poland tax file number	no	
Func_polish_national_id	detects Poland identity card	no	
Func_polish_passport_number	detects Polish passport number	no	
Func_polish_regon_number	detects Polish REGON number	no	
Func_portugal_eu_tax_file_number	detects Portugal Tax Identification Number	no	
Func_prince_edward_island_drivers_license_number	detects Prince Edward Island driver's license number	no	
Func_quebec_drivers_license_number	detects Quebec driver's license number	no	
Func_randomized_formatted_ssn	detects randomized formatted US SSN	yes	
Func_randomized_unformatted_ssn	detects randomized unformatted US SSN	yes	
Func_rhode_island_drivers_license_number	detects Rhode Island driver's license number	no	
Func_romania_eu_national_id_card	detects Romania personal numeric code (CNP)	no	
Func_saskatchewan_drivers_license_number	detects Saskatchewan driver's license number	no	
Func_slovakia_eu_national_id_card	detects Slovakia personal number	no	
Func_slovenia_eu_national_id_card	detects Slovenia Unique Master Citizen Number	no	
Func_slovenia_eu_tax_file_number	detects Slovenia tax file number	no	
Func_south_africa_identification_number	detects South Africa identification number	yes	
Func_south_carolina_drivers_license_number	detects South Carolina driver's license number	no	
Func_south_dakota_drivers_license_number	detects South Dakota driver's license number	no	

FUNCTION NAME	FUNCTION ACTION	IS A VALIDATOR	
Func_south_korea_resident_number	detects South Korea resident number	no	
Func_spain_eu_DL_and_NI_number_citizen	detects Spain DL and NI number citizen	no	
Func_spain_eu_DL_and_NI_number_foreigner	detects Spain DL and NI number foreigner	no	
Func_spain_eu_driver's_license_number	detects Spain driver's license number	no	
Func_spain_eu_tax_file_number	detects Spain tax file number	no	
Func_spanish_social_security_number	detects Spanish social security number	no	
Func_ssn	Function to detect non-randomized formatted US SSN	yes	
Func_sweden_eu_tax_file_number	detects Sweden tax file number	no	
Func_swedish_national_identifier	detects Swedish national identifier	yes	
Func_swiss_social_security_number_ahv	detects Swiss social security number AHV	no	
Func_taiwanese_national_id	detects Taiwanese national ID	no	
Func_tennessee_drivers_license_number	detects Tennessee driver's license number	no	
Func_texas_drivers_license_number	detects Texas driver's license number	no	
Func_Thai_Citizen_Id	detects Thai Citizen ID	no	
Func_Turkish_National_Id	detects Turkish National ID	yes	
Func_uk_drivers_license	detects UK driver's license	no	
Func_uk_eu_tax_file_number	detects UK unique taxpayer number	no	
Func_uk_nhs_number	detects UK NHS number	yes	
Func_uk_nino	detects UK NINO	no	

FUNCTION NAME	FUNCTION ACTION	IS A VALIDATOR	
Func_unformatted_canadian_sin	detects unformatted Canadian SIN	no	
Func_unformatted_itin	detects unformatted US ITIN	yes	
Func_unformatted_ssn	detects non-randomized unformatted US SSN	yes	
Func_usa_uk_passport	detects USA and UK passport	yes	
Func_utah_drivers_license_number	detects Utah driver's license number	no	
Func_vermont_drivers_license_number	detects Vermont driver's license number	no	
Func_virginia_drivers_license_number	detects Virginia driver's license number	no	
Func_washington_drivers_license_number	detects Washington driver's license number	no	
Func_west_virginia_drivers_license_number	detects West Virginia driver's license number	no	
Func_wisconsin_drivers_license_number	detects Wisconsin driver's license number	no	
Func_wyoming_drivers_license_number	detects Wyoming driver's license number	no	

Func_us_date

Func_us_date looks for dates in common U.S. formats. The common formats are "month/day/year", "month-day-year", and "month day year ". The names or abbreviations of months aren't case-sensitive.

Examples:

- December 2, 2016
- Dec 2, 2016
- dec 02 2016
- 12/2/2016
- 12/02/16
- Dec-2-2016
- 12-2-16

Accepted month names:

- English
  - January, February, march, April, may, June, July, August, September, October, November, December
  - Jan. Feb. Mar. Apr. May June July Aug. Sept. Oct. Nov. Dec.

Func_eu_date

Fund_eu_dates looks for dates in common E.U. formats (and most places outside the U.S.), such as "day/month/year", "day-month-year", and "day month year". The names or abbreviations of months aren't case-sensitive.

Examples:

- 2 Dec 2016
- 02 dec 2016
- 2 Dec 16
- 2/12/2016
- 02/12/16
- 2-Dec-2016
- 2-12-16

Accepted month names:

- English
  - January, February, march, April, may, June, July, August, September, October, November, December
  - Jan. Feb. Mar. Apr. May June July Aug. Sept. Oct. Nov. Dec.
- Dutch
  - januari, februari, maart, April, mei, juni, juli, augustus, September, ocktober, October, November, December
  - jan feb maart apr mei jun jul aug sep sept oct okt nov dec
- French
  - janvier, février, mars, avril, mai, juin juillet, août, septembre, octobre, novembre, décembre
  - janv. févr. mars avril mai juin juil. août sept. oct. nov. déc.
- German
  - jänuar, februar, märz, April, mai, juni juli, August, September, oktober, November, dezember
  - Jan./Jän. Feb. März Apr. Mai Juni Juli Aug. Sept. Okt. Nov. Dez.
- Italian
  - gennaio, febbraio, marzo, aprile, maggio, giugno, luglio, agosto, settembre, ottobre, novembre, dicembre
  - genn. febbr. mar. apr. magg. giugno luglio ag. sett. ott. nov. dic.
- Portuguese

- janeiro, fevereiro, março, marco, abril, maio, junho, julho, agosto, setembro, outubro, novembro, dezembro
- jan fev mar abr mai jun jul ago set out nov dez

- Spanish
  - enero, febrero, marzo, abril, mayo, junio, julio, agosto, septiembre, octubre, noviembre, diciembre
  - enero feb. marzo abr. mayo jun. jul. agosto sept./set. oct. nov. dic.

Func_eu_date1 (deprecated)

> **NOTE**
> This function is deprecated because it supports only Portuguese month names, which are now included in the `Func_eu_date` function above.

This function looks for a date in the format commonly used in Portuguese. The format for this function is the same as `Func_eu_date`, differing only in the language used.

Examples:

- 2 Dez 2016
- 02 dez 2016
- 2 Dez 16
- 2/12/2016
- 02/12/16
- 2-Dez-2016
- 2-12-16

Accepted month names:

- Portuguese
  - janeiro, fevereiro, março, marco, abril, maio, junho, julho, agosto, setembro, outubro, novembro, dezembro
  - jan fev mar abr mai jun jul ago set out nov dez

Func_eu_date2 (deprecated)

> **NOTE**
> This function is deprecated because it supports only Dutch month names, which are now included in the `Func_eu_date` function above.

This function looks for a date in the format commonly used in Dutch. The format for this function is the same as `Func_eu_date`, differing only in the language used.

Examples:

- 2 Mei 2016

- 02 mei 2016
- 2 Mei 16
- 2/12/2016
- 02/12/16
- 2-Mei-2016
- 2-12-16

Accepted month names:

- Dutch
  - januari, februari, maart, April, mei, juni, juli, augustus, September, ocktober, October, November, December
  - jan feb maart apr mei jun jul aug sep sept out okt nov dec

Func_expiration_date

Func_expiration_date looks for dates that are in formats commonly used by credit and debit cards. This function will match dates in format of "month/year", "month-year", "[month name] year", and "[month abbreviation] year". The names or abbreviations of months aren't case-sensitive.

Examples:

- MM/YY -- for example, 01/11 or 1/11
- MM/YYYY -- for example, 01/2011 or 1/2011
- MM-YY -- for example, 01-22 or 1-11
- MM-YYYY -- for example, 01-2000 or 1-2000

The following formats support YY or YYYY:

- Month-YYYY -- for example Jan-2010 or january-2010 or Jan-10 or january-10
- Month YYYY -- for example, 'january 2010' or 'Jan 2010' or 'january 10' or 'Jan 10'
- MonthYYYY -- for example, 'january2010' or 'Jan2010' or 'january10' or 'Jan10'
- Month/YYYY -- for example, 'january/2010' or 'Jan/2010' or 'january/10' or 'Jan/10'

Accepted month names:

- English
  - January, February, march, April, may, June, July, August, September, October, November, December
  - Jan Feb Mar Apr May June July Aug Sept Oct Nov Dec

Func_us_address

Func_us_address looks for a U.S. state name or postal abbreviation followed by a valid zip code. The zip code must be one of the correct zip codes associated with the U.S. state name or abbreviation. The U.S. state name and zip code cannot be separated by punctuation or letters.

Examples:

- Washington 98052
- Washington 98052-9998
- WA 98052
- WA 98052-9998

Customize a built-in sensitive information type

11/2/2020 • 9 minutes to read • Edit Online

When looking for sensitive information in content, you need to describe that information in what's called a *rule*. Data loss prevention (DLP) includes rules for the most-common sensitive information types that you can use right away. To use these rules, you have to include them in a policy. You might find that you want to adjust these built-in rules to meet your organization's specific needs, and you can do that by creating a custom sensitive information type. This topic shows you how to customize the XML file that contains the existing rule collection to detect a wider range of potential credit-card information.

You can take this example and apply it to other built-in sensitive information types. For a list of default sensitive information types and XML definitions, see Sensitive information type entity definitions.

Export the XML file of the current rules

To export the XML, you need to connect to the Security and Compliance Center via Remote PowerShell..

1. In the PowerShell, type the following to display your organization's rules on screen. If you haven't created your own, you'll only see the default, built-in rules, labeled "Microsoft Rule Package."

   ```
   Get-DlpSensitiveInformationTypeRulePackage
   ```

2. Store your organization's rules in a variable by typing the following. Storing something in a variable makes it easily available later in a format that works for remote PowerShell commands.

   ```
   $ruleCollections = Get-DlpSensitiveInformationTypeRulePackage
   ```

3. Make a formatted XML file with all that data by typing the following. (`Set-content` is the part of the cmdlet that writes the XML to the file.)

   ```
   Set-Content -path C:\custompath\exportedRules.xml -Encoding Byte -Value
   $ruleCollections.SerializedClassificationRuleCollection
   ```

   IMPORTANT

   Make sure that you use the file location where your rule pack is actually stored. `C:\custompath\` is a placeholder.

Find the rule that you want to modify in the XML

The cmdlets above exported the entire *rule collection*, which includes the default rules we provide. Next you'll need to look specifically for the Credit Card Number rule that you want to modify.

1. Use a text editor to open the XML file that you exported in the previous section.
2. Scroll down to the `<Rules>` tag, which is the start of the section that contains the DLP rules. Because this XML file contains the information for the entire rule collection, it contains other information at the top that you need to scroll past to get to the rules.
3. Look for *Func_credit_card* to find the Credit Card Number rule definition. In the XML, rule names can't

contain spaces, so the spaces are usually replaced with underscores, and rule names are sometimes abbreviated. An example of this is the U.S. Social Security number rule, which is abbreviated *SSN*. The Credit Card Number rule XML should look like the following code sample.

```
<Entity id="50842eb7-edc8-4019-85dd-5a5c1f2bb085"
       patternsProximity="300" recommendedConfidence="85">
      <Pattern confidenceLevel="85">
       <IdMatch idRef="Func_credit_card" />
        <Any minMatches="1">
          <Match idRef="Keyword_cc_verification" />
          <Match idRef="Keyword_cc_name" />
          <Match idRef="Func_expiration_date" />
        </Any>
      </Pattern>
    </Entity>
```

Now that you have located the Credit Card Number rule definition in the XML, you can customize the rule's XML to meet your needs. For a refresher on the XML definitions, see the Term glossary at the end of this topic.

Modify the XML and create a new sensitive information type

First, you need to create a new sensitive information type because you can't directly modify the default rules. You can do a wide variety of things with custom sensitive information types, which are outlined in Create a custom sensitive information type in Security & Compliance Center PowerShell. For this example, we'll keep it simple and only remove corroborative evidence and add keywords to the Credit Card Number rule.

All XML rule definitions are built on the following general template. You need to copy and paste the Credit Card Number definition XML in the template, modify some values (notice the ". . ." placeholders in the following example), and then upload the modified XML as a new rule that can be used in policies.

```
<?xml version="1.0" encoding="utf-16"?>
<RulePackage xmlns="https://schemas.microsoft.com/office/2011/mce">
  <RulePack id=". . .">
    <Version major="1" minor="0" build="0" revision="0" />
    <Publisher id=". . ." />
    <Details defaultLangCode=". . .">
      <LocalizedDetails langcode=" . . . ">
         <PublisherName>. . .</PublisherName>
         <Name>. . .</Name>
         <Description>. . .</Description>
      </LocalizedDetails>
    </Details>
  </RulePack>

 <Rules>
   <!-- Paste the Credit Card Number rule definition here.-->
      <LocalizedStrings>
         <Resource idRef=". . .">
           <Name default="true" langcode=" . . . ">. . .</Name>
           <Description default="true" langcode=". . ."> . . .</Description>
         </Resource>
      </LocalizedStrings>
   </Rules>
</RulePackage>
```

Now, you have something that looks similar to the following XML. Because rule packages and rules are identified by their unique GUIDs, you need to generate two GUIDs: one for the rule package and one to replace the GUID for the Credit Card Number rule. The GUID for the entity ID in the following code sample is the one for our built-in rule definition, which you need to replace with a new one. There are several ways to generate GUIDs, but you can do it easily in PowerShell by typing **[guid]::NewGuid()**.

```
<?xml version="1.0" encoding="utf-16"?>
<RulePackage xmlns="https://schemas.microsoft.com/office/2011/mce">
  <RulePack id="8aac8390-e99f-4487-8d16-7f0cdee8defc">
    <Version major="1" minor="0" build="0" revision="0" />
    <Publisher id="8d34806e-cd65-4178-ba0e-5d7d712e5b66" />
    <Details defaultLangCode="en">
      <LocalizedDetails langcode="en">
        <PublisherName>Contoso Ltd.</PublisherName>
        <Name>Financial Information</Name>
        <Description>Modified versions of the Microsoft rule package</Description>
      </LocalizedDetails>
    </Details>
  </RulePack>

 <Rules>
    <Entity id="db80b3da-0056-436e-b0ca-1f4cf7080d1f"
       patternsProximity="300" recommendedConfidence="85">
      <Pattern confidenceLevel="85">
        <IdMatch idRef="Func_credit_card" />
        <Any minMatches="1">
          <Match idRef="Keyword_cc_verification" />
          <Match idRef="Keyword_cc_name" />
          <Match idRef="Func_expiration_date" />
        </Any>
      </Pattern>
    </Entity>
      <LocalizedStrings>
         <Resource idRef="db80b3da-0056-436e-b0ca-1f4cf7080d1f">
<!-- This is the GUID for the preceding Credit Card Number entity because the following text is for that
Entity. -->
           <Name default="true" langcode="en-us">Modified Credit Card Number</Name>
           <Description default="true" langcode="en-us">Credit Card Number that looks for additional
keywords, and another version of Credit Card Number that doesn't require keywords (but has a lower
confidence level)</Description>
         </Resource>
      </LocalizedStrings>
   </Rules>
</RulePackage>
```

Remove the corroborative evidence requirement from a sensitive information type

Now that you have a new sensitive information type that you're able to upload to the Security & Compliance Center, the next step is to make the rule more specific. Modify the rule so that it only looks for a 16-digit number that passes the checksum but doesn't require additional (corroborative) evidence, like keywords. To do this, you need to remove the part of the XML that looks for corroborative evidence. Corroborative evidence is very helpful in reducing false positives. In this case there are usually certain keywords or an expiration date near the credit card number. If you remove that evidence, you should also adjust how confident you are that you found a credit card number by lowering the `confidenceLevel`, which is 85 in the example.

```
<Entity id="db80b3da-0056-436e-b0ca-1f4cf7080d1f" patternsProximity="300"
      <Pattern confidenceLevel="85">
        <IdMatch idRef="Func_credit_card" />
      </Pattern>
    </Entity>
```

Look for keywords that are specific to your organization

You might want to require corroborative evidence but want different or additional keywords, and perhaps you

want to change where to look for that evidence. You can adjust the `patternsProximity` to expand or shrink the window for corroborative evidence around the 16-digit number. To add your own keywords, you need to define a keyword list and reference it within your rule. The following XML adds the keywords "company card" and "Contoso card" so that any message that contains those phrases within 150 characters of a credit card number will be identified as a credit card number.

```
<Rules>
<! -- Modify the patternsProximity to be "150" rather than "300." -->
    <Entity id="db80b3da-0056-436e-b0ca-1f4cf7080d1f" patternsProximity="150" recommendedConfidence="85">
      <Pattern confidenceLevel="85">
        <IdMatch idRef="Func_credit_card" />
        <Any minMatches="1">
          <Match idRef="Keyword_cc_verification" />
          <Match idRef="Keyword_cc_name" />
<!-- Add the following XML, which references the keywords at the end of the XML sample. -->
          <Match idRef="My_Additional_Keywords" />
          <Match idRef="Func_expiration_date" />
        </Any>
      </Pattern>
    </Entity>
<!-- Add the following XML, and update the information inside the <Term> tags with the keywords that you
want to detect. -->
    <Keyword id="My_Additional_Keywords">
      <Group matchStyle="word">
        <Term caseSensitive="false">company card</Term>
        <Term caseSensitive="false">Contoso card</Term>
      </Group>
    </Keyword>
```

Upload your rule

To upload your rule, you need to do the following.

1. Save it as an .xml file with Unicode encoding. This is important because the rule won't work if the file is saved with a different encoding.
2. Connect to the Security and Compliance Center via Remote PowerShell.
3. In the PowerShell, type the following.

   ```
   New-DlpSensitiveInformationTypeRulePackage -FileData (Get-Content -Path
   "C:\custompath\MyNewRulePack.xml" -Encoding Byte)
   ```

   IMPORTANT

   Make sure that you use the file location where your rule pack is actually stored. `C:\custompath\` is a placeholder.

4. To confirm, type Y, and then press **Enter**.
5. Verify that your new rule was uploaded and its display name by typing:

   ```
   Get-DlpSensitiveInformationType
   ```

To start using the new rule to detect sensitive information, you need to add the rule to a DLP policy. To learn how to add the rule to a policy, see Create a DLP policy from a template.

Term glossary

These are the definitions for the terms you encountered during this procedure.

TERM	DEFINITION
Entity	Entities are what we call sensitive information types, such as credit card numbers. Each entity has a unique GUID as its ID. If you copy a GUID and search for it in the XML, you'll find the XML rule definition and all the localized translations of that XML rule. You can also find this definition by locating the GUID for the translation and then searching for that GUID.
Functions	The XML file references `Func_credit_card`, which is a function in compiled code. Functions are used to run complex regexes and verify that checksums match for our built-in rules.) Because this happens in the code, some of the variables don't appear in the XML file.
IdMatch	This is the identifier that the pattern is to trying to match—for example, a credit card number.
Keyword lists	The XML file also references `keyword_cc_verification` and `keyword_cc_name`, which are lists of keywords from which we are looking for matches within the `patternsProximity` for the entity. These aren't currently displayed in the XML.
Pattern	The pattern contains the list of what the sensitive type is looking for. This includes keywords, regexes, and internal functions, which perform tasks like verifying checksums. Sensitive information types can have multiple patterns with unique confidences. This is useful when creating a sensitive information type that returns a high confidence if corroborative evidence is found and a lower confidence if little or no corroborative evidence is found.
Pattern confidenceLevel	This is the level of confidence that the DLP engine found a match. This level of confidence is associated with a match for the pattern if the pattern's requirements are met. This is the confidence measure you should consider when using Exchange mail flow rules (also known as transport rules).
patternsProximity	When we find what looks like a credit card number pattern, `patternsProximity` is the proximity around that number where we'll look for corroborative evidence.

TERM	DEFINITION
recommendedConfidence	This is the confidence level we recommend for this rule. The recommended confidence applies to entities and affinities. For entities, this number is never evaluated against the `confidenceLevel` for the pattern. It's merely a suggestion to help you choose a confidence level if you want to apply one. For affinities, the `confidenceLevel` of the pattern must be higher than the `recommendedConfidence` number for a mail flow rule action to be invoked. The `recommendedConfidence` is the default confidence level used in mail flow rules that invokes an action. If you want, you can manually change the mail flow rule to be invoked based off the pattern's confidence level, instead.

For more information

- Sensitive information type entity definitions
- Create a custom sensitive information type
- Overview of data loss prevention policies

Create custom sensitive information types with Exact Data Match based classification

2/18/2021 • 22 minutes to read • Edit Online

Custom sensitive information types are used to help identify sensitive items so that you can prevent them from being inadvertently or inappropriately shared. You define a custom sensitive information type based on:

- patterns
- keyword evidence such as *employee*, *badge*, or *ID*
- character proximity to evidence in a particular pattern
- confidence levels

Such custom sensitive information types meet business needs for many organizations.

But what if you wanted a custom sensitive information type that uses exact data values, instead of one that found matches based on generic patterns? With Exact Data Match (EDM)-based classification, you can create a custom sensitive information type that is designed to:

- be dynamic and easily refreshed
- be more scalable
- result in fewer false-positives
- work with structured sensitive data
- handle sensitive information more securely
- be used with several Microsoft cloud services

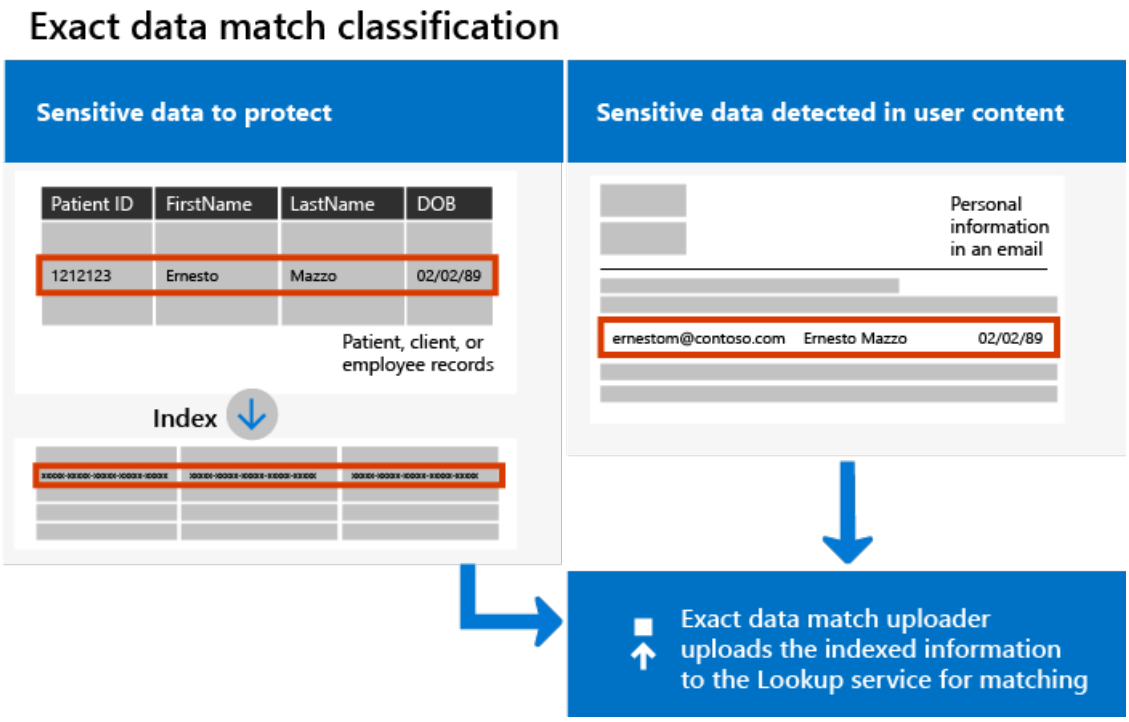

EDM-based classification enables you to create custom sensitive information types that refer to exact values in a database of sensitive information. The database can be refreshed daily, and contain up to 100 million rows of data. So as employees, patients, or clients come and go, and records change, your custom sensitive information types remain current and applicable. And, you can use EDM-based classification with policies, such as data loss prevention policies (DLP) or Microsoft Cloud App Security file policies.

NOTE

Microsoft 365 Information Protection now supports in preview double byte character set languages for:

- Chinese (simplified)
- Chinese (traditional)
- Korean
- Japanese

This support is available for sensitive information types. See, Information protection support for double byte character sets release notes (preview) for more information.

Required licenses and permissions

You must be a global admin, compliance administrator, or Exchange Online administrator to perform the tasks described in this article. To learn more about DLP permissions, see Permissions.

EDM-based classification is included in these subscriptions

- Office 365 E5
- Microsoft 365 E5
- Microsoft 365 E5 Compliance
- Microsoft E5/A5 Information Protection and Governance

Portal links for your subscription

PORTAL	WORLD WIDE/GCC	GCC-HIGH	DOD
Office SCC	protection.office.com	scc.office365.us	scc.protection.apps.mil
Microsoft 365 Security center	security.microsoft.com	security.microsoft.us	security.apps.mil
Microsoft 365 Compliance center	compliance.microsoft.com	compliance.microsoft.us	compliance.apps.mil

The work flow at a glance

PHASE	WHAT'S NEEDED
Part 1: Set up EDM-based classification (As needed) - Edit the database schema - Remove the schema	- Read access to the sensitive data - Database schema in XML format (example provided) - Rule package in XML format (example provided) - Admin permissions to the Security & Compliance Center (using PowerShell)
Part 2: Hash and upload the sensitive data (As needed) Refresh the data	- Custom security group and user account - Local admin access to machine with EDM Upload Agent - Read access to the sensitive data - Process and schedule for refreshing the data
Part 3: Use EDM-based classification with your Microsoft cloud services	- Microsoft 365 subscription with DLP - EDM-based classification feature enabled

Part 1: Set up EDM-based classification

Setting up and configuring EDM-based classification involves:

1. Saving sensitive data in .csv format
2. Define your sensitive information database schema
3. Create a rule package

Save sensitive data in .csv format

1. Identify the sensitive information you want to use. Export the data to an app, such as Microsoft Excel, and save the file in .csv format. The data file can include a maximum of:

   - Up to 100 million rows of sensitive data
   - Up to 32 columns (fields) per data source
   - Up to 5 columns (fields) marked as searchable

2. Structure the sensitive data in the .csv file such that the first row includes the names of the fields used for EDM-based classification. In your .csv file, you might have field names, such as "ssn", "birthdate", "firstname", "lastname". The column header names can't include spaces or underscores. For example, the sample .csv file that we use in this article is named *PatientRecords.csv*, and its columns include *PatientID*, *MRN*, *LastName*, *FirstName*, *SSN*, and more.

3. Pay attention to the format of the sensitive data fields. In particular, fields that may contain commas in their content (e.g. a street address that contains the value "Seattle,WA") would be parsed as two separate fields when parsed by the EDM tool. In order to avoid this, you need to ensure such fields are surrounded by single or double quotes in the sensitive data table. If fields with commas in them may also contain spaces, you would need to create a custom Sensitive Information Type that matches the corresponding format (e.g. a multi-word string with commas and spaces in it) to ensure the string is correctly matched when the document is scanned.

Define the schema for your database of sensitive information

If for business or technical reasons, you prefer not to use PowerShell or command line to create your schema and EDM sensitive info type patter (rule package), you can use the Exact Data Match Schema and Sensitive Information Type Wizard to create them. When you are done creating the schema and EDM sensitive info type pattern, return to complete all the steps necessary to make your EDM based sensitive information type available for use.

> **NOTE**
> The Exact Data Match Schema and Sensitive Information Type Wizard is only available for the World Wide and GCC clouds only.

1. Define the schema for the database of sensitive information in XML format (similar to our example below). Name this schema file **edm.xml**, and configure it such that for each column in the database, there is a line that uses the syntax:

   `\<Field name="" searchable=""/\>`.

   - Use column names for *Field name* values.
   - Use *searchable="true"* for the fields that you want to be searchable up to a maximum of 5 fields. At least one field must be searchable.

   As an example, the following XML file defines the schema for a patient records database, with five fields specified as searchable: *PatientID*, *MRN*, *SSN*, *Phone*, and *DOB*.

   (You can copy, modify, and use our example.)

```
<EdmSchema xmlns="http://schemas.microsoft.com/office/2018/edm">
      <DataStore name="PatientRecords" description="Schema for patient records" version="1">
            <Field name="PatientID" searchable="true" caseInsensitive="true" ignoredDelimiters="-
,/,*,#,^" />
            <Field name="MRN" searchable="true" />
            <Field name="FirstName" />
            <Field name="LastName" />
            <Field name="SSN" searchable="true" />
            <Field name="Phone" searchable="true" />
            <Field name="DOB" searchable="true" />
            <Field name="Gender" />
            <Field name="Address" />
      </DataStore>
</EdmSchema>
```

Configurable match using the caseInsensitive and ignoredDelimiters fields

The above XML sample makes use of the `caseInsensitive` and the `ignoredDelimiters` fields.

When you include the ***caseInsensitive*** field set to the value of `true` in your schema definition, EDM will not exclude an item based on case differences for `PatientID` field. So EDM will see, `PatientID` **FOO-1234** and **fOo-1234** as being identical.

When you include the ***ignoredDelimiters*** field with supported characters, EDM will ignore those characters in the `PatientID`. So EDM will see, `PatientID` **FOO-1234** and `PatientID` **FOO#1234** as being identical. The `ignoredDelimiters` flag supports any non-alphanumeric character, here are some examples:

- .
- -
- /
- _
- *
- ^
- #
- !
- ?
- [
-]
- {
- }
- \
- ~
- ;
- The `ignoredDelimiters` flag doesn't support:
- characters 0-9
- A-Z

- a-z
- "
- ,

In this example, where both `caseInsensitive` and `ignoredDelimiters` are used, EDM would see **FOO-1234** and **fOo#1234** as identical and classify the item as a patient record sensitive information type.

4. Connect to the Security & Compliance center using the procedures in Connect to Security & Compliance Center PowerShell.

5. To upload the database schema, run the following cmdlets, one at a time:

   ```
   $edmSchemaXml=Get-Content .\\edm.xml -Encoding Byte -ReadCount 0
   New-DlpEdmSchema -FileData $edmSchemaXml -Confirm:$true
   ```

   You will be prompted to confirm, as follows:

   > Confirm
   >
   > Are you sure you want to perform this action?
   >
   > New EDM Schema for the data store 'patientrecords' will be imported.
   >
   > [Y] Yes [A] Yes to All [N] No [L] No to All [?] Help (default is "Y"):

> **TIP**
> If you want your changes to occur without confirmation, in Step 5, use this cmdlet instead: New-DlpEdmSchema -FileData $edmSchemaXml

> **NOTE**
> It can take between 10-60 minutes to update the EDMSchema with additions. The update must complete before you execute steps that use the additions.

Set up a rule package

1. Create a rule package in XML format (with Unicode encoding), similar to the following example. (You can copy, modify, and use our example.)

   When you set up your rule package, make sure to correctly reference your .csv file and **edm.xml** file. You can copy, modify, and use our example. In this sample xml the following fields needs to be customized to create your EDM sensitive type:

   - **RulePack id & ExactMatch id**: Use New-GUID to generate a GUID.
   - **Datastore**: This field specifies EDM lookup data store to be used. You provide a data source name of a configured EDM Schema.
   - **idMatch**: This field points to the primary element for EDM.
     - Matches: Specifies the field to be used in exact lookup. You provide a searchable field name in EDM Schema for the DataStore.
     - Classification: This field specifies the sensitive type match that triggers EDM lookup. You can provide the Name or GUID of an existing built-in or custom sensitive information type. Be

aware that any string that matches the sensitive information type provided will be hashed and compared to every entry in the sensitive information table. In order to avoid causing performance issues, if you use a custom sensitive information type as the Classification element in EDM, avoid using one that will match a large percentage of content (such as "any number" or "any five-letter word") by adding supporting keywords or including formatting in the definition of the custom classification sensitive information type.

- **Match:** This field points to additional evidence found in proximity of idMatch.
  - Matches: You provide any field name in EDM Schema for DataStore.
- **Resource:** This section specifies the name and description for sensitive type in multiple locales.
  - idRef: You provide GUID for ExactMatch ID.
  - Name & descriptions: customize as required.

```
<RulePackage xmlns="http://schemas.microsoft.com/office/2018/edm">
  <RulePack id="fd098e03-1796-41a5-8ab6-198c93c62b11">
    <Version build="0" major="2" minor="0" revision="0" />
    <Publisher id="eb553734-8306-44b4-9ad5-c388ad970528" />
    <Details defaultLangCode="en-us">
      <LocalizedDetails langcode="en-us">
        <PublisherName>IP DLP</PublisherName>
        <Name>Health Care EDM Rulepack</Name>
        <Description>This rule package contains the EDM sensitive type for health care sensitive
types.</Description>
      </LocalizedDetails>
    </Details>
  </RulePack>
  <Rules>
    <ExactMatch id = "E1CC861E-3FE9-4A58-82DF-4BD259EAB371" patternsProximity = "300" dataStore
="PatientRecords" recommendedConfidence = "65" >
      <Pattern confidenceLevel="65">
        <idMatch matches = "SSN" classification = "U.S. Social Security Number (SSN)" />
      </Pattern>
      <Pattern confidenceLevel="75">
        <idMatch matches = "SSN" classification = "U.S. Social Security Number (SSN)" />
        <Any minMatches ="3" maxMatches ="6">
          <match matches="PatientID" />
          <match matches="MRN"/>
          <match matches="FirstName"/>
          <match matches="LastName"/>
          <match matches="Phone"/>
          <match matches="DOB"/>
        </Any>
      </Pattern>
    </ExactMatch>
    <LocalizedStrings>
      <Resource idRef="E1CC861E-3FE9-4A58-82DF-4BD259EAB371">
        <Name default="true" langcode="en-us">Patient SSN Exact Match.</Name>
        <Description default="true" langcode="en-us">EDM Sensitive type for detecting Patient SSN.
</Description>
      </Resource>
    </LocalizedStrings>
  </Rules>
</RulePackage>
```

2. Upload the rule package by running the following PowerShell cmdlets, one at a time:

```
$rulepack=Get-Content .\\rulepack.xml -Encoding Byte -ReadCount 0
New-DlpSensitiveInformationTypeRulePackage -FileData $rulepack
```

At this point, you have set up EDM-based classification. The next step is to hash the sensitive data, and then

upload the hashes for indexing.

Recall from the previous procedure that our PatientRecords schema defines five fields as searchable: *PatientID*, *MRN*, *SSN*, *Phone*, and *DOB*. Our example rule package includes those fields and references the database schema file (**edm.xml**), with one *ExactMatch* item per searchable field. Consider the following ExactMatch item:

```
<ExactMatch id = "E1CC861E-3FE9-4A58-82DF-4BD259EAB371" patternsProximity = "300" dataStore
="PatientRecords" recommendedConfidence = "65" >
      <Pattern confidenceLevel="65">
        <idMatch matches = "SSN" classification = "U.S. Social Security Number (SSN)" />
      </Pattern>
      <Pattern confidenceLevel="75">
        <idMatch matches = "SSN" classification = "U.S. Social Security Number (SSN)" />
        <Any minMatches ="3" maxMatches ="100">
          <match matches="PatientID" />
          <match matches="MRN"/>
          <match matches="FirstName"/>
          <match matches="LastName"/>
          <match matches="Phone"/>
          <match matches="DOB"/>
        </Any>
      </Pattern>
    </ExactMatch>
```

In this example, note that:

- The dataStore name references the .csv file we created earlier: **dataStore = "PatientRecords"**.
- The idMatch value references a searchable field that is listed in the database schema file: **idMatch matches = "SSN"**.
- The classification value references an existing or custom sensitive information type: **classification = "U.S. Social Security Number (SSN)"**. (In this case, we use the existing sensitive information type of U.S. Social Security Number.)

> **NOTE**
> It can take between 10-60 minutes to update the EDMSchema with additions. The update must complete before you execute steps that use the additions.

Editing the schema for EDM-based classification

If you want to make changes to your **edm.xml** file, such as changing which fields are used for EDM-based classification, follow these steps:

> **TIP**
> You can change your EDM schema and data file to take advantage of **configurable match**. When configured, EDM will ignore case differences and some delimiters when it evaluates an item. This makes defining your xml schema and your sensitive data files easier. To learn more see, Modify Exact Data Match schema to use configurable match.

1. Edit your **edm.xml** file (this is the file discussed in the Define the schema section of this article).
2. Connect to the Security & Compliance center using the procedures in Connect to Security & Compliance Center PowerShell.
3. To update your database schema, run the following cmdlets, one at a time:

```
$edmSchemaXml=Get-Content .\\edm.xml -Encoding Byte -ReadCount 0
Set-DlpEdmSchema -FileData $edmSchemaXml -Confirm:$true
```

You will be prompted to confirm, as follows:

> Confirm
>
> Are you sure you want to perform this action?
>
> EDM Schema for the data store 'patientrecords' will be updated.
>
> [Y] Yes [A] Yes to All [N] No [L] No to All [?] Help (default is "Y"):

> **TIP**
> If you want your changes to occur without confirmation, in Step 3, use this cmdlet instead: Set-DlpEdmSchema -FileData $edmSchemaXml

> **NOTE**
> It can take between 10-60 minutes to update the EDMSchema with additions. The update must complete before you execute steps that use the additions.

Removing the schema for EDM-based classification

(As needed) If you want to remove the schema you're using for EDM-based classification, follow these steps:

1. Connect to the Security & Compliance center using the procedures in Connect to Security & Compliance Center PowerShell.
2. Run the following PowerShell cmdlets, substituting the data store name of "patient records" with the one you want to remove:

```
Remove-DlpEdmSchema -Identity patientrecords
```

You will be prompted to confirm:

> Confirm
>
> Are you sure you want to perform this action?
>
> EDM Schema for the data store 'patientrecords' will be removed.
>
> [Y] Yes [A] Yes to All [N] No [L] No to All [?] Help (default is "Y"):

> **TIP**
> If you want your changes to occur without confirmation, in Step 2, use this cmdlet instead: Remove-DlpEdmSchema -Identity patientrecords -Confirm:$false

Part 2: Hash and upload the sensitive data

In this phase, you set up a custom security group and user account, and set up the EDM Upload Agent tool. Then, you use the tool to hash with salt value the sensitive data, and upload it.

The hashing and uploading can be done using one computer or you can separate the hashing step from the upload step for greater security.

If you want to hash and upload from one computer, you need to do it from a computer that can directly connect to your Microsoft 365 tenant. This requires that your clear text sensitive data files are on that computer for hashing.

If you do not want to expose your clear text sensitive data file, you can hash it on a computer in a secure location and then copy the hash file and the salt file to a computer that can directly connect to your Microsoft 365 tenant for upload. In this scenario, you will need the EDMUploadAgent on both computers.

> **IMPORTANT**
> If you used the Exact Data Match schema and sensitive information type wizard to create your schema and pattern files, you ***must** download the schema for this procedure.

Prerequisites

- a work or school account for Microsoft 365 that will be added to the **EDM_DataUploaders** security group
- a Windows 10 or Windows Server 2016 machine with .NET version 4.6.2 for running the EDMUploadAgent
- a directory on your upload machine for the:
  - EDMUploadAgent
  - your sensitive item file in csv format **PatientRecords.csv** in our examples
  - and the output hash and salt files
  - the datastore name from the **edm.xml** file, for this example its `PatientRecords`
- If you used the Exact Data Match schema and sensitive information type wizard you ***must*** download it

Set up the security group and user account

1. As a global administrator, go to the admin center using the appropriate link for your subscription and create a security group called **EDM_DataUploaders**.
2. Add one or more users to the **EDM_DataUploaders** security group. (These users will manage the database of sensitive information.)

Hash and upload from one computer

This computer must have direct access to your Microsoft 365 tenant.

> **NOTE**
> Before you begin this procedure, make sure that you are a member of the **EDM_DataUploaders** security group.

> **TIP**
> Optionally, you can run a validation against your csv file before uploading by running:
>
> `EdmUploadAgent.exe /ValidateData /DataFile [data file] /Schema [schema file]`
>
> For more information on all the EdmUploadAgent.exe >supported parameters run
>
> `EdmUploadAgent.exe /?`

Links to EDM upload agent by subscription type

- Commercial + GCC - most commercial customers should use this
- GCC-High - This is specifically for high security government cloud subscribers
- DoD - this is specifically for United States Department of Defense cloud customers

1. Create a working directory for the EDMUploadAgent. For example, **C:\EDM\Data**. Place the **PatientRecords.csv** file there.

2. Download and install the appropriate EDM Upload Agent for your subscription into the directory you created in step 1.

   > **NOTE**
   >
   > The EDMUploadAgent at the above links has been updated to automatically add a salt value to the hashed data. Alternately, you can provide your own salt value. Once you have used this version, you will not be able to use the previous version of the EDMUploadAgent.
   >
   > You can upload data with the EDMUploadAgent to any given data store only twice per day.

   > **TIP**
   >
   > To a get a list out of the supported command parameters, run the agent no arguments. For example 'EdmUploadAgent.exe'.

3. Authorize the EDM Upload Agent, open Command Prompt window (as an administrator), switch to the **C:\EDM\Data** directory and then run the following command:

   ```
   EdmUploadAgent.exe /Authorize
   ```

4. Sign in with your work or school account for Microsoft 365 that was added to the EDM_DataUploaders security group. Your tenant information is extracted from the user account to make the connection.

   OPTIONAL: If you used the Exact Data Match schema and sensitive information type wizard to create your schema and pattern files, run the following command in a Command Prompt window:

   ```
   EdmUploadAgent.exe /SaveSchema /DataStoreName <schema name> /OutputDir <path to output folder>
   ```

5. To hash and upload the sensitive data, run the following command in Command Prompt window:

   ```
   EdmUploadAgent.exe /UploadData /DataStoreName [DS Name] /DataFile [data file] /HashLocation [hash file location] /Schema [Schema file]
   ```

   Example: **EdmUploadAgent.exe /UploadData /DataStoreName PatientRecords /DataFile C:\Edm\Hash\PatientRecords.csv /HashLocation C:\Edm\Hash /Schema edm.xml**

   This will automatically add a randomly generated salt value to the hash for greater security. Optionally, if you want to use your own salt value, add the **/Salt** to the command. This value must be 64 characters in length and can only contain the a-z characters and 0-9 characters.

6. Check the upload status by running this command:

   ```
   EdmUploadAgent.exe /GetSession /DataStoreName \<DataStoreName\>
   ```

   Example: **EdmUploadAgent.exe /GetSession /DataStoreName PatientRecords**

   Look for the status to be in **ProcessingInProgress**. Check again every few minutes until the status changes to **Completed**. Once the status is completed, your EDM data is ready for use.

Separate Hash and upload

Perform the hash on a computer in a secure environment.

OPTIONAL: If you used the Exact Data Match schema and sensitive information type wizard to create your schema and pattern files, run the following command in a Command Prompt window:

```
EdmUploadAgent.exe /SaveSchema /DataStoreName <schema name> /OutputDir <path to output folder>
```

1. Run the following command in Command Prompt windows:

   ```
   EdmUploadAgent.exe /CreateHash /DataFile [data file] /HashLocation [hash file location] /Schema
   [Schema file] >
   ```

   For example:

   > **EdmUploadAgent.exe /CreateHash /DataFile C:\Edm\Data\PatientRecords.csv /HashLocation C:\Edm\Hash /Schema edm.xml**

   This will output a hashed file and a salt file with these extensions if you didn't specify the **/Salt** option:

   - .EdmHash
   - .EdmSalt

2. Copy these files in a secure fashion to the computer you will use to upload your sensitive items csv file (PatientRecords) to your tenant.

   To upload the hashed data, run the following command in Windows Command Prompt:

   ```
   EdmUploadAgent.exe /UploadHash /DataStoreName \<DataStoreName\> /HashFile \<HashedSourceFilePath\>
   ```

   For example:

   > **EdmUploadAgent.exe /UploadHash /DataStoreName PatientRecords /HashFile C:\Edm\Hash\PatientRecords.EdmHash**

   To verify that your sensitive data has been uploaded, run the following command in Command Prompt window:

   ```
   EdmUploadAgent.exe /GetDataStore
   ```

   You'll see a list of data stores and when they were last updated.

   If you want to see all the data uploads to a particular store, run the following command in a Windows command prompt:

   ```
   EdmUploadAgent.exe /GetSession /DataStoreName <DataStoreName>
   ```

   Proceed to set up your process and schedule for Refreshing your sensitive information database.

At this point, you are ready to use EDM-based classification with your Microsoft cloud services. For example, you can set up a DLP policy using EDM-based classification.

Refreshing your sensitive information database

You can refresh your sensitive information database daily, and the EDM Upload Tool can reindex the sensitive data and then reupload the indexed data.

1. Determine your process and frequency (daily or weekly) for refreshing the database of sensitive information.

2. Re-export the sensitive data to an app, such as Microsoft Excel, and save the file in .csv format. Keep the same file name and location you used when you followed the steps described in Hash and upload the sensitive data.

> **NOTE**
>
> If there are no changes to the structure (field names) of the .csv file, you won't need to make any changes to your database schema file when you refresh the data. But if you must make changes, make sure to edit the database schema and your rule package accordingly.

3. Use Task Scheduler to automate steps 2 and 3 in the Hash and upload the sensitive data procedure. You can schedule tasks using several methods:

METHOD	WHAT TO DO
Windows PowerShell	See the ScheduledTasks documentation and the example PowerShell script in this article
Task Scheduler API	See the Task Scheduler documentation
Windows user interface	In Windows, click **Start**, and type Task Scheduler. Then, in the list of results, right-click **Task Scheduler**, and choose **Run as administrator**.

Example PowerShell script for Task Scheduler

This section includes an example PowerShell script you can use to schedule your tasks for hashing data and uploading the hashed data:

To schedule hashing and upload in a combined step

```
param(\[string\]$dataStoreName,\[string\]$fileLocation)
\# Assuming current user is also the user context to run the task
$user = "$env:USERDOMAIN\\$env:USERNAME"
$edminstallpath = 'C:\\Program Files\\Microsoft\\EdmUploadAgent\\'
$edmuploader = $edminstallpath + 'EdmUploadAgent.exe'
$csvext = '.csv'
$schemaext = '.xml'
\# Assuming CSV file name is same as data store name
$dataFile = "$fileLocation\\$dataStoreName$csvext"
\# Assuming location to store hash file is same as the location of csv file
$hashLocation = $fileLocation
\# Assuming Schema file name is same as data store name
$schemaFile = "$fileLocation\\$dataStoreName$schemaext"
$uploadDataArgs = '/UploadData /DataStoreName ' + $dataStoreName + ' /DataFile ' + $dataFile + '
/HashLocation' + $hashLocation + ' /Schema ' + $schemaFile
\# Set up actions associated with the task
$actions = @()
$actions += New-ScheduledTaskAction -Execute $edmuploader -Argument $uploadDataArgs -WorkingDirectory
$edminstallpath
\# Set up trigger for the task
$trigger = New-ScheduledTaskTrigger -Weekly -DaysOfWeek Sunday -At 2am
\# Set up task settings
$principal = New-ScheduledTaskPrincipal -UserId $user -LogonType S4U -RunLevel Highest
$settings = New-ScheduledTaskSettingsSet -RunOnlyIfNetworkAvailable -StartWhenAvailable -WakeToRun
\# Create the scheduled task
$scheduledTask = New-ScheduledTask -Action $actions -Principal $principal -Trigger $trigger -Settings
$settings
\# Get credentials to run the task
$creds = Get-Credential -UserName $user -Message "Enter credentials to run the task"
$password=\[Runtime.InteropServices.Marshal\]::PtrToStringAuto(\
[Runtime.InteropServices.Marshal\]::SecureStringToBSTR($creds.Password))
\# Register the scheduled task
$taskName = 'EDMUpload\_' + $dataStoreName
Register-ScheduledTask -TaskName $taskName -InputObject $scheduledTask -User $user -Password $password
```

To schedule hashing and upload as separate steps

```
param(\[string\]$dataStoreName,\[string\]$fileLocation)
\# Assuming current user is also the user context to run the task
$user = "$env:USERDOMAIN\\$env:USERNAME"
$edminstallpath = 'C:\\Program Files\\Microsoft\\EdmUploadAgent\\'
$edmuploader = $edminstallpath + 'EdmUploadAgent.exe'
$csvext = '.csv'
$edmext = '.EdmHash'
$schemaext = '.xml'
\# Assuming CSV file name is same as data store name
$dataFile = "$fileLocation\\$dataStoreName$csvext"
$hashFile = "$fileLocation\\$dataStoreName$edmext"
\# Assuming Schema file name is same as data store name
$schemaFile = "$fileLocation\\$dataStoreName$schemaext "

\# Assuming location to store hash file is same as the location of csv file
$hashLocation = $fileLocation
$createHashArgs = '/CreateHash' + ' /DataFile ' + $dataFile + ' /HashLocation ' + $hashLocation + ' /Schema
' + $schemaFile
$uploadHashArgs = '/UploadHash /DataStoreName ' + $dataStoreName + ' /HashFile ' + $hashFile
\# Set up actions associated with the task
$actions = @()
$actions += New-ScheduledTaskAction -Execute $edmuploader -Argument $createHashArgs -WorkingDirectory
$edminstallpath
$actions += New-ScheduledTaskAction -Execute $edmuploader -Argument $uploadHashArgs -WorkingDirectory
$edminstallpath
\# Set up trigger for the task
$trigger = New-ScheduledTaskTrigger -Weekly -DaysOfWeek Sunday -At 2am
\# Set up task settings
$principal = New-ScheduledTaskPrincipal -UserId $user -LogonType S4U -RunLevel Highest
$settings = New-ScheduledTaskSettingsSet -RunOnlyIfNetworkAvailable -StartWhenAvailable -WakeToRun
\# Create the scheduled task
$scheduledTask = New-ScheduledTask -Action $actions -Principal $principal -Trigger $trigger -Settings
$settings
\# Get credentials to run the task
$creds = Get-Credential -UserName $user -Message "Enter credentials to run the task"
$password=\[Runtime.InteropServices.Marshal\]::PtrToStringAuto(\
[Runtime.InteropServices.Marshal\]::SecureStringToBSTR($creds.Password))
\# Register the scheduled task
$taskName = 'EDMUpload\_' + $dataStoreName
Register-ScheduledTask -TaskName $taskName -InputObject $scheduledTask -User $user -Password $password
```

Part 3: Use EDM-based classification with your Microsoft cloud services

These locations are support EDM sensitive information types:

- DLP for Exchange Online (email)
- OneDrive for Business (files)
- Microsoft Teams (conversations)
- DLP for SharePoint (files)
- Microsoft Cloud App Security DLP policies

EDM sensitive information types for following scenarios are currently in development, but not yet available:

- Auto-classification of sensitivity labels and retention labels

To create a DLP policy with EDM

1. Go to the Security & Compliance Center using the appropriate link for your subscription.
2. Choose **Data loss prevention** > **Policy**.
3. Choose **Create a policy** > **Custom** > **Next**.

4. On the **Name your policy** tab, specify a name and description, and then choose **Next**.
5. On the **Choose locations** tab, select **Let me choose specific locations**, and then choose **Next**.
6. In the **Status** column, select **Exchange email, OneDrive accounts, Teams chat and channel message**, and then choose **Next**.
7. On the **Policy settings** tab, choose **Use advanced settings**, and then choose **Next**.
8. Choose **+ New rule**.
9. In the **Name** section, specify a name and description for the rule.
10. In the **Conditions** section, in the **+ Add a condition** list, choose **Content contains sensitive type**.

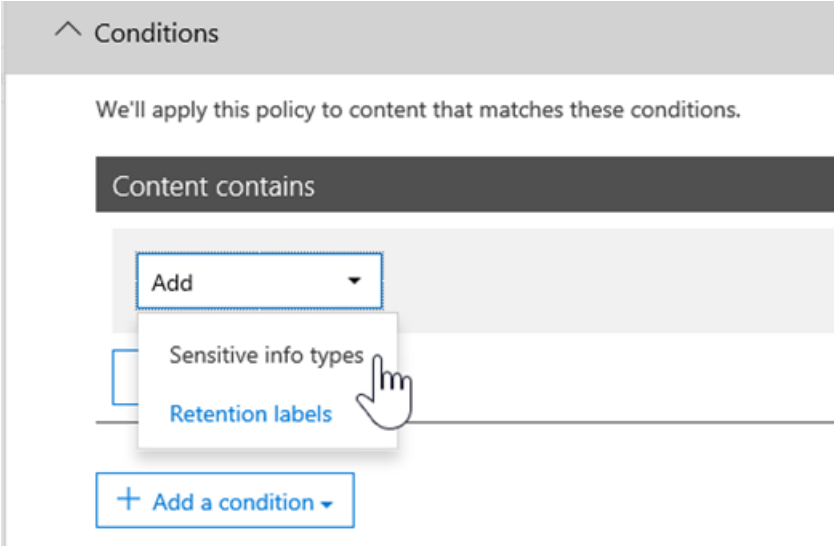

11. Search for the sensitive information type you created when you set up your rule package, and then choose **+ Add**.
    Then choose **Done**.
12. Finish selecting options for your rule, such as **User notifications**, **User overrides**, **Incident reports**, and so on, and then choose **Save**.
13. On the **Policy settings** tab, review your rules, and then choose **Next**.
14. Specify whether to turn on the policy right away, test it out, or keep it turned off. Then choose **Next**.
15. On the **Review your settings** tab, review your policy. Make any needed changes. When you're ready, choose **Create**.

> **NOTE**
> Allow approximately one hour for your new DLP policy to work its way through your data center.

Related articles

- Sensitive information type-entity definitions
- Learn about sensitive information types
- Overview of DLP policies
- Microsoft Cloud App Security
- New-DlpEdmSchema
- Modify Exact Data Match schema to use configurable match

Use the Exact Data Match Schema and Sensitive Information Type Wizard

2/18/2021 • 2 minutes to read • Edit Online

Creating a custom sensitive information type with Exact Data Match (EDM) based classification involves many steps. You can use this wizard to create your schema and sensitive information type pattern (rule package) files to help simplify the process.

> **NOTE**
> The Exact Data Match Schema and Sensitive Information Type Wizard is only available for the World Wide and GCC clouds only.

This wizard can be used instead of the:

- Define the schema for your database of sensitive information
- Set up a pattern (rule package)

steps in Part 1: Set up EDM-based classification.

Pre-requisites

1. Familiarize yourself with the steps to create a custom sensitive information type with EDM work flow at a glance.
2. Perform the steps in the Save sensitive data in .csv format section.

Use the exact data match schema and sensitive information type pattern wizard

1. In the Microsoft 365 Compliance center for your tenant go to **Data classification** > **Exact data matches**.
2. Choose **Create EDM schema** to open the schema wizard configuration flyout.

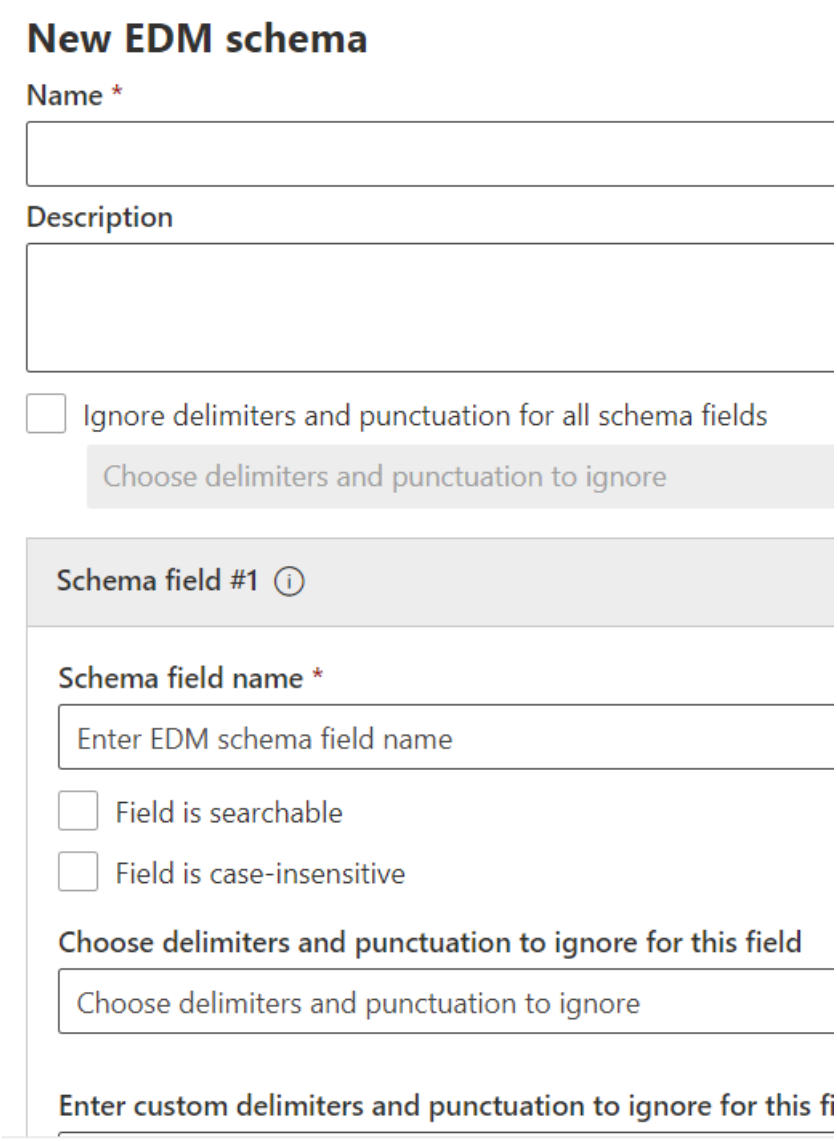

3. Fill in an appropriate **Name** and **Description**.
4. Choose **Ignore delimiters and punctuations for all schema fields** if you want that behavior. To learn more about configuring EDM to ignore case or delimitere, see Creating a custom sensitive information type with Exact Data Match (EDM) based classification.
5. Fill in your desired values for your **Schema field #1** and add more fields as needed.

> **IMPORTANT**
> At least one, but no more than five of your schema fields must be designated as searchable.

6. Choose save. Your schema will now be listed.
7. Choose **EDM sensitive info types** and **Create EDM sensitive info type** to open the sensitive info type configuration wizard.
8. Choose **Choose an existing EDM schema** and choose the schema you created in steps 2-6 from the list.
9. Choose **Next** and choose **Create pattern**.
10. Choose the **Confidence level** and **Primary element**. To learn more about configuring a pattern, see Create a custom sensitive information type in the Compliance Center

11. Choose the **Primary element's sensitive info type** to associate it with. See Sensitive Information Type Entity Definitions to learn more about the available sensitive information types.
12. Choose **Done**.
13. Choose your desired **Confidence level and character proximity**. This will be the default value for the whole EDM sensitive info type
14. Choose **Create pattern** if you want to creaet additional patterns for your EDM sensitive info type.
15. Choose **Next** and fill in a **Name** and **Description for admins**.
16. Review and choose **Submit**.

You can delete or edit the sensitive information type pattern by selecting it which surfaces the edit and delete controls.

> **IMPORTANT**
> If you want to remove a schema, and it is already associated with an EDM sensitive info type, you must first delete the EDM sensitive info type, then you can delete the schema.

Post steps

After you have used this wizard to create your EDM schema and pattern (rule package) files, you still have to perform the steps in Part 2: Hash and upload the sensitive data before you can use the EDM custom sensitive information type.

Modify Exact Data Match schema to use configurable match

2/18/2021 • 2 minutes to read • Edit Online

Exact Data Match (EDM) based classification enables you to create custom sensitive information types that refer to exact values in a database of sensitive information. When you need to allow for variants of a exact string, you can use *configurable match* to tell Microsoft 365 to ignore case and some delimiters.

> **IMPORTANT**
> Use this procedure to modify an existing EDM schema and data file.

1. Uninstall the **EdmUploadAgent.exe** from the computer that you use to connect to Microsoft 365 for EDM schema and data file upload purposes.
2. Download the appropriate **EdmUploadAgent.exe** file for your subscription using the links below:
   - Commercial + GCC - most commercial customers should use this
   - GCC-High - This is specifically for high security government cloud subscribers
   - DoD - this is specifically for United States Department of Defense cloud customers
3. Authorize the EDM Upload Agent, open Command Prompt window (as an administrator) and run the following command:

   ```
   EdmUploadAgent.exe /Authorize
   ```

4. If you don't have a current copy of the existing schema, you'll need to download a copy of the existing schema, run this command:

   ```
   EdmUploadAgent.exe /SaveSchema /DataStoreName <dataStoreName> [/OutputDir [Output dir location]]
   ```

5. Customize the schema so each column utilizes “caseInsensitive” and / or “ignoredDelimiters”. The default value for “caseInsensitive” is “false” and for “ignoredDelimiters”, it is an empty string.

> **NOTE**
> The underlying custom sensitive information type or built in sensitive information type used to detect the general regex pattern must support detection of the variations inputs listed with ignoredDelimiters. For example, the built in U.S. social security number (SSN) sensitive information type can detect variations in the data that include dashes, spaces, or lack of spaces between the grouped numbers that make up the SSN. As a result, the only delimiters that are relevant to include in EDM’s ignoredDelimiters for SSN data are: dash and space.

Here is a sample schema that simulates case insensitive match by creating the extra columns needed to recognize case variations in the sensitive data.

```
<EdmSchema xmlns="http://schemas.microsoft.com/office/2018/edm">
  <DataStore name="PatientRecords" description="Schema for patient records policy" version="1">
           <Field name="PolicyNumber" searchable="true" />
           <Field name="PolicyNumberLowerCase" searchable="true" />
           <Field name="PolicyNumberUpperCase" searchable="true" />
           <Field name="PolicyNumberCapitalLetters" searchable="true" />
  </DataStore>
</EdmSchema>
```

In the above example, the variations of the original `PolicyNumber` column will no longer be needed if both `caseInsensitive` and `ignoredDelimiters` are added.

To update this schema so that EDM uses configurable match use the `caseInsensitive` and `ignoredDelimiters` flags. Here's how that looks:

```
<EdmSchema xmlns="http://schemas.microsoft.com/office/2018/edm">
  <DataStore name="PatientRecords" description="Schema for patient records policy" version="1">
         <Field name="PolicyNumber" searchable="true" caseInsensitive="true" ignoredDelimiters="-,/,*,#,^"
/>
  </DataStore>
</EdmSchema>
```

The `ignoredDelimiters` flag supports any non-alphanumeric character, here are some examples:

- .
- -
- /
- _
- *
- ^
- #
- !
- ?
- [
-]
- {
- }
- \
- ~
- ;

The `ignoredDelimiters` flag doesn't support:

- characters 0-9
- A-Z
- a-z
- "
- ,

6. Connect to the Security & Compliance center using the procedures in Connect to Security & Compliance Center PowerShell.

7. Update your schema by running these cmdlets one at a time:

```
$edmSchemaXml=Get-Content .\\edm.xml -Encoding Byte -ReadCount 0
```

```
Set-DlpEdmSchema -FileData $edmSchemaXml -Confirm:$true
```

8. If necessary, update the data file to match the new schema version

> **TIP**
>
> Optionally, you can run a validation against your csv file before uploading by running:
>
> ```
> EdmUploadAgent.exe /ValidateData /DataFile [data file] [schema file]
> ```
>
> For more information on all the EdmUploadAgent.exe >supported parameters run
>
> ```
> EdmUploadAgent.exe /?
> ```

9. Open Command Prompt window (as an administrator) and run the following command to hash and upload your sensitive data:

   ```
   EdmUploadAgent.exe /UploadData /DataStoreName [DS Name] /DataFile [data file] /HashLocation [hash file location] /Salt [custom salt] /Schema [Schema file]
   ```

Related articles

- Create a custom sensitive information type with Exact Data Match based classification
- Sensitive information type-entity definitions
- Custom sensitive information types
- Overview of DLP policies
- Microsoft Cloud App Security
- New-DlpEdmSchema

Get started with custom sensitive information types

2/18/2021 • 5 minutes to read • Edit Online

If the pre-configured sensitive information types don't meet your needs, you can create your own custom sensitive information types that you fully define or you can copy one of the pre-configured ones and modify it.

The custom sensitive information types that you create by using this method are added to the rule package named `Microsoft.SCCManaged.CustomRulePack`.

There are two ways to create a new sensitive information type:

- from scratch where you fully define all elements
- copy and modify an existing sensitive information type

Before you begin

- You should be familiar with sensitive information types and what they are composed of. See, Learn about sensitive information types. It is critical to understand the roles of:
  - regular expressions - Microsoft 365 sensitive information types uses the Boost.RegEx 5.1.3 engine
  - keyword lists - you can create your own as you define your sensitive information type or choose from existing keyword lists
  - keyword dictionary
  - functions
  - confidence levels
- You must have Global admin or Compliance admin permissions to create, test, and deploy a custom sensitive information type through the UI. See About admin roles in Office 365.
- Your organization must have a subscription, such as Office 365 Enterprise, that includes Data Loss Prevention (DLP). See Messaging Policy and Compliance ServiceDescription.

> **IMPORTANT**
> Microsoft Customer Service & Support can't assist with creating custom classifications or regular expression patterns. Support engineers can provide limited support for the feature, such as, providing sample regular expression patterns for testing purposes, or assisting with troubleshooting an existing regular expression pattern that's not triggering as expected, but can't provide assurances that any custom content-matching development will fulfill your requirements or obligations.

Create a custom sensitive information type

Use this procedure to create a new sensitive information type that you fully define.

1. In the Compliance Center, go to **Data classification** > **Sensitive info types** and choose **Create info type**.
2. Fill in values for **Name** and **Description** and choose **Next**.
3. Choose **Create pattern**. You can create multiple patterns, each with different elements and confidence levels, as you define your new sensitive information type.
4. Choose the default confidence level for the pattern. The values are **Low confidence**, **Medium confidence**, and **High confidence**.
5. Choose and define **Primary element**. The primary element can be a **Regular expression** with an optional

validator, a **Keyword list**, a **Keyword dictionary**, or one of the pre-configured **Functions**. For more information on DLP functions, see What the DLP functions look for.

6. Fill in a value for **Character proximity**.
7. (Optional) Add supporting elements if you have any. Supporting elements can be a regular expression with an optional validator, a keyword list, a keyword dictionary or one of the pre-defined functions.
8. (Optional) Add additional checks from the list of available checks
9. Choose **Create**.
10. Choose **Next**.
11. Choose the **recommended confidence level** for this sensitive information type.
12. Check your setting and choose **Submit**.

> **IMPORTANT**
> Microsoft 365 uses the search crawler to identify and classify sensitive information in SharePoint Online and OneDrive for Business sites. To identify your new custom sensitive information type in existing content, the content must be re-crawled. Content is crawled based on a schedule, but you can manually re-crawl content for a site collection, list, or library. For more information, see Manually request crawling and re-indexing of a site, a library or a list.

13. On the **Data classification** page, you'll see all the sensitive information types listed. Choose **Refresh** and then browse for or use the search tool to find the sensitive information type you just created.

Test a sensitive information type

You can test any sensitive information type in the list. We suggest that you test every sensitive information type that you create before using it in a policy.

1. Prepare two files, like a Word document. One with content that matches the elements you specified in your sensitive information type and one that doesn't match.
2. In the Compliance Center, go to **Data classification** > **Sensitive info types** and choose the sensitive information type from the list to open the details pane and choose **Test**.
3. Upload a file and choose **Test**.
4. On the **Matches results** page, review the results and choose **Finish**.

Modify custom sensitive information types in the Compliance Center

1. In the Compliance Center, go to **Data classification** > **Sensitive info types** and choose the sensitive information type from the list that you want to modify choose **Edit**.
2. You can add other patterns, with unique primary and supporting elements, confidence levels, character proximity, and additional checks or edit/remove the existing ones. For more information, see Create a custom sensitive information type.

Remove custom sensitive information types in the Compliance Center

> **NOTE**
> You can only remove custom sensitive information types; you can't remove built-in sensitive information types.

IMPORTANT

Before your remove a custom sensitive information type, verify that no DLP policies or Exchange mail flow rules (also known as transport rules) still reference the sensitive information type.

1. In the Compliance Center, go to **Data classification** > **Sensitive info types** and choose the sensitive information type from the list that you want to remove.
2. In the fly-out that opens, choose **Delete**.

Copy and modify a sensitive information type

Use this procedure to create a new sensitive information type that is based on an existing sensitive information type.

1. In the Compliance Center, go to **Data classification** > **Sensitive info types** and choose the sensitive information type that you want to copy.
2. In the flyout, choose **Copy**.
3. Choose **Refresh** in the list of sensitive information types and either browse or search for the copy you just made. Partial sting searches work, so you could just search for `copy` and search would return all the sensitive information types with the word `copy` in the name.
4. Fill in values for **Name** and **Description** and choose **Next**.
5. Choose your sensitive information type copy and choose **Edit**.
6. Give your new sensitive information type a new **Name** and **Description**.
7. You can choose to edit or remove the existing patterns and add new ones. Choose the default confidence level for the new pattern. The values are **Low confidence**, **Medium confidence**, and **High confidence**.
8. Choose and define **Primary element**. The primary element can be a **Regular expression**, a **Keyword list**, a **Keyword dictionary**, or one of the pre-configured **Functions**. See, What the DLP functions look for.
9. Fill in a value for **Character proximity**.
10. (Optional) If you have **Supporting elements** or any **Additional checks** add them. If needed you can group your **Supporting elements**.
11. Choose **Create**.
12. Choose **Next**.
13. Choose the **recommended confidence level** for this sensitive information type.
14. Check your setting and choose **Submit**.

You can also create custom sensitive information types by using PowerShell and Exact Data Match capabilities. To learn more about those methods, see:

- Create a custom sensitive information type in Security & Compliance Center PowerShell
- Create a custom sensitive information type for DLP with Exact Data Match (EDM)

NOTE

Microsoft 365 Information Protection supports, in preview, double byte character set languages for:

- Chinese (simplified)
- Chinese (traditional)
- Korean
- Japanese

This support is available for sensitive information types. See, Information protection support for double byte character sets release notes (preview) for more information.

Create a custom sensitive information type using PowerShell

2/18/2021 • 27 minutes to read • Edit Online

This topic shows you how to use PowerShell to create an XML *rule package* file that defines your own custom sensitive information types. You need to know how to create a regular expression. As an example, this topic creates a custom sensitive information type that identifies an employee ID. You can use this example XML as a starting point for your own XML file. If you are new to sensitive information types, see Learn about sensitive information types.

After you've created a well-formed XML file, you can upload it to Microsoft 365 by using Microsoft 365 PowerShell. Then you're ready to use your custom sensitive information type in your policies and test that it's detecting the sensitive information as you intended.

> **NOTE**
> If you don't need the fine grained control that PowerShell provides, you can create custom sensitive information types in the Compliance center. For more information, see Create a custom sensitive information type.

Important disclaimer

Due to the variances in customer environments and content match requirements, Microsoft Support cannot assist in providing custom content-matching definitions; e.g., defining custom classifications or regular expression (also known as RegEx) patterns. For custom content-matching development, testing, and debugging, Microsoft 365 customers will need to rely upon internal IT resources, or use an external consulting resource such as Microsoft Consulting Services (MCS). Support engineers can provide limited support for the feature, but cannot provide assurances that any custom content-matching development will fulfill the customer's requirements or obligations. As an example of the type of support that can be provided, sample regular expression patterns may be provided for testing purposes. Or, support can assist with troubleshooting an existing RegEx pattern which is not triggering as expected with a single specific content example.

See Potential validation issues to be aware of in this topic.

For more information about the Boost.RegEx (formerly known as RegEx++) engine that's used for processing the text, see Boost.Regex 5.1.3.

Sample XML of a rule package

Here's the sample XML of the rule package that we'll create in this topic. Elements and attributes are explained in the sections below.

```
<?xml version="1.0" encoding="UTF-16"?>
<RulePackage xmlns="http://schemas.microsoft.com/office/2011/mce">
<RulePack id="DAD86A92-AB18-43BB-AB35-96F7C594ADAA">
 <Version build="0" major="1" minor="0" revision="0"/>
 <Publisher id="619DD8C3-7B80-4998-A312-4DF0402BAC04"/>
 <Details defaultLangCode="en-us">
  <LocalizedDetails langcode="en-us">
   <PublisherName>Contoso</PublisherName>
   <Name>Employee ID Custom Rule Pack</Name>
   <Description>
    This rule package contains the custom Employee ID entity
```

```
    This rule package contains the custom Employee ID entity.
    </Description>
   </LocalizedDetails>
  </Details>
</RulePack>
<Rules>
<!-- Employee ID -->
 <Entity id="E1CC861E-3FE9-4A58-82DF-4BD259EAB378" patternsProximity="300" recommendedConfidence="70">
  <Pattern confidenceLevel="60">
   <IdMatch idRef="Regex_employee_id"/>
  </Pattern>
  <Pattern confidenceLevel="70">
   <IdMatch idRef="Regex_employee_id"/>
   <Match idRef="Func_us_date"/>
  </Pattern>
  <Pattern confidenceLevel="80">
   <IdMatch idRef="Regex_employee_id"/>
   <Match idRef="Func_us_date"/>
   <Any minMatches="1">
    <Match idRef="Keyword_badge" minCount="2"/>
    <Match idRef="Keyword_employee"/>
   </Any>
   <Any minMatches="0" maxMatches="0">
    <Match idRef="Keyword_false_positives_local"/>
    <Match idRef="Keyword_false_positives_intl"/>
   </Any>
  </Pattern>
 </Entity>
 <Regex id="Regex_employee_id">(\s)(\d{9})(\s)</Regex>
 <Keyword id="Keyword_employee">
  <Group matchStyle="word">
   <Term>Identification</Term>
   <Term>Contoso Employee</Term>
  </Group>
 </Keyword>
 <Keyword id="Keyword_badge">
  <Group matchStyle="string">
   <Term>card</Term>
   <Term>badge</Term>
   <Term caseSensitive="true">ID</Term>
  </Group>
 </Keyword>
 <Keyword id="Keyword_false_positives_local">
  <Group matchStyle="word">
   <Term>credit card</Term>
   <Term>national ID</Term>
  </Group>
 </Keyword>
 <Keyword id="Keyword_false_positives_intl">
  <Group matchStyle="word">
   <Term>identity card</Term>
   <Term>national ID</Term>
   <Term>EU debit card</Term>
  </Group>
 </Keyword>
 <LocalizedStrings>
  <Resource idRef="E1CC861E-3FE9-4A58-82DF-4BD259EAB378">
   <Name default="true" langcode="en-us">Employee ID</Name>
   <Description default="true" langcode="en-us">
   A custom classification for detecting Employee IDs.
   </Description>
   <Description default="false" langcode="de-de">
   Description for German locale.
   </Description>
  </Resource>
 </LocalizedStrings>
</Rules>
</RulePackage>
```

What are your key requirements? [Rule, Entity, Pattern elements]

Before you get started, it's helpful to understand the basic structure of the XML schema for a rule, and how you can use this structure to define your custom sensitive information type so that it will identify the right content.

A rule defines one or more entities (sensitive information types), and each entity defines one or more patterns. A pattern is what a policy looks for when it evaluates content such as email and documents.

In this topic, the XML markup uses rule to mean the patterns that define an entity, also known as a sensitive information type. So in this topic, when you see rule, think entity or sensitive information type, not conditions and actions.

Simplest scenario: entity with one pattern

Here's the simplest scenario. You want your policy to identify content that contains your organization's employee ID, which is formatted as a nine-digit number. So the pattern refers to a regular expression contained in the rule that identifies nine-digit numbers. Any content containing a nine-digit number satisfies the pattern.

However, while simple, this pattern may identify many false positives by matching content that contains any nine-digit number that is not necessarily an employee ID.

More common scenario: entity with multiple patterns

For this reason, it's more common to define an entity by using more than one pattern, where the patterns identify supporting evidence (such as a keyword or date) in addition to the entity (such as a nine-digit number).

For example, to increase the likelihood of identifying content that contains an employee ID, you can define another pattern that also identifies a hire date, and define yet another pattern that identifies both a hire date and a keyword (such as "employee ID"), in addition to the nine-digit number.

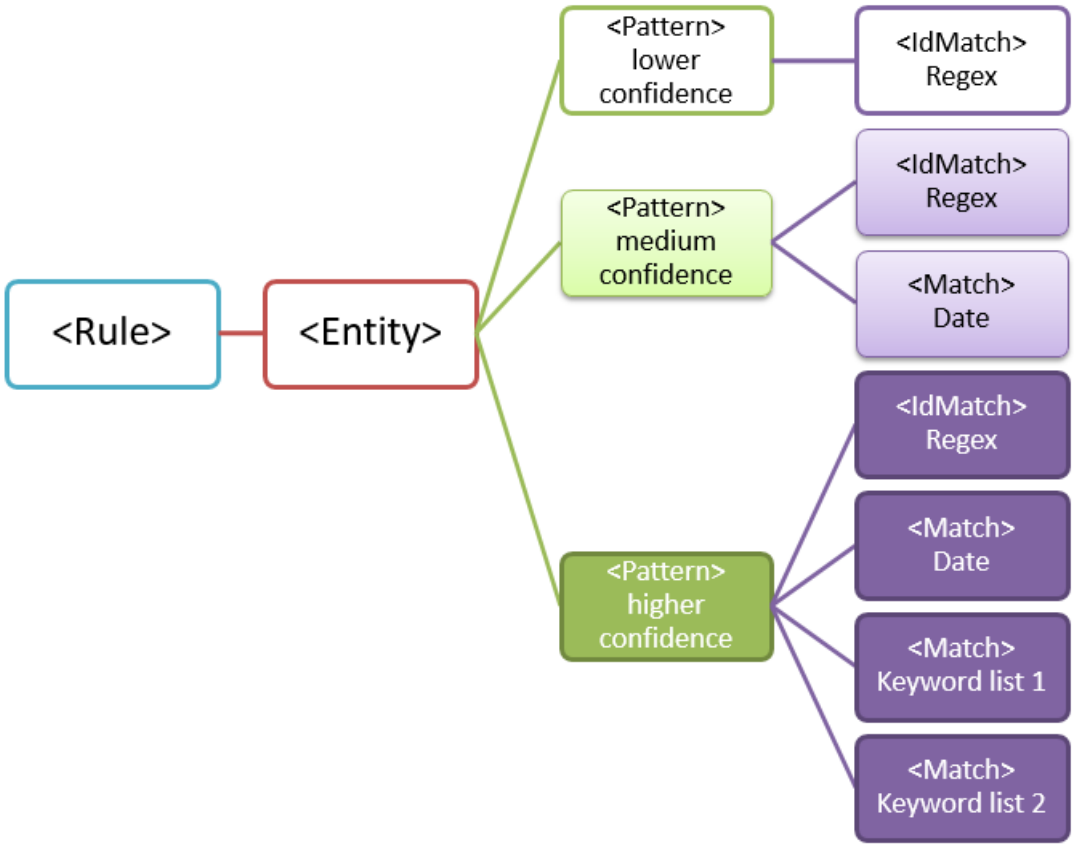

Note a couple of important aspects of this structure:

- Patterns that require more evidence have a higher confidence level. This is useful because when you later use this sensitive information type in a policy, you can use more restrictive actions (such as block content) with only the higher-confidence matches, and you can use less restrictive actions (such as send notification) with the lower-confidence matches.
- The supporting IdMatch and Match elements reference regexes and keywords that are actually children of the Rule element, not the Pattern. These supporting elements are referenced by the Pattern but included in the Rule. This means that a single definition of a supporting element, like a regular expression or a keyword list, can be referenced by multiple entities and patterns.

What entity do you need to identify? [Entity element, id attribute]

An entity is a sensitive information type, such as a credit card number, that has a well-defined pattern. Each entity has a unique GUID as its ID.

Name the entity and generate its GUID

1. In your XML editor of choice, add the Rules and Entity elements.
2. Add a comment that contains the name of your custom entity — in this example, Employee ID. Later, you'll add the entity name to the localized strings section, and that name is what appears in the UI when you create a policy.
3. Generate a GUID for your entity. There are several ways to generate GUIDs, but you can do it easily in PowerShell by typing **[guid]::NewGuid()**. Later, you'll also add the entity GUID to the localized strings section.

```
<Rules>
  <!-- Employee ID -->
  <Entity id="E1CC861E-3FE9-4A58-82DF-4BD259EAB378">
  </Entity>
</Rules>
```

Add comment and name entity.

Generate GUID and add as id attribute.

What pattern do you want to match? [Pattern element, IdMatch element, Regex element]

The pattern contains the list of what the sensitive information type is looking for. This can include regexes, keywords, and built-in functions (which perform tasks like running regexes to find dates or addresses). Sensitive information types can have multiple patterns with unique confidences.

What all of the below patterns have in common is that they all reference the same regular expression, which looks for a nine-digit number (\d{9}) surrounded by white space (\s) ... (\s). This regular expression is referenced by the IdMatch element and is the common requirement for all patterns that look for the Employee ID entity. IdMatch is the identifier that the pattern is to trying to match, such as Employee ID or credit card number or social security number. A Pattern element must have exactly one IdMatch element.

```
<Rules>

  <!-- Employee ID -->
  <Entity id="E1CC861E-3FE9-4A58-82DF-4BD259EAB378">

    <Pattern>
      <IdMatch idRef="Regex_employee_id" />
    </Pattern>

    <Pattern>
      <IdMatch idRef="Regex_employee_id" />
    </Pattern>

    <Pattern>
      <IdMatch idRef="Regex_employee_id" />
    </Pattern>

  </Entity>

  <Regex id="Regex_employee_id">(\s)(\d{9})(\s)</Regex>

</Rules>
```

Each Pattern has exactly one required IdMatch element, which references the same regex.

Regex defines the identifier the pattern is trying to match.

When satisfied, a pattern returns a count and confidence level, which you can use in the conditions in your policy. When you add a condition for detecting a sensitive information type to a policy, you can edit the count and confidence level as shown here. Confidence level (also called match accuracy) is explained later in this topic.

When content contains sensitive information *

Sensitive information type	Instance count min	Instance count max	Match accuracy min	Match accuracy max
U.S. Individual Taxpayer Identification Number (ITIN)	1	9	75	100
U.S. Social Security Number (SSN)	1	9	75	100
U.S. / U.K. Passport Number	1	9	75	100

When you create your regular expression, keep in mind that there are potential issues to be aware of. For example, if you write and upload a regex that identifies too much content, this can impact performance. To learn more about these potential issues, see the later section Potential validation issues to be aware of.

Do you want to require additional evidence? [Match element,

minCount attribute]

In addition to the IdMatch, a pattern can use the Match element to require additional supporting evidence, such as a keyword, regex, date, or address.

A Pattern can include multiple Match elements; they can be included directly in the Pattern element or combined by using the Any element. Match elements are joined by an implicit AND operator; all Match elements must be satisfied for the pattern to be matched. You can use the Any element to introduce AND or OR operators (more on that in a later section).

You can use the optional minCount attribute to specify how many instances of a match need to be found for each of the Match elements. For example, you can specify that a pattern is satisfied only when at least two keywords from a keyword list are found.


```
<Rules>

  <!-- Employee ID -->
  <Entity id="E1CC861E-3FE9-4A58-82DF-4BD259EAB378">

    <Pattern>
      <IdMatch idRef="Regex_employee_id" />
    </Pattern>

    <Pattern>
      <IdMatch idRef="Regex_employee_id" />
    </Pattern>

    <Pattern>
      <IdMatch idRef="Regex_employee_id" />
      <Match idRef="Keyword_employee" />
      <Match idRef="Keyword_badge" minCount="2" />
    </Pattern>

  </Entity>

  <Regex id="Regex_employee_id">(\s)(\d{9})(\s)</Regex>
```

Keywords [Keyword, Group, and Term elements, matchStyle and caseSensitive attributes]

When you identify sensitive information, like an employee ID, you often want to require keywords as corroborative evidence. For example, in addition to matching a nine-digit number, you may want to look for words like "card", "badge", or "ID". To do this, you use the Keyword element. The Keyword element has an ID attribute that can be referenced by multiple Match elements in multiple patterns or entities.

Keywords are included as a list of Term elements in a Group element. The Group element has a matchStyle attribute with two possible values:

- **matchStyle="word"** Word match identifies whole words surrounded by white space or other delimiters. You should always use word unless you need to match parts of words or match words in Asian languages.
- **matchStyle="string"** String match identifies strings no matter what they're surrounded by. For example, "id" will match "bid" and "idea". Use string only when you need to match Asian words or if your keyword may be included as part of other strings.

Finally, you can use the caseSensitive attribute of the Term element to specify that the content must match the keyword exactly, including lower- and upper-case letters.

```
<Rules>

    <!-- Employee ID -->
    <Entity id="E1CC861E-3FE9-4A58-82DF-4BD259EAB378" >

        <Pattern>
            <IdMatch idRef="Regex_employee_id"/>
        </Pattern>

        <Pattern>
            <IdMatch idRef="Regex_employee_id"/>
        </Pattern>

        <Pattern>
            <IdMatch idRef="Regex_employee_id"/>
            <Match idRef="Keyword_badge"/>
            <Match idRef="Keyword_employee"/>
        </Pattern>

    </Entity>
    <Regex id="Regex_employee_id">(\s)(\d{9})(\s)</Regex>

    <Keyword id="Keyword_employee">
        <Group matchStyle="word">
            <Term>Identification</Term>
            <Term>Contoso Employee</Term>
        </Group>
    </Keyword>

    <Keyword id="Keyword_badge">
        <Group matchStyle="string">
            <Term>card</Term>
            <Term>badge</Term>
            <Term caseSensitive="true">ID</Term>
        </Group>
    </Keyword>

</Rules>
```

Match element uses idRef to reference Keyword element defined outside entity.

matchStyle can be either word or string.

A Term can be made case sensitive.

Regular expressions [Regex element]

In this example, the employee ID entity already uses the IdMatch element to reference a regex for the pattern — a nine-digit number surrounded by whitespace. In addition, a pattern can use a Match element to reference an additional Regex element to identify corroborative evidence, such as a five- or nine-digit number in the format of a US zip code.

Additional patterns such as dates or addresses [built-in functions]

In addition to the built-in sensitive information types, sensitive information types can also use built-in functions that can identify corroborative evidence such as a US date, EU date, expiration date, or US address. Microsoft 365 does not support uploading your own custom functions, but when you create a custom sensitive information type, your entity can reference the built-in functions.

For example, an employee ID badge has a hire date on it, so this custom entity can use the built-in function `Func_us_date` to identify a date in the format commonly used in the US.

For more information, see What the DLP functions look for.

```
<Rules>

  <!-- Employee ID -->
  <Entity id="E1CC861E-3FE9-4A58-82DF-4BD259EAB378">

    <Pattern>
      <IdMatch idRef="Regex_employee_id" />
    </Pattern>

    <Pattern>
      <IdMatch idRef="Regex_employee_id" />
      <Match idRef="Func_us_date" />
    </Pattern>

    <Pattern>
      <IdMatch idRef="Regex_employee_id" />
      <Match idRef="Keyword_employee" />
      <Match idRef="Keyword_badge" />
    </Pattern>

  </Entity>

  <Regex id="Regex_employee_id">(\s)(\d{9})(\s)</Regex>

  <Keyword id="Keyword_employee">
    <Group matchStyle="word">
      <Term>Identification</Term>
      <Term>Contoso Employee</Term>
    </Group>
  </Keyword>

  <Keyword id="Keyword_badge">
    <Group matchStyle="string">
      <Term>card</Term>
      <Term>badge</Term>
      <Term>ID</Term>
    </Group>
  </Keyword>

</Rules>
```

Patterns can reference the built-in DLP functions to find corroborative evidence.

Because the functions are built in, they don't require a separate element in the Rules element.

Different combinations of evidence [Any element, minMatches and maxMatches attributes]

In a Pattern element, all IdMatch and Match elements are joined by an implicit AND operator — all of the matches must be satisfied before the pattern can be satisfied. However, you can create more flexible matching logic by using the Any element to group Match elements. For example, you can use the Any element to match all, none, or an exact subset of its children Match elements.

The Any element has optional minMatches and maxMatches attributes that you can use to define how many of the children Match elements must be satisfied before the pattern is matched. Note that these attributes define the number of Match elements that must be satisfied, not the number of instances of evidence found for the matches. To define a minimum number of instances for a specific match, such as two keywords from a list, use the minCount attribute for a Match element (see above).

Match at least one child Match element

If you want to require that only a minimum number of Match elements must be met, you can use the minMatches attribute. In effect, these Match elements are joined by an implicit OR operator. This Any element is satisfied if a US-formatted date or a keyword from either list is found.

```
<Any minMatches="1" >
     <Match idRef="Func_us_date" />
     <Match idRef="Keyword_employee" />
     <Match idRef="Keyword_badge" />
</Any>
```

Match an exact subset of any children Match elements

If you want to require that an exact number of Match elements must be met, you can set minMatches and maxMatches to the same value. This Any element is satisfied only if exactly one date or keyword is found — any more than that, and the pattern won't be matched.

```
<Any minMatches="1" maxMatches="1" >
     <Match idRef="Func_us_date" />
     <Match idRef="Keyword_employee" />
     <Match idRef="Keyword_badge" />
</Any>
```

Match none of children Match elements

If you want to require the absence of specific evidence for a pattern to be satisfied, you can set both minMatches and maxMatches to 0. This can be useful if you have a keyword list or other evidence that are likely to indicate a false positive.

For example, the employee ID entity looks for the keyword "card" because it might refer to an "ID card". However, if card appears only in the phrase "credit card", "card" in this content is unlikely to mean "ID card". So you can add "credit card" as a keyword to a list of terms that you want to exclude from satisfying the pattern.

```
<Any minMatches="0" maxMatches="0" >
    <Match idRef="Keyword_false_positives_local" />
    <Match idRef="Keyword_false_positives_intl" />
</Any>
```

Match a number of unique terms

If you want to match a number of unique terms, use the *uniqueResults* parameter, set to *true*, as shown in the following example:

```
<Pattern confidenceLevel="75">
    <IdMatch idRef="Salary_Revision_terms" />
    <Match idRef=" Salary_Revision_ID " minCount="3" uniqueResults="true" />
</Pattern>
```

In this example, a pattern is defined for salary revision using at least three unique matches.

How close to the entity must the other evidence be? [patternsProximity attribute]

Your sensitive information type is looking for a pattern that represents an employee ID, and as part of that pattern it's also looking for corroborative evidence like a keyword such as "ID". It makes sense that the closer together this evidence is, the more likely the pattern is to be an actual employee ID. You can determine how close other evidence in the pattern must be to the entity by using the required patternsProximity attribute of the Entity element.

```
<!-- Employee ID -->
<Entity id="E1CC861E-3FE9-4A58-82DF-4BD259EAB378" patternsProximity="300" >

  <Pattern confidenceLevel="75">
    <IdMatch idRef="Regex_employee_id" />
    <Match idRef="Keyword_employee" />
  </Pattern>

</Entity>
```

Attribute value defines the distance in Unicode characters from IdMatch to all other Match elements in the Pattern.

For each pattern in the entity, the patternsProximity attribute value defines the distance (in Unicode characters) from the IdMatch location for all other Matches specified for that Pattern. The proximity window is anchored by the IdMatch location, with the window extending to the left and right of the IdMatch.

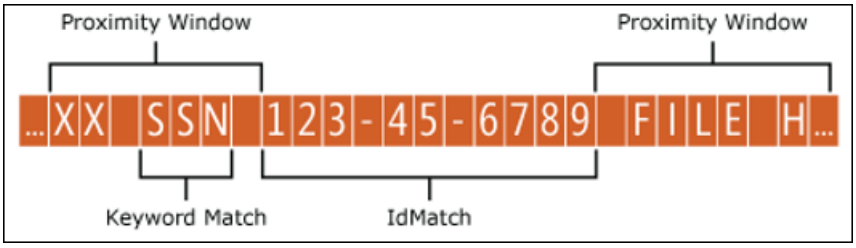

The example below illustrates how the proximity window affects the pattern matching where IdMatch element for the employee ID custom entity requires at least one corroborating match of keyword or date. Only ID1 matches because for ID2 and ID3, either no or only partial corroborating evidence is found within the proximity window.

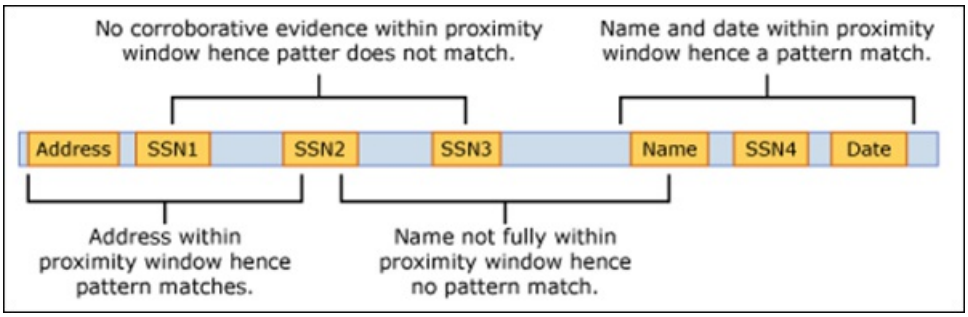

Note that for email, the message body and each attachment are treated as separate items. This means that the proximity window does not extend beyond the end of each of these items. For each item (attachment or body), both the idMatch and corroborative evidence needs to reside in that item.

What are the right confidence levels for different patterns? [confidenceLevel attribute, recommendedConfidence attribute]

The more evidence that a pattern requires, the more confidence you have that an actual entity (such as employee ID) has been identified when the pattern is matched. For example, you have more confidence in a pattern that requires a nine-digit ID number, hire date, and keyword in close proximity, than you do in a pattern that requires only a nine-digit ID number.

The Pattern element has a required confidenceLevel attribute. You can think of the value of confidenceLevel (an integer between 1 and 100) as a unique ID for each pattern in an entity — the patterns in an entity must have different confidence levels that you assign. The precise value of the integer doesn't matter — simply pick numbers that make sense to your compliance team. After you upload your custom sensitive information type and then create a policy, you can reference these confidence levels in the conditions of the rules that you create.

```
<!-- Employee ID -->
<Entity id="E1CC861E-3FE9-4A58-82DF-4BD259EAB378" patternsProximity="300">

  <Pattern confidenceLevel="60">
    <IdMatch idRef="Regex_employee_id" />
  </Pattern>

  <Pattern confidenceLevel="70">
    <IdMatch idRef="Regex_employee_id" />
    <Match idRef="Func_us_date" />
  </Pattern>

  <Pattern confidenceLevel="80">
    <IdMatch idRef="Regex_employee_id" />
    <Match idRef="Func_us_date" />
    <Any minMatches="1" >
        <Match idRef="Keyword_badge" />
        <Match idRef="Keyword_employee" />
    </Any>
    <Any maxMatches="0" >
        <Match idRef="Keyword_false_positives_local" />
        <Match idRef="Keyword_false_positives_intl" />
    </Any>
  </Pattern>

</Entity>
```

The more evidence the Pattern requires, the higher the confidence level.

In addition to confidenceLevel for each Pattern, the Entity has a recommendedConfidence attribute. The recommended confidence attribute can be thought of as the default confidence level for the rule. When you create a rule in a policy, if you don't specify a confidence level for the rule to use, that rule will match based on the recommended confidence level for the entity. Please note that the recommendedConfidence attribute is mandatory for each Entity ID in the Rule Package, if missing you won't be able to save policies that use the Sensitive Information Type.

Do you want to support other languages in the UI of the Compliance center? [LocalizedStrings element]

If your compliance team uses the Microsoft 365 Compliance center to create polices policies in different locales and in different languages, you can provide localized versions of the name and description of your custom sensitive information type. When your compliance team uses Microsoft 365 in a language that you support, they'll see the localized name in the UI.

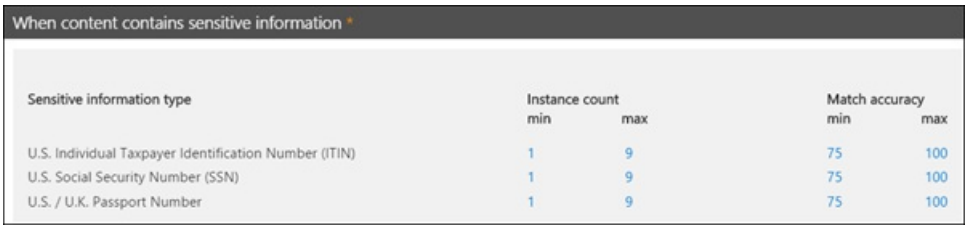

The Rules element must contain a LocalizedStrings element, which contains a Resource element that references the GUID of your custom entity. In turn, each Resource element contains one or more Name and Description elements that each use the langcode attribute to provide a localized string for a specific language.

```
  <LocalizedStrings>
    <Resource idRef="E1CC861E-3FE9-4A58-82DF-4BD259EAB378">
      <Name default="true" langcode="en-us">Employee ID</Name>
      <Description default="true" langcode="en-us">
        A custom classification for detecting Employee IDs.
      </Description>
      <Name default="true" langcode="de-de">Name for German locale</Name>
      <Description default="true" langcode="de-de">
        Description for German locale.
      </Description>
    </Resource>

  </LocalizedStrings>
 </Rules>
</RulePackage>
```

Resource element references GUID of the entity.

Name and Description elements use langcode attribute to specify strings for a locale.

Note that you use localized strings only for how your custom sensitive information type appears in the UI of the Compliance center. You can't use localized strings to provide different localized versions of a keyword list or regular expression.

Other rule package markup [RulePack GUID]

Finally, the beginning of each RulePackage contains some general information that you need to fill in. You can use the following markup as a template and replace the ". . ." placeholders with your own info.

Most importantly, you'll need to generate a GUID for the RulePack. Above, you generated a GUID for the entity; this is a second GUID for the RulePack. There are several ways to generate GUIDs, but you can do it easily in PowerShell by typing [guid]::NewGuid().

The Version element is also important. When you upload your rule package for the first time, Microsoft 365 notes the version number. Later, if you update the rule package and upload a new version, make sure to update the version number or Microsoft 365 won't deploy the rule package.

```
<?xml version="1.0" encoding="utf-16"?>
<RulePackage xmlns="http://schemas.microsoft.com/office/2011/mce">
  <RulePack id=". . .">
    <Version major="1" minor="0" build="0" revision="0" />
    <Publisher id=". . ." />
    <Details defaultLangCode=". . .">
      <LocalizedDetails langcode=" . . . ">
         <PublisherName>. . .</PublisherName>
         <Name>. . .</Name>
         <Description>. . .</Description>
      </LocalizedDetails>
    </Details>
  </RulePack>

  <Rules>
  . . .
  </Rules>
</RulePackage>
```

When complete, your RulePack element should look like this.

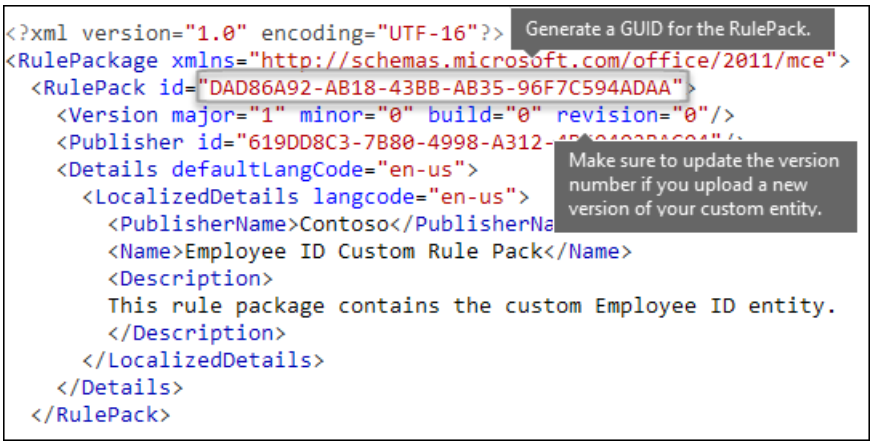


```
<?xml version="1.0" encoding="UTF-16"?>
<RulePackage xmlns="http://schemas.microsoft.com/office/2011/mce">
  <RulePack id="DAD86A92-AB18-43BB-AB35-96F7C594ADAA">
    <Version major="1" minor="0" build="0" revision="0"/>
    <Publisher id="619DD8C3-7B80-4998-A312-
    <Details defaultLangCode="en-us">
      <LocalizedDetails langcode="en-us">
        <PublisherName>Contoso</PublisherNa
        <Name>Employee ID Custom Rule Pack</Name>
        <Description>
        This rule package contains the custom Employee ID entity.
        </Description>
      </LocalizedDetails>
    </Details>
  </RulePack>
```

Changes for Exchange Online

Previously, you might have used Exchange Online PowerShell to import your custom sensitive information types for DLP. Now your custom sensitive information types can be used in both the Exchange admin center and the Compliance center. As part of this improvement, you should use Compliance center PowerShell to import your custom sensitive information types — you can't import them from the Exchange PowerShell anymore. Your custom sensitive information types will continue to work just like before; however, it may take up to one hour for changes made to custom sensitive information types in the Compliance center to appear in the Exchange admin center.

Note that in the Compliance center, you use the **New-DlpSensitiveInformationTypeRulePackage** cmdlet to upload a rule package. (Previously, in the Exchange admin center, you used the **ClassificationRuleCollection**` cmdlet.)

Upload your rule package

To upload your rule package, do the following steps:

1. Save it as an .xml file with Unicode encoding.
2. Connect to Compliance center PowerShell
3. Use the following syntax:

   ```
   New-DlpSensitiveInformationTypeRulePackage -FileData (Get-Content -Path "PathToUnicodeXMLFile" -
   Encoding Byte -ReadCount 0)
   ```

   This example uploads the Unicode XML file named MyNewRulePack.xml from C:\My Documents.

   ```
   New-DlpSensitiveInformationTypeRulePackage -FileData (Get-Content -Path "C:\My
   Documents\MyNewRulePack.xml" -Encoding Byte -ReadCount 0)
   ```

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

   > **NOTE**
   > The maximum number of rule packages supported is 10, but each package can contain the definition of multiple sensitive information types.

4. To verify that you've successfully created a new sensitive information type, do any of the following steps:

   - Run the Get-DlpSensitiveInformationTypeRulePackage cmdlet to verify the new rule package is

listed:

```
Get-DlpSensitiveInformationTypeRulePackage
```

- Run the Get-DlpSensitiveInformationType cmdlet to verify the sensitive information type is listed:

  ```
  Get-DlpSensitiveInformationType
  ```

  For custom sensitive information types, the Publisher property value will be something other than Microsoft Corporation.

- Replace <Name> with the Name value of the sensitive information type (example: Employee ID) and run the Get-DlpSensitiveInformationType cmdlet:

  ```
  Get-DlpSensitiveInformationType -Identity "<Name>"
  ```

Potential validation issues to be aware of

When you upload your rule package XML file, the system validates the XML and checks for known bad patterns and obvious performance issues. Here are some known issues that the validation checks for — a regular expression:

- Cannot begin or end with alternator "|", which matches everything because it's considered an empty match.

  For example, "|a" or "b|" will not pass validation.

- Cannot begin or end with a ".{0,m}" pattern, which has no functional purpose and only impairs performance.

  For example, ".{0,50}ASDF" or "ASDF.{0,50}" will not pass validation.

- Cannot have ".{0,m}" or ".{1,m}" in groups, and cannot have ".*" or ".+" in groups.

  For example, "(.{0,50000})" will not pass validation.

- Cannot have any character with "{0,m}" or "{1,m}" repeaters in groups.

  For example, "(a*)" will not pass validation.

- Cannot begin or end with ".{1,m}"; instead, use just "."

  For example, ".{1,m}asdf" will not pass validation; instead, use just ".asdf".

- Cannot have an unbounded repeater (such as "*" or "+") on a group.

  For example, "(xx)*" and "(xx)+" will not pass validation.

- Keywords have a maximum of 50 characters in Length. If you have a keyword within a Group exceeding this, a suggested solution is to create the Group of terms as a Keyword Dictionary and reference the GUID of the Keyword Dictionary within the XML structure as part of the Entity for Match or idMatch in the file.

- Each Custom Sensitive Information Type can have a maximum of 2048 keywords total.

- The maximum size of Keyword Dictionaries in a single tenant is 100 kilobytes compressed. Reference the same dictionary as many times as necessary when creating custom sensitive information types. Start with creating custom keyword lists in the sensitive information type and use keyword dictionaries if you have more than 2048 keywords in a keyword list or a keyword is larger than 50 characters in length.

- Ensure each Entity element contains a recommendedConfidence attribute.
- When using the PowerShell Cmdlet there is a maximum return size of the Deserialized Data of approximately 1 megabyte. This will affect the size of your rule pack XML file. Keep the uploaded file limited to a 770 kilobyte maximum as a suggested limit for consistent results without error when processing.
- The XML structure does not require formatting characters such as spaces, tabs, or carriage return/linefeed entries. Take note of this when optimizing for space on uploads. Tools such as Microsoft Visual Code provide join line features to compact the XML file.

If a custom sensitive information type contains an issue that may affect performance, it won't be uploaded and you may see one of these error messages:

- **Generic quantifiers which match more content than expected (e.g., '+', '*')**
- **Lookaround assertions**
- **Complex grouping in conjunction with general quantifiers**

Recrawl your content to identify the sensitive information

Microsoft 365 uses the search crawler to identify and classify sensitive information in site content. Content in SharePoint Online and OneDrive for Business sites is recrawled automatically whenever it's updated. But to identify your new custom type of sensitive information in all existing content, that content must be recrawled.

In Microsoft 365, you can't manually request a recrawl of an entire tenant, but you can do this for a site collection, list, or library — see Manually request crawling and re-indexing of a site, a library or a list.

Remove a custom sensitive information type

> **NOTE**
> Before your remove a custom sensitive information type, verify that no DLP policies or Exchange mail flow rules (also known as transport rules) still reference the sensitive information type.

In Compliance center PowerShell, there are two methods to remove custom sensitive information types:

- **Remove individual custom sensitive information types**: Use the method documented in Modify a custom sensitive information type. You export the custom rule package that contains the custom sensitive information type, remove the sensitive information type from the XML file, and import the updated XML file back into the existing custom rule package.
- **Remove a custom rule package and all custom sensitive information types that it contains**: This method is documented in this section.

1. Connect to Compliance center PowerShell
2. To remove a custom rule package, use the Remove-DlpSensitiveInformationTypeRulePackage cmdlet:

   ```
   Remove-DlpSensitiveInformationTypeRulePackage -Identity "RulePackageIdentity"
   ```

   You can use the Name value (for any language) or the `RulePack id` (GUID) value to identify the rule package.

   This example removes the rule package named "Employee ID Custom Rule Pack".

```
Remove-DlpSensitiveInformationTypeRulePackage -Identity "Employee ID Custom Rule Pack"
```

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

3. To verify that you've successfully removed a custom sensitive information type, do any of the following steps:

   - Run the Get-DlpSensitiveInformationTypeRulePackage cmdlet and verify the rule package is no longer listed:

     ```
     Get-DlpSensitiveInformationTypeRulePackage
     ```

   - Run the Get-DlpSensitiveInformationType cmdlet to verify the sensitive information types in the removed rule package are no longer listed:

     ```
     Get-DlpSensitiveInformationType
     ```

     For custom sensitive information types, the Publisher property value will be something other than Microsoft Corporation.

   - Replace <Name> with the Name value of the sensitive information type (for example, Employee ID) and run the Get-DlpSensitiveInformationType cmdlet to verify the sensitive information type is no longer listed:

     ```
     Get-DlpSensitiveInformationType -Identity "<Name>"
     ```

Modify a custom sensitive information type

In Compliance center PowerShell, modifying a custom sensitive information type requires you to:

1. Export the existing rule package that contains the custom sensitive information type to an XML file (or use the existing XML file if you have it).
2. Modify the custom sensitive information type in the exported XML file.
3. Import the updated XML file back into the existing rule package.

To connect to Compliance Center PowerShell, see Connect to Compliance Center PowerShell.

Step 1: Export the existing rule package to an XML file

> **NOTE**
> If you have a copy of the XML file (for example, you just created and imported it), you can skip to the next step to modify the XML file.

1. If you don't already know it, run the Get-DlpSensitiveInformationTypeRulePackage cmdlet to find the name of the custom rule package:

```
Get-DlpSensitiveInformationTypeRulePackage
```

> **NOTE**
> The built-in rule package that contains the built-in sensitive information types is named Microsoft Rule Package. The rule package that contains the custom sensitive information types that you created in the Compliance center UI is named Microsoft.SCCManaged.CustomRulePack.

2. Use the Get-DlpSensitiveInformationTypeRulePackage cmdlet to store the custom rule package to a variable:

   ```
   $rulepak = Get-DlpSensitiveInformationTypeRulePackage -Identity "RulePackageName"
   ```

   For example, if the name of the rule package is "Employee ID Custom Rule Pack", run the following cmdlet:

   ```
   $rulepak = Get-DlpSensitiveInformationTypeRulePackage -Identity "Employee ID Custom Rule Pack"
   ```

3. Use the Set-Content cmdlet to export the custom rule package to an XML file:

   ```
   Set-Content -Path "XMLFileAndPath" -Encoding Byte -Value
   $rulepak.SerializedClassificationRuleCollection
   ```

   This example export the rule package to the file named ExportedRulePackage.xml in the C:\My Documents folder.

   ```
   Set-Content -Path "C:\My Documents\ExportedRulePackage.xml" -Encoding Byte -Value
   $rulepak.SerializedClassificationRuleCollection
   ```

Step 2: Modify the sensitive information type in the exported XML file

Sensitive information types in the XML file and other elements in the file are described earlier in this topic.

Step 3: Import the updated XML file back into the existing rule package

To import the updated XML back into the existing rule package, use the Set-DlpSensitiveInformationTypeRulePackage cmdlet:

```
Set-DlpSensitiveInformationTypeRulePackage -FileData ([Byte[]]$(Get-Content -Path "C:\My Documents\External
Sensitive Info Type Rule Collection.xml" -Encoding Byte -ReadCount 0))
```

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

Reference: Rule package XML schema definition

You can copy this markup, save it as an XSD file, and use it to validate your rule package XML file.

```
<?xml version="1.0" encoding="utf-8"?>
<xs:schema xmlns:mce="http://schemas.microsoft.com/office/2011/mce"
           targetNamespace="http://schemas.microsoft.com/office/2011/mce"
           xmlns:xs="https://www.w3.org/2001/XMLSchema"
           elementFormDefault="qualified"
           attributeFormDefault="unqualified"
           id="RulePackageSchema">
  <!-- Use include if this schema has the same target namespace as the schema being referenced, otherwise
use import -->
  <xs:element name="RulePackage" type="mce:RulePackageType"/>
  <xs:simpleType name="LangType">
```

```
<xs:simpleType name="LangType">
  <xs:union memberTypes="xs:language">
    <xs:simpleType>
      <xs:restriction base="xs:string">
        <xs:enumeration value=""/>
      </xs:restriction>
    </xs:simpleType>
  </xs:union>
</xs:simpleType>
<xs:simpleType name="GuidType" final="#all">
  <xs:restriction base="xs:token">
    <xs:pattern value="[0-9a-fA-F]{8}\-([0-9a-fA-F]{4}\-){3}[0-9a-fA-F]{12}"/>
  </xs:restriction>
</xs:simpleType>
<xs:complexType name="RulePackageType">
  <xs:sequence>
    <xs:element name="RulePack" type="mce:RulePackType"/>
    <xs:element name="Rules" type="mce:RulesType">
      <xs:key name="UniqueRuleId">
        <xs:selector xpath="mce:Entity|mce:Affinity|mce:Version/mce:Entity|mce:Version/mce:Affinity"/>
        <xs:field xpath="@id"/>
      </xs:key>
      <xs:key name="UniqueProcessorId">
        <xs:selector xpath="mce:Regex|mce:Keyword|mce:Fingerprint"></xs:selector>
        <xs:field xpath="@id"/>
      </xs:key>
      <xs:key name="UniqueResourceIdRef">
        <xs:selector xpath="mce:LocalizedStrings/mce:Resource"/>
        <xs:field xpath="@idRef"/>
      </xs:key>
      <xs:keyref name="ReferencedRuleMustExist" refer="mce:UniqueRuleId">
        <xs:selector xpath="mce:LocalizedStrings/mce:Resource"/>
        <xs:field xpath="@idRef"/>
      </xs:keyref>
      <xs:keyref name="RuleMustHaveResource" refer="mce:UniqueResourceIdRef">
        <xs:selector xpath="mce:Entity|mce:Affinity|mce:Version/mce:Entity|mce:Version/mce:Affinity"/>
        <xs:field xpath="@id"/>
      </xs:keyref>
    </xs:element>
  </xs:sequence>
</xs:complexType>
<xs:complexType name="RulePackType">
  <xs:sequence>
    <xs:element name="Version" type="mce:VersionType"/>
    <xs:element name="Publisher" type="mce:PublisherType"/>
    <xs:element name="Details" type="mce:DetailsType">
      <xs:key name="UniqueLangCodeInLocalizedDetails">
        <xs:selector xpath="mce:LocalizedDetails"/>
        <xs:field xpath="@langcode"/>
      </xs:key>
      <xs:keyref name="DefaultLangCodeMustExist" refer="mce:UniqueLangCodeInLocalizedDetails">
        <xs:selector xpath="."/>
        <xs:field xpath="@defaultLangCode"/>
      </xs:keyref>
    </xs:element>
    <xs:element name="Encryption" type="mce:EncryptionType" minOccurs="0" maxOccurs="1"/>
  </xs:sequence>
  <xs:attribute name="id" type="mce:GuidType" use="required"/>
</xs:complexType>
<xs:complexType name="VersionType">
  <xs:attribute name="major" type="xs:unsignedShort" use="required"/>
  <xs:attribute name="minor" type="xs:unsignedShort" use="required"/>
  <xs:attribute name="build" type="xs:unsignedShort" use="required"/>
  <xs:attribute name="revision" type="xs:unsignedShort" use="required"/>
</xs:complexType>
<xs:complexType name="PublisherType">
  <xs:attribute name="id" type="mce:GuidType" use="required"/>
</xs:complexType>
<xs:complexType name="LocalizedDetailsType">
```

```
    <xs:sequence>
      <xs:element name="PublisherName" type="mce:NameType"/>
      <xs:element name="Name" type="mce:RulePackNameType"/>
      <xs:element name="Description" type="mce:OptionalNameType"/>
    </xs:sequence>
    <xs:attribute name="langcode" type="mce:LangType" use="required"/>
  </xs:complexType>
  <xs:complexType name="DetailsType">
    <xs:sequence>
      <xs:element name="LocalizedDetails" type="mce:LocalizedDetailsType" maxOccurs="unbounded"/>
    </xs:sequence>
    <xs:attribute name="defaultLangCode" type="mce:LangType" use="required"/>
  </xs:complexType>
  <xs:complexType name="EncryptionType">
    <xs:sequence>
      <xs:element name="Key" type="xs:normalizedString"/>
      <xs:element name="IV" type="xs:normalizedString"/>
    </xs:sequence>
  </xs:complexType>
  <xs:simpleType name="RulePackNameType">
    <xs:restriction base="xs:token">
      <xs:minLength value="1"/>
      <xs:maxLength value="64"/>
    </xs:restriction>
  </xs:simpleType>
  <xs:simpleType name="NameType">
    <xs:restriction base="xs:normalizedString">
      <xs:minLength value="1"/>
      <xs:maxLength value="256"/>
    </xs:restriction>
  </xs:simpleType>
  <xs:simpleType name="OptionalNameType">
    <xs:restriction base="xs:normalizedString">
      <xs:minLength value="0"/>
      <xs:maxLength value="256"/>
    </xs:restriction>
  </xs:simpleType>
  <xs:simpleType name="RestrictedTermType">
    <xs:restriction base="xs:string">
      <xs:minLength value="1"/>
      <xs:maxLength value="100"/>
    </xs:restriction>
  </xs:simpleType>
  <xs:complexType name="RulesType">
    <xs:sequence>
      <xs:choice maxOccurs="unbounded">
        <xs:element name="Entity" type="mce:EntityType"/>
        <xs:element name="Affinity" type="mce:AffinityType"/>
        <xs:element name="Version" type="mce:VersionedRuleType"/>
      </xs:choice>
      <xs:choice minOccurs="0" maxOccurs="unbounded">
        <xs:element name="Regex" type="mce:RegexType"/>
        <xs:element name="Keyword" type="mce:KeywordType"/>
        <xs:element name="Fingerprint" type="mce:FingerprintType"/>
        <xs:element name="ExtendedKeyword" type="mce:ExtendedKeywordType"/>
      </xs:choice>
      <xs:element name="LocalizedStrings" type="mce:LocalizedStringsType"/>
    </xs:sequence>
  </xs:complexType>
  <xs:complexType name="EntityType">
    <xs:sequence>
      <xs:element name="Pattern" type="mce:PatternType" maxOccurs="unbounded"/>
      <xs:element name="Version" type="mce:VersionedPatternType" minOccurs="0" maxOccurs="unbounded" />
    </xs:sequence>
    <xs:attribute name="id" type="mce:GuidType" use="required"/>
    <xs:attribute name="patternsProximity" type="mce:ProximityType" use="required"/>
    <xs:attribute name="recommendedConfidence" type="mce:ProbabilityType"/>
    <xs:attribute name="workload" type="mce:WorkloadType"/>
  </xs:complexType>
```

```
  <xs:complexType name="PatternType">
    <xs:sequence>
      <xs:element name="IdMatch" type="mce:IdMatchType"/>
      <xs:choice minOccurs="0" maxOccurs="unbounded">
        <xs:element name="Match" type="mce:MatchType"/>
        <xs:element name="Any" type="mce:AnyType"/>
      </xs:choice>
    </xs:sequence>
    <xs:attribute name="confidenceLevel" type="mce:ProbabilityType" use="required"/>
  </xs:complexType>
  <xs:complexType name="AffinityType">
    <xs:sequence>
      <xs:element name="Evidence" type="mce:EvidenceType" maxOccurs="unbounded"/>
      <xs:element name="Version" type="mce:VersionedEvidenceType" minOccurs="0" maxOccurs="unbounded" />
    </xs:sequence>
    <xs:attribute name="id" type="mce:GuidType" use="required"/>
    <xs:attribute name="evidencesProximity" type="mce:ProximityType" use="required"/>
    <xs:attribute name="thresholdConfidenceLevel" type="mce:ProbabilityType" use="required"/>
    <xs:attribute name="workload" type="mce:WorkloadType"/>
  </xs:complexType>
  <xs:complexType name="EvidenceType">
    <xs:sequence>
      <xs:choice maxOccurs="unbounded">
        <xs:element name="Match" type="mce:MatchType"/>
        <xs:element name="Any" type="mce:AnyType"/>
      </xs:choice>
    </xs:sequence>
    <xs:attribute name="confidenceLevel" type="mce:ProbabilityType" use="required"/>
  </xs:complexType>
  <xs:complexType name="IdMatchType">
    <xs:attribute name="idRef" type="xs:string" use="required"/>
  </xs:complexType>
  <xs:complexType name="MatchType">
    <xs:attribute name="idRef" type="xs:string" use="required"/>
    <xs:attribute name="minCount" type="xs:positiveInteger" use="optional"/>
    <xs:attribute name="uniqueResults" type="xs:boolean" use="optional"/>
  </xs:complexType>
  <xs:complexType name="AnyType">
    <xs:sequence>
      <xs:choice maxOccurs="unbounded">
        <xs:element name="Match" type="mce:MatchType"/>
        <xs:element name="Any" type="mce:AnyType"/>
      </xs:choice>
    </xs:sequence>
    <xs:attribute name="minMatches" type="xs:nonNegativeInteger" default="1"/>
    <xs:attribute name="maxMatches" type="xs:nonNegativeInteger" use="optional"/>
  </xs:complexType>
  <xs:simpleType name="ProximityType">
    <xs:union>
      <xs:simpleType>
        <xs:restriction base='xs:string'>
          <xs:enumeration value="unlimited"/>
        </xs:restriction>
      </xs:simpleType>
      <xs:simpleType>
        <xs:restriction base="xs:positiveInteger">
          <xs:minInclusive value="1"/>
        </xs:restriction>
      </xs:simpleType>
    </xs:union>
  </xs:simpleType>
  <xs:simpleType name="ProbabilityType">
    <xs:restriction base="xs:integer">
      <xs:minInclusive value="1"/>
      <xs:maxInclusive value="100"/>
    </xs:restriction>
  </xs:simpleType>
  <xs:simpleType name="WorkloadType">
    <xs:restriction base="xs:string">
```

```
      <xs:enumeration value="Exchange"/>
      <xs:enumeration value="Outlook"/>
    </xs:restriction>
  </xs:simpleType>
  <xs:simpleType name="EngineVersionType">
    <xs:restriction base="xs:token">
      <xs:pattern value="^\d{2}\.01?\.\d{3,4}\.\d{1,3}$"/>
    </xs:restriction>
  </xs:simpleType>
  <xs:complexType name="VersionedRuleType">
    <xs:choice maxOccurs="unbounded">
      <xs:element name="Entity" type="mce:EntityType"/>
      <xs:element name="Affinity" type="mce:AffinityType"/>
    </xs:choice>
    <xs:attribute name="minEngineVersion" type="mce:EngineVersionType" use="required" />
  </xs:complexType>
  <xs:complexType name="VersionedPatternType">
    <xs:sequence>
      <xs:element name="Pattern" type="mce:PatternType" maxOccurs="unbounded"/>
    </xs:sequence>
    <xs:attribute name="minEngineVersion" type="mce:EngineVersionType" use="required" />
  </xs:complexType>
  <xs:complexType name="VersionedEvidenceType">
    <xs:sequence>
      <xs:element name="Evidence" type="mce:EvidenceType" maxOccurs="unbounded"/>
    </xs:sequence>
    <xs:attribute name="minEngineVersion" type="mce:EngineVersionType" use="required" />
  </xs:complexType>
  <xs:simpleType name="FingerprintValueType">
    <xs:restriction base="xs:string">
      <xs:minLength value="2732"/>
      <xs:maxLength value="2732"/>
    </xs:restriction>
  </xs:simpleType>
  <xs:complexType name="FingerprintType">
    <xs:simpleContent>
      <xs:extension base="mce:FingerprintValueType">
        <xs:attribute name="id" type="xs:token" use="required"/>
        <xs:attribute name="threshold" type="mce:ProbabilityType" use="required"/>
        <xs:attribute name="shingleCount" type="xs:positiveInteger" use="required"/>
        <xs:attribute name="description" type="xs:string" use="optional"/>
      </xs:extension>
    </xs:simpleContent>
  </xs:complexType>
  <xs:complexType name="RegexType">
    <xs:simpleContent>
      <xs:extension base="xs:string">
        <xs:attribute name="id" type="xs:token" use="required"/>
      </xs:extension>
    </xs:simpleContent>
  </xs:complexType>
  <xs:complexType name="KeywordType">
    <xs:sequence>
      <xs:element name="Group" type="mce:GroupType" maxOccurs="unbounded"/>
    </xs:sequence>
    <xs:attribute name="id" type="xs:token" use="required"/>
  </xs:complexType>
  <xs:complexType name="GroupType">
    <xs:sequence>
      <xs:choice>
        <xs:element name="Term" type="mce:TermType" maxOccurs="unbounded"/>
      </xs:choice>
    </xs:sequence>
    <xs:attribute name="matchStyle" default="word">
      <xs:simpleType>
        <xs:restriction base="xs:NMTOKEN">
          <xs:enumeration value="word"/>
          <xs:enumeration value="string"/>
        </xs:restriction>
```

```
      </xs:simpleType>
    </xs:attribute>
  </xs:complexType>
  <xs:complexType name="TermType">
    <xs:simpleContent>
      <xs:extension base="mce:RestrictedTermType">
        <xs:attribute name="caseSensitive" type="xs:boolean" default="false"/>
      </xs:extension>
    </xs:simpleContent>
  </xs:complexType>
  <xs:complexType name="ExtendedKeywordType">
    <xs:simpleContent>
      <xs:extension base="xs:string">
        <xs:attribute name="id" type="xs:token" use="required"/>
      </xs:extension>
    </xs:simpleContent>
  </xs:complexType>
  <xs:complexType name="LocalizedStringsType">
    <xs:sequence>
      <xs:element name="Resource" type="mce:ResourceType" maxOccurs="unbounded">
      <xs:key name="UniqueLangCodeUsedInNamePerResource">
        <xs:selector xpath="mce:Name"/>
        <xs:field xpath="@langcode"/>
      </xs:key>
      <xs:key name="UniqueLangCodeUsedInDescriptionPerResource">
        <xs:selector xpath="mce:Description"/>
        <xs:field xpath="@langcode"/>
      </xs:key>
    </xs:element>
    </xs:sequence>
  </xs:complexType>
  <xs:complexType name="ResourceType">
    <xs:sequence>
      <xs:element name="Name" type="mce:ResourceNameType" maxOccurs="unbounded"/>
      <xs:element name="Description" type="mce:DescriptionType" minOccurs="0" maxOccurs="unbounded"/>
    </xs:sequence>
    <xs:attribute name="idRef" type="mce:GuidType" use="required"/>
  </xs:complexType>
  <xs:complexType name="ResourceNameType">
    <xs:simpleContent>
      <xs:extension base="xs:string">
        <xs:attribute name="default" type="xs:boolean" default="false"/>
        <xs:attribute name="langcode" type="mce:LangType" use="required"/>
      </xs:extension>
    </xs:simpleContent>
  </xs:complexType>
  <xs:complexType name="DescriptionType">
    <xs:simpleContent>
      <xs:extension base="xs:string">
        <xs:attribute name="default" type="xs:boolean" default="false"/>
        <xs:attribute name="langcode" type="mce:LangType" use="required"/>
      </xs:extension>
    </xs:simpleContent>
  </xs:complexType>
</xs:schema>
```

More information

- Overview of data loss prevention policies
- Sensitive information type entity definitions
- What the DLP functions look for

Create a keyword dictionary

2/18/2021 • 6 minutes to read • Edit Online

Data loss prevention (DLP) can identify, monitor, and protect your sensitive items. Identifying sensitive items sometimes requires looking for keywords, particularly when identifying generic content (such as healthcare-related communication), or inappropriate or explicit language. Although you can create keyword lists in sensitive information types, keyword lists are limited in size and require modifying XML to create or edit them. Keyword dictionaries provide simpler management of keywords and at a much larger scale, supporting up to 1MB of terms (post compression) in the dictionary and support any language. The tenant limit is also 1MB after compression. 1MB of post compression limit means that all dictionaries combined across a tenant can have close to 1 million character.

NOTE

Microsoft 365 Information Protection now supports in preview double byte character set languages for:

- Chinese (simplified)
- Chinese (traditional)
- Korean
- Japanese

This support is available for sensitive information types. See, Information protection support for double byte character sets release notes (preview) for more information.

Basic steps to creating a keyword dictionary

The keywords for your dictionary could come from a variety of sources, most commonly from a file (such as a .csv or .txt list) imported in the service or by PowerShell cmdlet, from a list you enter directly in the PowerShell cmdlet, or from an existing dictionary. When you create a keyword dictionary, you follow the same core steps:

1. Use the **Security & Compliance Center** (https://protection.office.com) or connect to **Security & Compliance Center PowerShell**.
2. **Define or load your keywords from your intended source**. The wizard and the cmdlet both accept a comma-separated list of keywords to create a custom keyword dictionary, so this step will vary slightly depending on where your keywords come from. Once loaded, they're encoded and converted to a byte array before they're imported.
3. **Create your dictionary**. Choose a name and description and create your dictionary.

Create a keyword dictionary using the Security & Compliance Center

Use the following steps to create and import keywords for a custom dictionary:

1. Connect to the Security & Compliance Center (https://protection.office.com).
2. Navigate to **Classifications > Sensitive info types**.
3. Select **Create** and enter a **Name** and **Description** for your sensitive info type, then select **Next**
4. Select **Add an element**, then select **Dictionary (Large keywords)** in the **Detect content containing** drop-down list.

5. Select **Add a dictionary**
6. Under the Search control, select **You can create new keyword dictionaries here**.
7. Enter a **Name** for your custom dictionary.
8. Select **Import**, and select either **From text** or **From csv** depending on your keyword file type.
9. In the file dialog, select the keyword file from your local PC or network file share, then select **Open**.
10. Select **Save**, then select your custom dictionary from the **Keyword dictionaries** list.
11. Select **Add**, then select **Next**.
12. Review and finalize your sensitive info type selections, then select **Finish**.

Create a keyword dictionary from a file using PowerShell

Often when you need to create a large dictionary, it's to use keywords from a file or a list exported from some other source. In this case, you'll create a keyword dictionary containing a list of inappropriate language to screen in external email. You must first connect to Security & Compliance Center PowerShell.

1. Copy the keywords into a text file and make sure that each keyword is on a separate line.
2. Save the text file with Unicode encoding. In Notepad > **Save As** > **Encoding** > **Unicode**.
3. Read the file into a variable by running this cmdlet:

   ```
   $fileData = Get-Content <filename> -Encoding Byte -ReadCount 0
   ```

4. Create the dictionary by running this cmdlet:

   ```
   New-DlpKeywordDictionary -Name <name> -Description <description> -FileData $fileData
   ```

Modifying an existing keyword dictionary

You might need to modify keywords in one of your keyword dictionaries, or modify one of the built-in dictionaries. Currently, your can only update a custom keyword dictionary using PowerShell.

For example, we'll modify some terms in PowerShell, save the terms locally where you can modify them in an editor, and then update the previous terms in place.

First, retrieve the dictionary object:

```
$dict = Get-DlpKeywordDictionary -Name "Diseases"
```

Printing `$dict` will show the various variables. The keywords themselves are stored in an object on the backend, but `$dict.KeywordDictionary` contains a string representation of them, which you'll use to modify the dictionary.

Before you modify the dictionary, you need to turn the string of terms back into an array using the `.split(',')` method. Then you'll clean up the unwanted spaces between the keywords with the `.trim()` method, leaving just the keywords to work with.

```
$terms = $dict.KeywordDictionary.split(',').trim()
```

Now you'll remove some terms from the dictionary. Because the example dictionary has only a few keywords, you could just as easily skip to exporting the dictionary and editing it in Notepad, but dictionaries generally contain a large amount of text, so you'll first learn this way to edit them easily in PowerShell.

In the last step, you saved the keywords to an array. There are several ways to remove items from an array, but as a straightforward approach, you'll create an array of the terms you want to remove from the dictionary, and then copy only the dictionary terms to it that aren't in the list of terms to remove.

Run the command `$terms` to show the current list of terms. The output of the command looks like this:

`aarskog's syndrome` `abandonment` `abasia` `abderhalden-kaufmann-lignac` `abdominalgia` `abduction contracture` `abetalipoproteinemia` `abiotrophy` `ablatio` `ablation` `ablepharia` `abocclusion` `abolition` `aborter` `abortion` `abortus` `aboulomania` `abrami's disease`

Run this command to specify the terms that you want to remove:

```
$termsToRemove = @('abandonment', 'ablatio')
```

Run this command to actually remove the terms from the list:

```
$updatedTerms = $terms | Where-Object{ $_ -notin $termsToRemove }
```

Run the command `$updatedTerms` to show the updated list of terms. The output of the command looks like this (the specified terms have been removed):

`aarskog's syndrome` `abasia` `abderhalden-kaufmann-lignac` `abdominalgia` `abduction contracture` `abetalipo proteinemia` `abiotrophy` `ablation` `ablepharia` `abocclusion` `abolition` `aborter` `abortion` `abortus` `aboulomania` `abrami's disease`

```
Now save the dictionary locally and add a few more terms. You could add the terms right here in PowerShell,
but you'll still need to export the file locally to ensure it's saved with Unicode encoding and contains the
BOM.

Save the dictionary locally by running the following:

```powershell
Set-Content $updatedTerms -Path "C:\myPath\terms.txt"
```

Now simply open the file, add your additional terms, and save with Unicode encoding (UTF-16). Now you'll upload the updated terms and update the dictionary in place.

```
PS> Set-DlpKeywordDictionary -Identity "Diseases" -FileData (Get-Content -Path "C:myPath\terms.txt" -
Encoding Byte -ReadCount 0)
```

Now the dictionary has been updated in place. Note that the `Identity` field takes the name of the dictionary. If you wanted to also change the name of your dictionary using the `set-` cmdlet, you would just need to add the `-Name` parameter to what's above with your new dictionary name.

Using keyword dictionaries in custom sensitive information types and DLP policies

Keyword dictionaries can be used as part of the match requirements for a custom sensitive information type, or as a sensitive information type themselves. Both require you to create a custom sensitive information type.
```

Follow the instructions in the linked article to create a sensitive information type. Once you have the XML, you'll need the GUID identifier for the dictionary to use it.

```
<Entity id="9e5382d0-1b6a-42fd-820e-44e0d3b15b6e" patternsProximity="300" recommendedConfidence="75">
    <Pattern confidenceLevel="75">
        <IdMatch idRef=". . ."/>
    </Pattern>
</Entity>
```

To get the identity of your dictionary, run this command and copy the **Identity** property value:

```
Get-DlpKeywordDictionary -Name "Diseases"
```

The output of the command looks like this:

```
RunspaceId : 138e55e7-ea1e-4f7a-b824-79f2c4252255 Identity : 8d2d44b0-91f4-41f2-94e0-21c1c5b5fc9f
Name : Diseases Description : Names of diseases and injuries from ICD-10-CM lexicon
KeywordDictionary : aarskog's syndrome, abandonment, abasia, abderhalden-kaufmann-lignac, abdominalgia,
abduction contracture, abetalipo
proteinemia, abiotrophy, ablatio, ablation, ablepharia, abocclusion, abolition, aborter, abortion, abortus,
aboulomania,
abrami's disease, abramo IsValid : True ObjectState : Unchanged
```

Paste the identity into your custom sensitive information type's XML and upload it. Now your dictionary will appear in your list of sensitive information types and you can use it right in your policy, specifying how many keywords are required to match.

```
<Entity id="d333c6c2-5f4c-4131-9433-db3ef72a89e8" patternsProximity="300" recommendedConfidence="85">
      <Pattern confidenceLevel="85">
        <IdMatch idRef="8d2d44b0-91f4-41f2-94e0-21c1c5b5fc9f" />
      </Pattern>
    </Entity>
    <LocalizedStrings>
      <Resource idRef="d333c6c2-5f4c-4131-9433-db3ef72a89e8">
        <Name default="true" langcode="en-us">Diseases</Name>
        <Description default="true" langcode="en-us">Detects various diseases</Description>
      </Resource>
    </LocalizedStrings>
```

Document Fingerprinting

11/2/2020 • 6 minutes to read • Edit Online

Information workers in your organization handle many kinds of sensitive information during a typical day. In the Security & Compliance Center, Document Fingerprinting makes it easier for you to protect this information by identifying standard forms that are used throughout your organization. This topic describes the concepts behind Document Fingerprinting and how to create one by using PowerShell.

Basic scenario for Document Fingerprinting

Document Fingerprinting is a Data Loss Prevention (DLP) feature that converts a standard form into a sensitive information type, which you can use in the rules of your DLP policies. For example, you can create a document fingerprint based on a blank patent template and then create a DLP policy that detects and blocks all outgoing patent templates with sensitive content filled in. Optionally, you can set up policy tips to notify senders that they might be sending sensitive information, and the sender should verify that the recipients are qualified to receive the patents. This process works with any text-based forms used in your organization. Additional examples of forms that you can upload include:

- Government forms
- Health Insurance Portability and Accountability Act (HIPAA) compliance forms
- Employee information forms for Human Resources departments
- Custom forms created specifically for your organization

Ideally, your organization already has an established business practice of using certain forms to transmit sensitive information. After you upload an empty form to be converted to a document fingerprint and set up a corresponding policy, the DLP detects any documents in outbound mail that match that fingerprint.

How Document Fingerprinting works

You've probably already guessed that documents don't have actual fingerprints, but the name helps explain the feature. In the same way that a person's fingerprints have unique patterns, documents have unique word patterns. When you upload a file, DLP identifies the unique word pattern in the document, creates a document fingerprint based on that pattern, and uses that document fingerprint to detect outbound documents containing the same pattern. That's why uploading a form or template creates the most effective type of document fingerprint. Everyone who fills out a form uses the same original set of words and then adds his or her own words to the document. As long as the outbound document isn't password protected and contains all the text from the original form, DLP can determine if the document matches the document fingerprint.

> **IMPORTANT**
> For now, DLP can use document fingerprinting as a detection method in Exchange online only.

The following example shows what happens if you create a document fingerprint based on a patent template, but you can use any form as a basis for creating a document fingerprint.

Example of a patent document matching a document fingerprint of a patent template

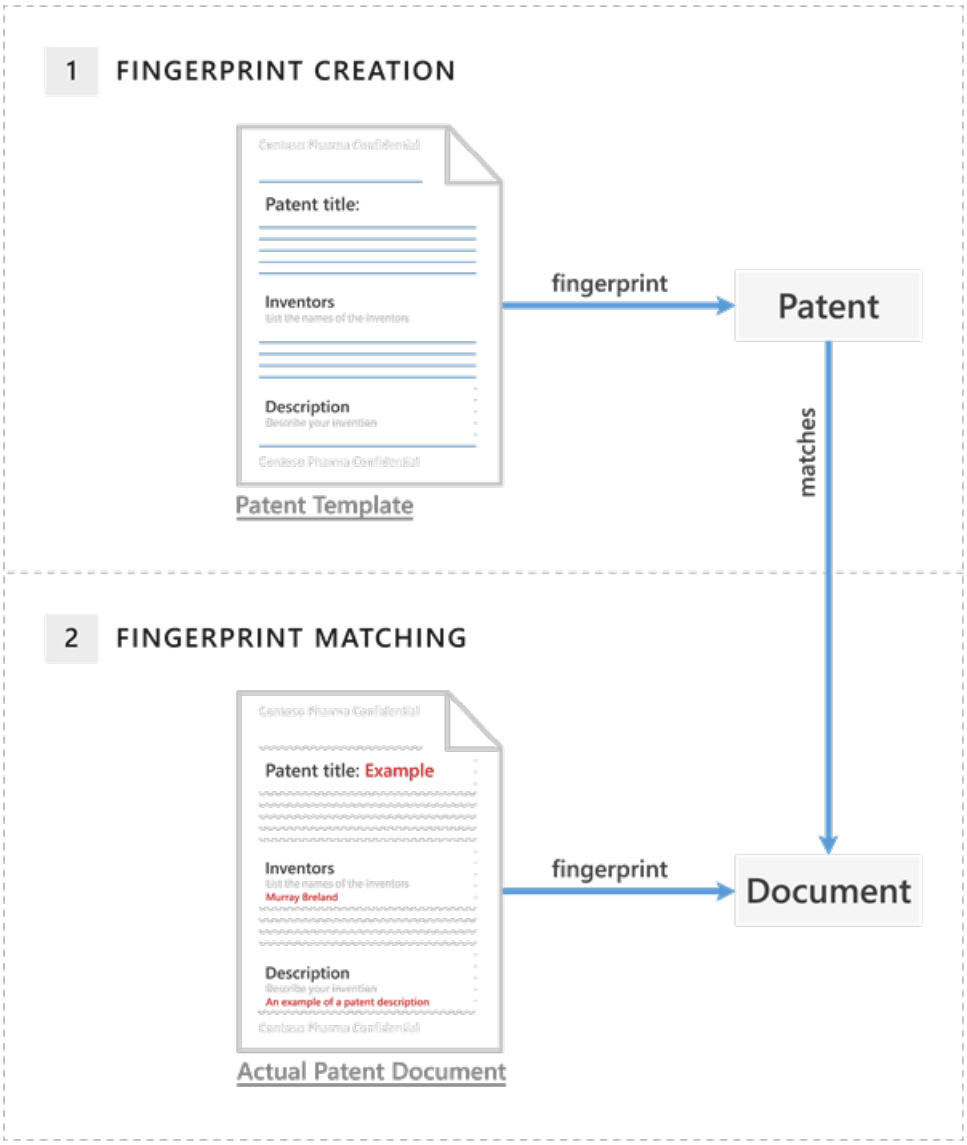

The patent template contains the blank fields "Patent title," "Inventors," and "Description" and descriptions for each of those fields—that's the word pattern. When you upload the original patent template, it's in one of the supported file types and in plain text. DLP converts this word pattern into a document fingerprint, which is a small Unicode XML file containing a unique hash value representing the original text, and the fingerprint is saved as a data classification in Active Directory. (As a security measure, the original document itself isn't stored on the service; only the hash value is stored, and the original document can't be reconstructed from the hash value.) The patent fingerprint then becomes a sensitive information type that you can associate with a DLP policy. After you associate the fingerprint with a DLP policy, DLP detects any outbound emails containing documents that match the patent fingerprint and deals with them according to your organization's policy.

For example, you might want to set up a DLP policy that prevents regular employees from sending outgoing messages containing patents. DLP will use the patent fingerprint to detect patents and block those emails. Alternatively, you might want to let your legal department to be able to send patents to other organizations because it has a business need for doing so. You can allow specific departments to send sensitive information by creating exceptions for those departments in your DLP policy, or you can allow them to override a policy tip with a business justification.

Supported file types

Document Fingerprinting supports the same file types that are supported in mail flow rules (also known as transport rules). For a list of supported file types, see Supported file types for mail flow rule content inspection. One quick note about file types: neither mail flow rules nor Document Fingerprinting supports the .dotx file type, which can be confusing because that's a template file in Word. When you see the word "template" in this and other Document Fingerprinting topics, it refers to a document that you have established as a standard form, not

the template file type.

Limitations of document fingerprinting

Document Fingerprinting won't detect sensitive information in the following cases:

- Password protected files
- Files that contain only images
- Documents that don't contain all the text from the original form used to create the document fingerprint

Use PowerShell to create a classification rule package based on document fingerprinting

Note that you can currently create a document fingerprint only by using PowerShell in the Security & Compliance Center. To connect, see Connect to Security & Compliance Center PowerShell.

DLP uses classification rule packages to detect sensitive content. To create a classification rule package based on a document fingerprint, use the **New-DlpFingerprint** and **New-DlpSensitiveInformationType** cmdlets. Because the results of **New-DlpFingerprint** aren't stored outside the data classification rule, you always run **New-DlpFingerprint** and **New-DlpSensitiveInformationType** or **Set-DlpSensitiveInformationType** in the same PowerShell session. The following example creates a new document fingerprint based on the file C:\My Documents\Contoso Employee Template.docx. You store the new fingerprint as a variable so you can use it with the **New-DlpSensitiveInformationType** cmdlet in the same PowerShell session.

```
$Employee_Template = Get-Content "C:\My Documents\Contoso Employee Template.docx" -Encoding byte -ReadCount
0
$Employee_Fingerprint = New-DlpFingerprint -FileData $Employee_Template -Description "Contoso Employee
Template"
```

Now, let's create a new data classification rule named "Contoso Employee Confidential" that uses the document fingerprint of the file C:\My Documents\Contoso Customer Information Form.docx.

```
$Customer_Form = Get-Content "C:\My Documents\Contoso Customer Information Form.docx" -Encoding byte -
ReadCount 0
$Customer_Fingerprint = New-DlpFingerprint -FileData $Customer_Form -Description "Contoso Customer
Information Form"
New-DlpSensitiveInformationType -Name "Contoso Customer Confidential" -Fingerprints $Customer_Fingerprint -
Description "Message contains Contoso customer information."
```

You can now use the **Get-DlpSensitiveInformationType** cmdlet to find all DLP data classification rule packages, and in this example, "Contoso Customer Confidential" is part of the data classification rule packages list.

Finally, add the "Contoso Customer Confidential" data classification rule package to a DLP policy in the Security & Compliance Center. This example adds a rule to an existing DLP policy named "ConfidentialPolicy".

```
New-DlpComplianceRule -Name "ContosoConfidentialRule" -Policy "ConfidentialPolicy" -
ContentContainsSensitiveInformation @{Name="Contoso Customer Confidential"} -BlockAccess $True
```

You can also use the data classification rule package in mail flow rules in Exchange Online, as shown in the following example. To run this command, you first need to Connect to Exchange Online PowerShell. Also note that it takes time for the rule package to sync from the Security & Compliance Center to the Exchange admin center.

```
New-TransportRule -Name "Notify :External Recipient Contoso confidential" -NotifySender NotifyOnly -Mode
Enforce -SentToScope NotInOrganization -MessageContainsDataClassification @{Name=" Contoso Customer
Confidential"}
```

DLP now detects documents that match the Contoso Customer Form.docx document fingerprint.

For syntax and parameter information, see:

- New-DlpFingerprint
- New-DlpSensitiveInformationType
- Remove-DlpSensitiveInformationType
- Set-DlpSensitiveInformationType
- Get-DlpSensitiveInformationType

What the DLP functions look for

11/2/2020 • 10 minutes to read • Edit Online

Data loss prevention (DLP) policies can use sensitive information types to identify sensitive items. Credit card number and EU debit card number are examples of sensitive information types. Sensitive information types look for specific patterns. Sensitive information types validate the data by looking at it's format, it's checksums, and looks for relevant keywords or other information. Some of this functionality is performed by internal functions. For example, the Credit Card Number sensitive information type uses a function to look for dates that are formatted like an expiration date. This helps to corroborate that a number is a credit card number.

This article explains what these functions look for, to help you understand how the predefined sensitive information types work. For more information, see Sensitive information type entity definitions

Table of functions

FUNCTION NAME	FUNCTION ACTION	IS A VALIDATOR	
Func_Argentina_Unique_Tax_Key	detects and validates Argentina Unique tax key	no	
Func_aba_routing	detects ABA routing number	yes	
Func_alabama_drivers_license_number	detects Alabama driver's license number	no	
Func_alaska_delaware_oregon_drivers_license_number	detects Alaska, Delaware, Oregon driver's license number	no	
Func_alaska_drivers_license_number	detects Alaska driver's license number	no	
Func_alberta_drivers_license_number	detects Alberta driver's license number	no	
Func_Argentina_Unique_Tax_Key	detects Argentina Unique tax key	no	
Func_arizona_drivers_license_number	detects Arizona driver's license number	no	
Func_arkansas_drivers_license_number	detects Arkansas driver's license number	no	
Func_australian_business_number	detects Australia business number	no	
Func_Australian_Company_Number	detects Australia company number	no	

FUNCTION NAME	FUNCTION ACTION	IS A VALIDATOR	
Func_australian_medical_account_number	detects Australia medical account number	no	
Func_australian_tax_file_number	detects Australia tax file number	yes	
Func_austria_eu_ssn_or_equivalent	detects Austria social security number	no	
Func_austria_eu_tax_file_number	detects Austria tax file number	no	
Func_Austria_Value_Added_Tax	detects Austria Value Added Tax	no	
Func_belgium_national_number	detects Belgium national number	no	
Func_belgium_value_added_tax_number	detects Belgium value added tax number	no	
Func_brazil_cnpj	detects Brazil legal entity number (CNPJ)	yes	
Func_brazil_cpf	detects Brazil CPF	yes	
Func_brazil_rg	detects Brazil RG	no	
Func_british_columbia_drivers_license_number	detects British Columbia driver's license number	no	
Func_bulgaria_eu_national_id_card	detects Bulgaria uniform civil number	no	
Func_california_drivers_license_number	detects California driver's license number	no	
Func_canadian_sin	detects Canada sin	yes	
Func_chile_id_card	detects Chile ID card	no	
Func_china_resident_id	detects China-resident ID	no	
Func_colorado_drivers_license_number	detects Colorado driver's license number	no	
Func_connecticut_drivers_license_number	detects Connecticut driver's license number	no	
Func_credit_card	detects credit card	yes	no
Func_croatia_id_card	detects Croatia ID card	no	

FUNCTION NAME	FUNCTION ACTION	IS A VALIDATOR	
Func_croatia_oib_number	detects Croatia OIB number	no	
Func_cyprus_eu_tax_file_number	detects Cyprus tax file number	no	
Func_czech_id_card	detects Czech ID card	no	
Func_czech_id_card_new_format	detects Czech ID card in new format	no	
Func_dea_number	detects DEA number	yes	
Func_denmark_eu_tax_file_number	detects Denmark personal identification number	no	
Func_district_of_columbia_drivers_license_number	detects District of Columbia driver's license number	no	
Func_estonia_eu_national_id_card	detects Estonia Personal Identification Code	no	
Func_eu_debit_card	detects EU debit card	no	
Func_finnish_national_id	detects Finnish national ID	no	
Func_florida_drivers_license_number	detects Florida driver's license number	no	
Func_florida_maryland_michigan_minnesota_drivers_license_number	detects Florida, Maryland, Michigan, Minnesota driver's license number	no	
Func_formatted_itin	detects formatted US ITIN	yes	
Func_fr_insee	detects France INSEE	no	
Func_fr_passport	detects France passport	no	
Func_france_eu_tax_file_number	detects France tax file number	no	
Func_france_value_added_tax_number	detects France value added tax number	no	
Func_french_drivers_license	detects French driver's license	no	
Func_french_insee	detects French INSEE	no	
Func_georgia_drivers_license_number	detects Georgia driver's license number	no	

FUNCTION NAME	FUNCTION ACTION	IS A VALIDATOR	
Func_german_drivers_license	detects Germany driver's license	no	
Func_german_passport	detects Germany passport	no	
Func_german_passport_data	detects Germany passport	no	
Func_germany_eu_tax_file_number	detects Germany tax file number	no	
Func_germany_value_added_tax_number	detects Germany value added tax number	no	
Func_greece_eu_ssn	detects Greece sin (AMKA)	no	
Func_hawaii_drivers_license_number	detects Hawaii driver's license number	no	
Func_hong_kong_id_card	detects Hong Kong ID card	no	
Func_hungarian_value_added_tax_number	detects Hungary value added tax number	no	
Func_hungary_eu_national_id_card	detects Hungary personal identification number	no	
Func_hungary_eu_ssn_or_equivalent	detects Hungary social security number	no	
Func_hungary_eu_tax_file_number	detects Hungary tax file number	no	
Func_iban	detects IBAN	yes	
Func_idaho_drivers_license_number	detects Idaho driver's license number	no	
Func_illinois_drivers_license_number	detects Illinois driver's license number	no	
Func_india_aadhaar	detects India aadhaar	yes	
Func_indiana_drivers_license_number	detects Indiana driver's license number	no	
Func_iowa_drivers_license_number	detects Iowa driver's license number	no	
Func_ireland_pps	detects Ireland PPS	no	
Func_israeli_national_id_number	detects Israel national ID number	no	

FUNCTION NAME	FUNCTION ACTION	IS A VALIDATOR	
Func_italy_eu_national_id_card	detects Italy fiscal code	no	
Func_italy_value_added_tax_number	detects Italy value added tax number	no	
Func_japanese_my_number_corporate	detects Japan my number corporate	yes	
Func_japanese_my_number_personal	detects Japan my number personal	yes	
Func_jp_bank_account	detects Japan bank account	no	
Func_jp_bank_account_branch_code	detects Japan bank account branch code	no	
Func_jp_drivers_license_number	detects Japan driver's license number	no	
Func_jp_passport	detects Japan passport	no	
Func_jp_resident_registration_number	detects Japan-resident registration number	no	
Func_jp_sin	detects Japan SIN	no	
Func_jp_sin_pre_1997	detects Japan sin pre 1997	no	
Func_kansas_drivers_license_number	detects Kansas driver's license number	no	
Func_kentucky_drivers_license_number	detects Kentucky driver's license number	no	
Func_kentucky_massachusetts_virginia_drivers_license_number	detects Kentucky, Massachusetts, Virginia driver's license number	no	
Func_latvia_eu_national_id_card	detects Latvia personal code	no	
Func_lithuania_eu_tax_file_number	detects Lithuania personal code	no	
Func_louisiana_drivers_license_number	detects Louisiana driver's license number	no	
Func_luxemburg_eu_tax_file_number	detects Luxemburg national identification number (natural persons)	no	

FUNCTION NAME	FUNCTION ACTION	IS A VALIDATOR	
Func_luxemburg_eu_tax_file_number_non_natural	detects Luxemburg national identification number (non-natural persons)	no	
Func_maine_drivers_license_number	detects Maine driver's license number	no	
Func_manitoba_drivers_license_number	detects Manitoba driver's license number	no	
Func_maryland_drivers_license_number	detects Maryland driver's license number	no	
Func_massachusetts_drivers_license_number	detects Massachusetts driver's license number	no	
Func_mexico_population_registry_code	detects Mexico population registry code	no	
Func_michigan_minnesota_drivers_license_number	detects Michigan, Minnesota driver's license number	no	
Func_minnesota_drivers_license_number	detects Minnesota driver's license number	no	
Func_mississippi_oklahoma_drivers_license_number	detects Mississippi, Oklahoma driver's license number	no	
Func_missouri_drivers_license_number	detects Missouri driver's license number	no	
Func_montana_drivers_license_number	detects Montana driver's license number	no	
Func_nebraska_drivers_license_number	detects Nebraska driver's license number	no	
Func_netherlands_bsn	detects Netherlands BSN	no	
Func_netherlands_eu_tax_file_number	detects Netherlands tax file number	no	
Func_netherlands_value_added_tax_number	detects Netherlands value added tax number	no	
Func_nevada_drivers_license_number	detects Nevada driver's license number	no	
Func_new_brunswick_drivers_license_number	detects New Brunswick driver's license number	no	

FUNCTION NAME	FUNCTION ACTION	IS A VALIDATOR	
Func_new_hampshire_drivers_license_number	detects New Hampshire driver's license number	no	
Func_new_jersey_drivers_license_number	detects New Jersey driver's license number	no	
Func_new_mexico_drivers_license_number	detects New Mexico driver's license number	no	
Func_new_york_drivers_license_number	detects New York driver's license number	no	
Func_new_zealand_bank_account_number	detects New Zealand bank account number	no	
Func_new_zealand_inland_revenue_number	detects New Zealand inland revenue number	no	
Func_new_zealand_ministry_of_health_number	detects New Zealand ministry of health number	no	
Func_newfoundland_labrador_drivers_license_number	detects Newfoundland Labrador driver's license number	no	
Func_newzealand_driver_license_number	detects New Zealand driver license number	no	
Func_newzealand_social_welfare_number	detects New Zealand social welfare number	no	
Func_north_carolina_drivers_license_number	detects North Carolina driver's license number	no	
Func_north_dakota_drivers_license_number	detects North Dakota driver's license number	no	
Func_norway_id_number	detects Norway ID number	no	
Func_nova_scotia_drivers_license_number	detects Nova Scotia driver's license number	no	
Func_ohio_drivers_license_number	detects Ohio driver's license number	no	
Func_ontario_drivers_license_number	detects Ontario driver's license number	no	
Func_pennsylvania_drivers_license_number	detects Pennsylvania driver's license number	no	
Func_pesel_identification_number	detects Poland National ID (PESEL)	no	

FUNCTION NAME	FUNCTION ACTION	IS A VALIDATOR	
Func_poland_eu_tax_file_number	detects Poland tax file number	no	
Func_polish_national_id	detects Poland identity card	no	
Func_polish_passport_number	detects Polish passport number	no	
Func_polish_regon_number	detects Polish REGON number	no	
Func_portugal_eu_tax_file_number	detects Portugal Tax Identification Number	no	
Func_prince_edward_island_drivers_license_number	detects Prince Edward Island driver's license number	no	
Func_quebec_drivers_license_number	detects Quebec driver's license number	no	
Func_randomized_formatted_ssn	detects randomized formatted US SSN	yes	
Func_randomized_unformatted_ssn	detects randomized unformatted US SSN	yes	
Func_rhode_island_drivers_license_number	detects Rhode Island driver's license number	no	
Func_romania_eu_national_id_card	detects Romania personal numeric code (CNP)	no	
Func_saskatchewan_drivers_license_number	detects Saskatchewan driver's license number	no	
Func_slovakia_eu_national_id_card	detects Slovakia personal number	no	
Func_slovenia_eu_national_id_card	detects Slovenia Unique Master Citizen Number	no	
Func_slovenia_eu_tax_file_number	detects Slovenia tax file number	no	
Func_south_africa_identification_number	detects South Africa identification number	yes	
Func_south_carolina_drivers_license_number	detects South Carolina driver's license number	no	
Func_south_dakota_drivers_license_number	detects South Dakota driver's license number	no	

FUNCTION NAME	FUNCTION ACTION	IS A VALIDATOR	
Func_south_korea_resident_number	detects South Korea resident number	no	
Func_spain_eu_DL_and_NI_number_citizen	detects Spain DL and NI number citizen	no	
Func_spain_eu_DL_and_NI_number_foreigner	detects Spain DL and NI number foreigner	no	
Func_spain_eu_driver's_license_number	detects Spain driver's license number	no	
Func_spain_eu_tax_file_number	detects Spain tax file number	no	
Func_spanish_social_security_number	detects Spanish social security number	no	
Func_ssn	Function to detect non-randomized formatted US SSN	yes	
Func_sweden_eu_tax_file_number	detects Sweden tax file number	no	
Func_swedish_national_identifier	detects Swedish national identifier	yes	
Func_swiss_social_security_number_ahv	detects Swiss social security number AHV	no	
Func_taiwanese_national_id	detects Taiwanese national ID	no	
Func_tennessee_drivers_license_number	detects Tennessee driver's license number	no	
Func_texas_drivers_license_number	detects Texas driver's license number	no	
Func_Thai_Citizen_Id	detects Thai Citizen ID	no	
Func_Turkish_National_Id	detects Turkish National ID	yes	
Func_uk_drivers_license	detects UK driver's license	no	
Func_uk_eu_tax_file_number	detects UK unique taxpayer number	no	
Func_uk_nhs_number	detects UK NHS number	yes	
Func_uk_nino	detects UK NINO	no	

FUNCTION NAME	FUNCTION ACTION	IS A VALIDATOR	
Func_unformatted_canadian_sin	detects unformatted Canadian SIN	no	
Func_unformatted_itin	detects unformatted US ITIN	yes	
Func_unformatted_ssn	detects non-randomized unformatted US SSN	yes	
Func_usa_uk_passport	detects USA and UK passport	yes	
Func_utah_drivers_license_number	detects Utah driver's license number	no	
Func_vermont_drivers_license_number	detects Vermont driver's license number	no	
Func_virginia_drivers_license_number	detects Virginia driver's license number	no	
Func_washington_drivers_license_number	detects Washington driver's license number	no	
Func_west_virginia_drivers_license_number	detects West Virginia driver's license number	no	
Func_wisconsin_drivers_license_number	detects Wisconsin driver's license number	no	
Func_wyoming_drivers_license_number	detects Wyoming driver's license number	no	

Func_us_date

Func_us_date looks for dates in common U.S. formats. The common formats are "month/day/year", "month-day-year", and "month day year ". The names or abbreviations of months aren't case-sensitive.

Examples:

- December 2, 2016
- Dec 2, 2016
- dec 02 2016
- 12/2/2016
- 12/02/16
- Dec-2-2016
- 12-2-16

Accepted month names:

- English
  - January, February, march, April, may, June, July, August, September, October, November, December
  - Jan. Feb. Mar. Apr. May June July Aug. Sept. Oct. Nov. Dec.

Func_eu_date

Fund_eu_dates looks for dates in common E.U. formats (and most places outside the U.S.), such as "day/month/year", "day-month-year", and "day month year". The names or abbreviations of months aren't case-sensitive.

Examples:

- 2 Dec 2016
- 02 dec 2016
- 2 Dec 16
- 2/12/2016
- 02/12/16
- 2-Dec-2016
- 2-12-16

Accepted month names:

- English
  - January, February, march, April, may, June, July, August, September, October, November, December
  - Jan. Feb. Mar. Apr. May June July Aug. Sept. Oct. Nov. Dec.
- Dutch
  - januari, februari, maart, April, mei, juni, juli, augustus, September, ocktober, October, November, December
  - jan feb maart apr mei jun jul aug sep sept oct okt nov dec
- French
  - janvier, février, mars, avril, mai, juin juillet, août, septembre, octobre, novembre, décembre
  - janv. févr. mars avril mai juin juil. août sept. oct. nov. déc.
- German
  - jänuar, februar, märz, April, mai, juni juli, August, September, oktober, November, dezember
  - Jan./Jän. Feb. März Apr. Mai Juni Juli Aug. Sept. Okt. Nov. Dez.
- Italian
  - gennaio, febbraio, marzo, aprile, maggio, giugno, luglio, agosto, settembre, ottobre, novembre, dicembre
  - genn. febbr. mar. apr. magg. giugno luglio ag. sett. ott. nov. dic.
- Portuguese

    - janeiro, fevereiro, março, marco, abril, maio, junho, julho, agosto, setembro, outubro, novembro, dezembro
    - jan fev mar abr mai jun jul ago set out nov dez
- Spanish
    - enero, febrero, marzo, abril, mayo, junio, julio, agosto, septiembre, octubre, noviembre, diciembre
    - enero feb. marzo abr. mayo jun. jul. agosto sept./set. oct. nov. dic.

Func_eu_date1 (deprecated)

> **NOTE**
> This function is deprecated because it supports only Portuguese month names, which are now included in the `Func_eu_date` function above.

This function looks for a date in the format commonly used in Portuguese. The format for this function is the same as `Func_eu_date`, differing only in the language used.

Examples:

- 2 Dez 2016
- 02 dez 2016
- 2 Dez 16
- 2/12/2016
- 02/12/16
- 2-Dez-2016
- 2-12-16

Accepted month names:

- Portuguese
    - janeiro, fevereiro, março, marco, abril, maio, junho, julho, agosto, setembro, outubro, novembro, dezembro
    - jan fev mar abr mai jun jul ago set out nov dez

Func_eu_date2 (deprecated)

> **NOTE**
> This function is deprecated because it supports only Dutch month names, which are now included in the `Func_eu_date` function above.

This function looks for a date in the format commonly used in Dutch. The format for this function is the same as `Func_eu_date`, differing only in the language used.

Examples:

- 2 Mei 2016

- 02 mei 2016
- 2 Mei 16
- 2/12/2016
- 02/12/16
- 2-Mei-2016
- 2-12-16

Accepted month names:

- Dutch
  - januari, februari, maart, April, mei, juni, juli, augustus, September, ocktober, October, November, December
  - jan feb maart apr mei jun jul aug sep sept out okt nov dec

Func_expiration_date

Func_expiration_date looks for dates that are in formats commonly used by credit and debit cards. This function will match dates in format of "month/year", "month-year", "[month name] year", and "[month abbreviation] year". The names or abbreviations of months aren't case-sensitive.

Examples:

- MM/YY -- for example, 01/11 or 1/11
- MM/YYYY -- for example, 01/2011 or 1/2011
- MM-YY -- for example, 01-22 or 1-11
- MM-YYYY -- for example, 01-2000 or 1-2000

The following formats support YY or YYYY:

- Month-YYYY -- for example Jan-2010 or january-2010 or Jan-10 or january-10
- Month YYYY -- for example, 'january 2010' or 'Jan 2010' or 'january 10' or 'Jan 10'
- MonthYYYY -- for example, 'january2010' or 'Jan2010' or 'january10' or 'Jan10'
- Month/YYYY -- for example, 'january/2010' or 'Jan/2010' or 'january/10' or 'Jan/10'

Accepted month names:

- English
  - January, February, march, April, may, June, July, August, September, October, November, December
  - Jan Feb Mar Apr May June July Aug Sept Oct Nov Dec

Func_us_address

Func_us_address looks for a U.S. state name or postal abbreviation followed by a valid zip code. The zip code must be one of the correct zip codes associated with the U.S. state name or abbreviation. The U.S. state name and zip code cannot be separated by punctuation or letters.

Examples:

- Washington 98052
- Washington 98052-9998
- WA 98052
- WA 98052-9998

Learn about trainable classifiers

2/18/2021 • 5 minutes to read • Edit Online

Classifying and labeling content so it can be protected and handled properly is the starting place for the information protection discipline. Microsoft 365 has three ways to classify content.

Manually

This method requires human judgment and action. An admin may either use the pre-existing labels and sensitive information types or create their own and then publish them. Users and admins apply them to content as they encounter it. You can then protect the content and manage its disposition.

Automated pattern matching

This category of classification mechanisms include finding content by:

- Keywords or metadata values (keyword query language).
- Using previously identified patterns of sensitive information like social security, credit card or bank account numbers (Sensitive information type entity definitions).
- Recognizing an item because it's a variation on a template (document finger printing).
- Using the presence of exact strings (exact data match).

Sensitivity and retention labels can then be automatically applied to make the content available for use in data loss prevention (DLP) and auto-apply polices for retention labels.

Classifiers

This classification method is particularly well suited to content that isn't easily identified by either the manual or automated pattern matching methods. This method of classification is more about training a classifier to identify an item based on what the item is, not by elements that are in the item (pattern matching). A classifier learns how to identify a type of content by looking at hundreds of examples of the content you're interested in classifying. You start by feeding it examples that are definitely in the category. Once it processes those, you test it by giving it a mix of both matching and non-matching examples. The classifier then makes predictions as to whether any given item falls into the category you're building. You then confirm its results, sorting out the true positives, true negatives, false positives, and false negatives to help increase the accuracy of its predictions.

When you publish the classifier, it sorts through items in locations like SharePoint Online, Exchange, and OneDrive, and classifies the content. After you publish the classifier, you can continue to train it using a feedback process that is similar to the initial training process.

Where you can use trainable classifiers

Both built-in classifiers and trainable classifiers are available as a condition for Office autolabeling with sensitivity labels, auto-apply retention label policy based on a condition and in communication compliance.

Sensitivity labels can use classifiers as conditions, see Apply a sensitivity label to content automatically.

> **IMPORTANT**
> Classifiers only work with items that are not encrypted and are in English.

Types of classifiers

- **pre-trained classifiers** - Microsoft has created and pre-trained a number of classifiers that you can start using without training them. These classifiers will appear with the status of `Ready to use`.
- **custom classifiers** - If you have classification needs that extend beyond what the pre-trained classifiers cover, you can create and train your own classifiers.

Pre-trained classifiers

Microsoft 365 comes with five pre-trained classifiers:

Caution

We are deprecating the **Offensive Language** pre-trained classifier because it has been producing a high number of false positives. Don't use it and if you are currently using it, you should move your business processes off of it. We recommend using the **Threat**, **Profanity**, and **Harassment** pre-trained classifiers instead.

- **Resumes**: detects items that are textual accounts of an applicant's personal, educational, professional qualifications, work experience, and other personally identifying information
- **Source Code**: detects items that contain a set of instructions and statements written in the top 25 used computer programming languages on GitHub
  - ActionScript
  - C
  - C#
  - C++
  - Clojure
  - CoffeeScript
  - Go
  - Haskell
  - Java
  - JavaScript
  - Lua
  - MATLAB
  - Objective-C
  - Perl
  - PHP
  - Python
  - R
  - Ruby
  - Scala
  - Shell
  - Swift
  - Tex
  - Vim Script

> **NOTE**
> Source Code is trained to detect when the bulk of the text is source code. It does not detect source code text that is interspersed with plain text.

- **Harassment**: detects a specific category of offensive language text items related to offensive conduct

targeting one or multiple individuals based on the following traits: race, ethnicity, religion, national origin, gender, sexual orientation, age, disability

- **Profanity**: detects a specific category of offensive language text items that contain expressions that embarrass most people
- **Threat**: detects a specific category of offensive language text items related to threats to commit violence or do physical harm or damage to a person or property

These appear in the **Microsoft 365 compliance center** > **Data classification** > **Trainable classifiers** view with the status of `Ready to use`.

> **IMPORTANT**
> Please note that the offensive language, harassment, profanity, and threat classifiers only work with searchable text are not exhaustive or complete. Further, language and cultural standards continually change, and in light of these realities, Microsoft reserves the right to update these classifiers in its discretion. While the classifiers may assist your organization in monitoring offensive and other language used, the classifiers do not address consequences of such language and are not intended to provide your organization's sole means of monitoring or responding to the use of such language. Your organization, and not Microsoft or its subsidiaries, remains responsible for all decisions related to monitoring, enforcement, blocking, removal and retention of any content identified by a pre-trained classifier.

Custom classifiers

When the pre-trained classifiers don't meet your needs, you can create and train your own classifiers. There's significantly more work involved with creating your own, but they'll be much better tailored to your organizations needs.

For example you could create trainable classifiers for:

- Legal documents - such as attorney client privilege, closing sets, statement of work
- Strategic business documents - like press releases, merger and acquisition, deals, business or marketing plans, intellectual property, patents, design docs
- Pricing information - like invoices, price quotes, work orders, bidding documents
- Financial information - such as organizational investments, quarterly or annual results

Process flow for creating custom classifiers

Creating and publishing a classifier for use in compliance solutions, such as retention policies and communication supervision, follows this flow. For more detail on creating a custom trainable classifier see, Creating a custom classifier.

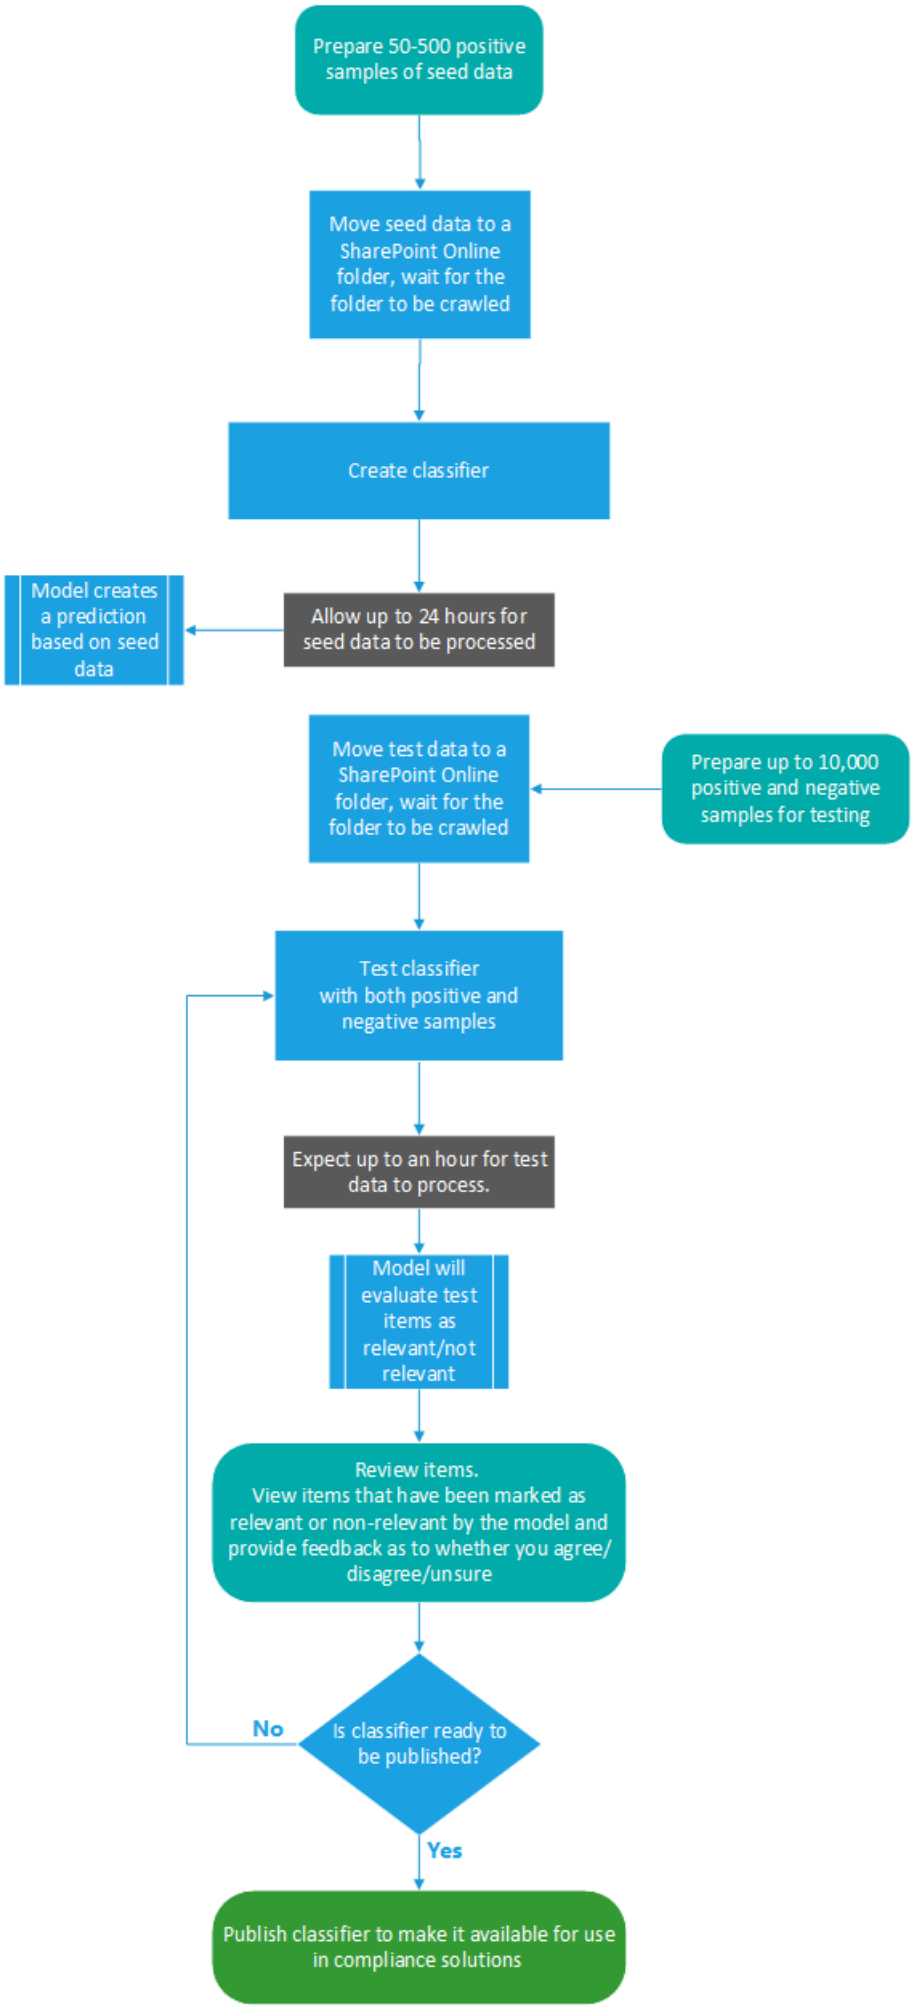

Retraining classifiers

You can help improve the accuracy of all custom classifiers and some pre-trained classifiers by providing them with feedback on the accuracy of the classification that they perform. This is called retraining and follow this workflow.

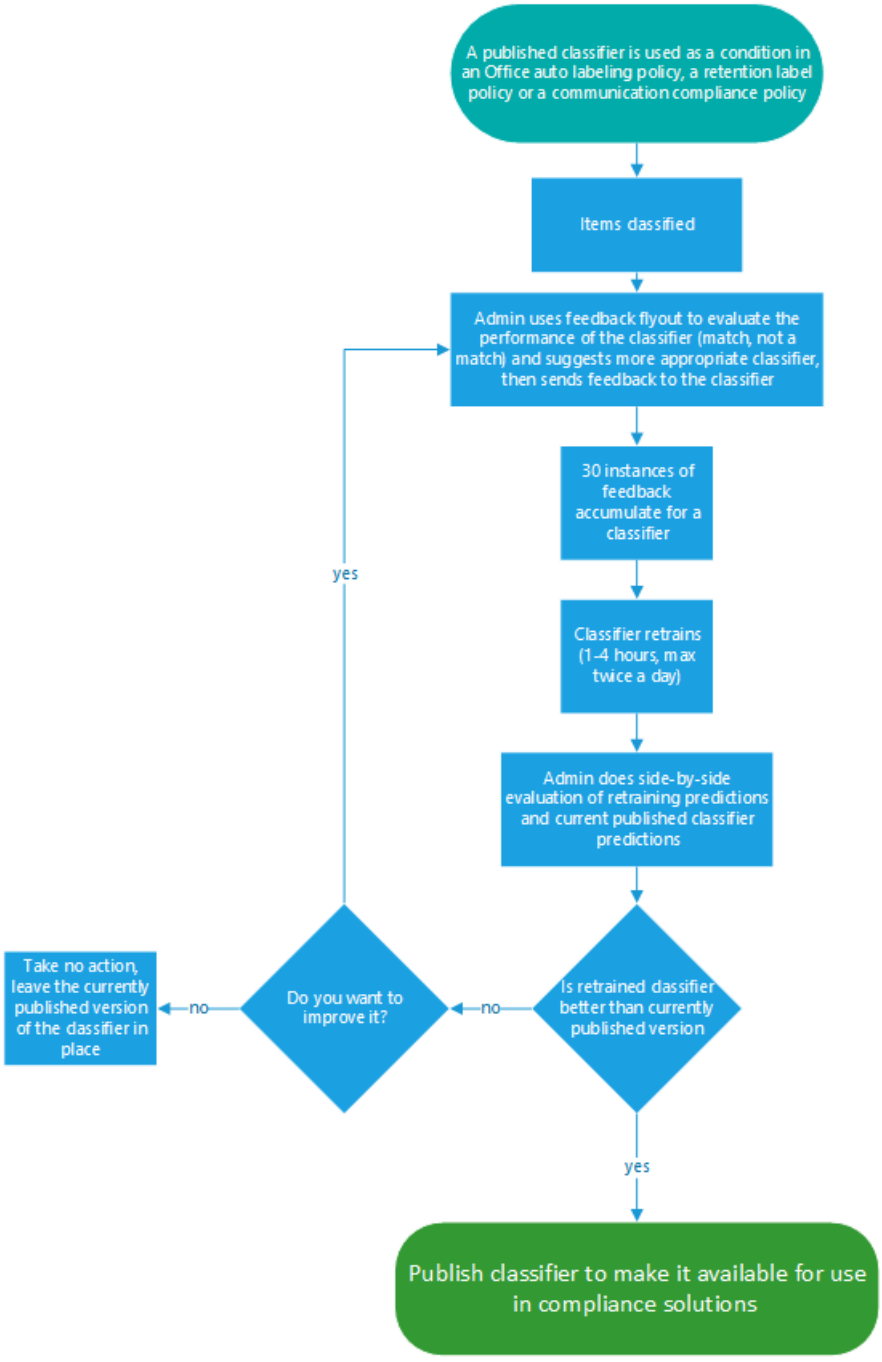

See also

- Retention labels
- Data loss prevention (DLP)
- Sensitivity labels
- Sensitive information type entity definitions
- Document finger printing
- Exact data match

Get started with trainable classifiers

2/18/2021 • 6 minutes to read • Edit Online

A Microsoft 365 trainable classifier is a tool you can train to recognize various types of content by giving it samples to look at. Once trained, you can use it to identify item for application of Office sensitivity labels, Communications compliance policies, and retention label policies.

Creating a custom trainable classifier first involves giving it samples that are human picked and positively match the category. Then, after it has processed those, you test the classifiers ability to predict by giving it a mix of positive and negative samples. This article shows you how to create and train a custom classifier and how to improve the performance of custom trainable classifiers and pre-trained classifiers over their lifetime through retraining.

To learn more about the different types of classifiers, see Learn about trainable classifiers.

Prerequisites

Licensing requirements

Classifiers are a Microsoft 365 E5, or E5 Compliance feature. You must have one of these subscriptions to make use of them.

Permissions

To access classifiers in the UI:

- the Global admin needs to opt in for the tenant to create custom classifiers.
- Compliance Administrator role is required to train a classifier.

You'll need accounts with these permissions to use classifiers in these scenarios:

- Retention label policy scenario: Record Management and Retention Management roles
- Sensitivity label policy scenario: Security Administrator, Compliance Administrator, Compliance Data Administrator
- Communication compliance policy scenario: Insider Risk Management Admin, Supervisory Review Administrator

> **IMPORTANT**
> By default, only the user who creates a custom classifier can train and review predictions made by that classifier.

Prepare for a custom trainable classifier

It's helpful to understand what's involved in creating a custom trainable classifier before you dive in.

Timeline

This timeline reflects a sample deployment of trainable classifiers.

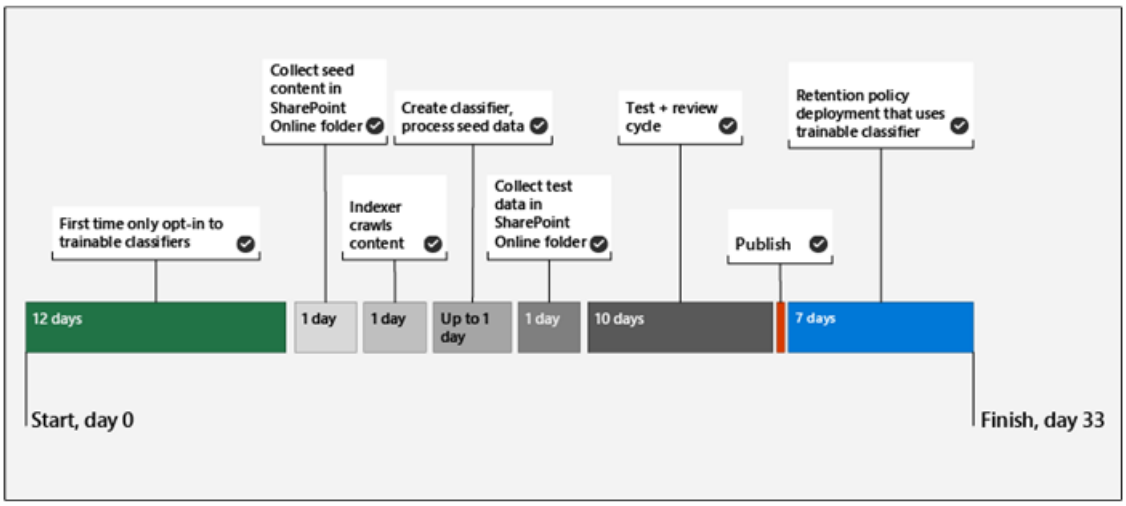

TIP

Opt-in is required the first time for trainable classifiers. It takes twelve days for Microsoft 365 to complete a baseline evaluation of your organizations content. Contact your global administrator to kick off the opt-in process.

Overall workflow

To understand more about the overall workflow of creating custom trainable classifiers, see Process flow for creating customer trainable classifiers.

Seed content

When you want a trainable classifier to independently and accurately identify an item as being in particular category of content, you first have to present it with many samples of the type of content that are in the category. This feeding of samples to the trainable classifier is known as *seeding*. Seed content is selected by a human and is judged to represent the category of content.

TIP

You need to have at least 50 positive samples and as many as 500. The trainable classifier will process up to the 500 most recent created samples (by file created date/time stamp). The more samples you provide, the more accurate the predictions the classifier will make.

Testing content

Once the trainable classifier has processed enough positive samples to build a prediction model, you need to test the predictions it makes to see if the classifier can correctly distinguish between items that match the category and items that don't. You do this by selecting another, hopefully larger, set of human picked content that consists of samples that should fall into the category and samples that won't. You should test with different data than the initial seed data you first provided. Once it processes those, you manually go through the results and verify whether each prediction is correct, incorrect, or you aren't sure. The trainable classifier uses this feedback to improve its prediction model.

TIP

For best results, have at least 200 items in your test sample set with an even distribution of positive and negative matches.

How to create a trainable classifier

1. Collect between 50-500 seed content items. These must be only samples that strongly represent the type of content you want the trainable classifier to positively identify as being in the classification category. See, Default crawled file name extensions and parsed file types in SharePoint Server for the supported file types.

   > **IMPORTANT**
   > The seed and test sample items must not be encrypted and they must be in English.

   > **IMPORTANT**
   > Make sure the items in your seed set are **strong** examples of the category. The trainable classifier initially builds its model based on what you seed it with. The classifier assumes all seed samples are strong positives and has no way of knowing if a sample is a weak or negative match to the category.

2. Place the seed content in a SharePoint Online folder that is dedicated to holding *the seed content only*. Make note of the site, library, and folder URL.

   > **TIP**
   > If you create a new site and folder for your seed data, allow at least an hour for that location to be indexed before creating the trainable classifier that will use that seed data.

3. Sign in to Microsoft 365 compliance center with compliance admin or security admin role access and open **Microsoft 365 compliance center** or **Microsoft 365 security center** > **Data classification**.
4. Choose the **Trainable classifiers** tab.
5. Choose **Create trainable classifier**.
6. Fill in appropriate values for the `Name` and `Description` fields of the category of items you want this trainable classifier to identify.
7. Pick the SharePoint Online site, library, and folder URL for the seed content site from step 2. Choose `Add`.
8. Review the settings and choose `Create trainable classifier`.
9. Within 24 hours the trainable classifier will process the seed data and build a prediction model. The classifier status is `In progress` while it processes the seed data. When the classifier is finished processing the seed data, the status changes to `Need test items`.
10. You can now view the details page by choosing the classifier.

Training process

Add items to test the classifier

Step	Items	Status
Add items to test the classifier	-	Ready to test
Review items to improve the classifier accuracy	-	Not available
Publish the classifier	-	Not available

Add items to test

11. Collect at least 200 test content items (10,000 max) for best results. These should be a mix of items that are strong positives, strong negatives and some that are a little less obvious in their nature. See, Default crawled file name extensions and parsed file types in SharePoint Server for the supported file types.

    IMPORTANT
    The sample items must not be encrypted and they must be in English.

12. Place the test content in a SharePoint Online folder that is dedicated to holding *the test content only*. Make note of the SharePoint Online site, library, and folder URL.

    TIP
    If you create a new site and folder for your test data, allow at least an hour for that location to be indexed before creating the trainable classifier that will use that seed data.

13. Choose `Add items to test`.
14. Pick the SharePoint Online site, library, and folder URL for the test content site from step 12. Choose `Add` .
15. Finish the wizard by choosing `Done`. Your trainable classifier will take up to an hour to process the test files.
16. When the trainable classifier is done processing your test files, the status on the details page will change to `Ready to review`. If you need to increase the test sample size, choose `Add items to test` and allow the trainable classifier to process the additional items.

Training process

Review items to generate classifier accuracy

Step	Items	Status
Add items to test the classifier	24	In progress
Review items to improve the classifier accuracy	0	Ready to review
Publish the classifier	-	Not available

17. Choose `Tested items to review` tab to review items.

18. Microsoft 365 will present 30 items at a time. Review them and in the `We predict this item is "Relevant". Do you agree?` box choose either `Yes` or `No` or `Not sure, skip to next item`. Model accuracy is automatically updated after every 30 items.

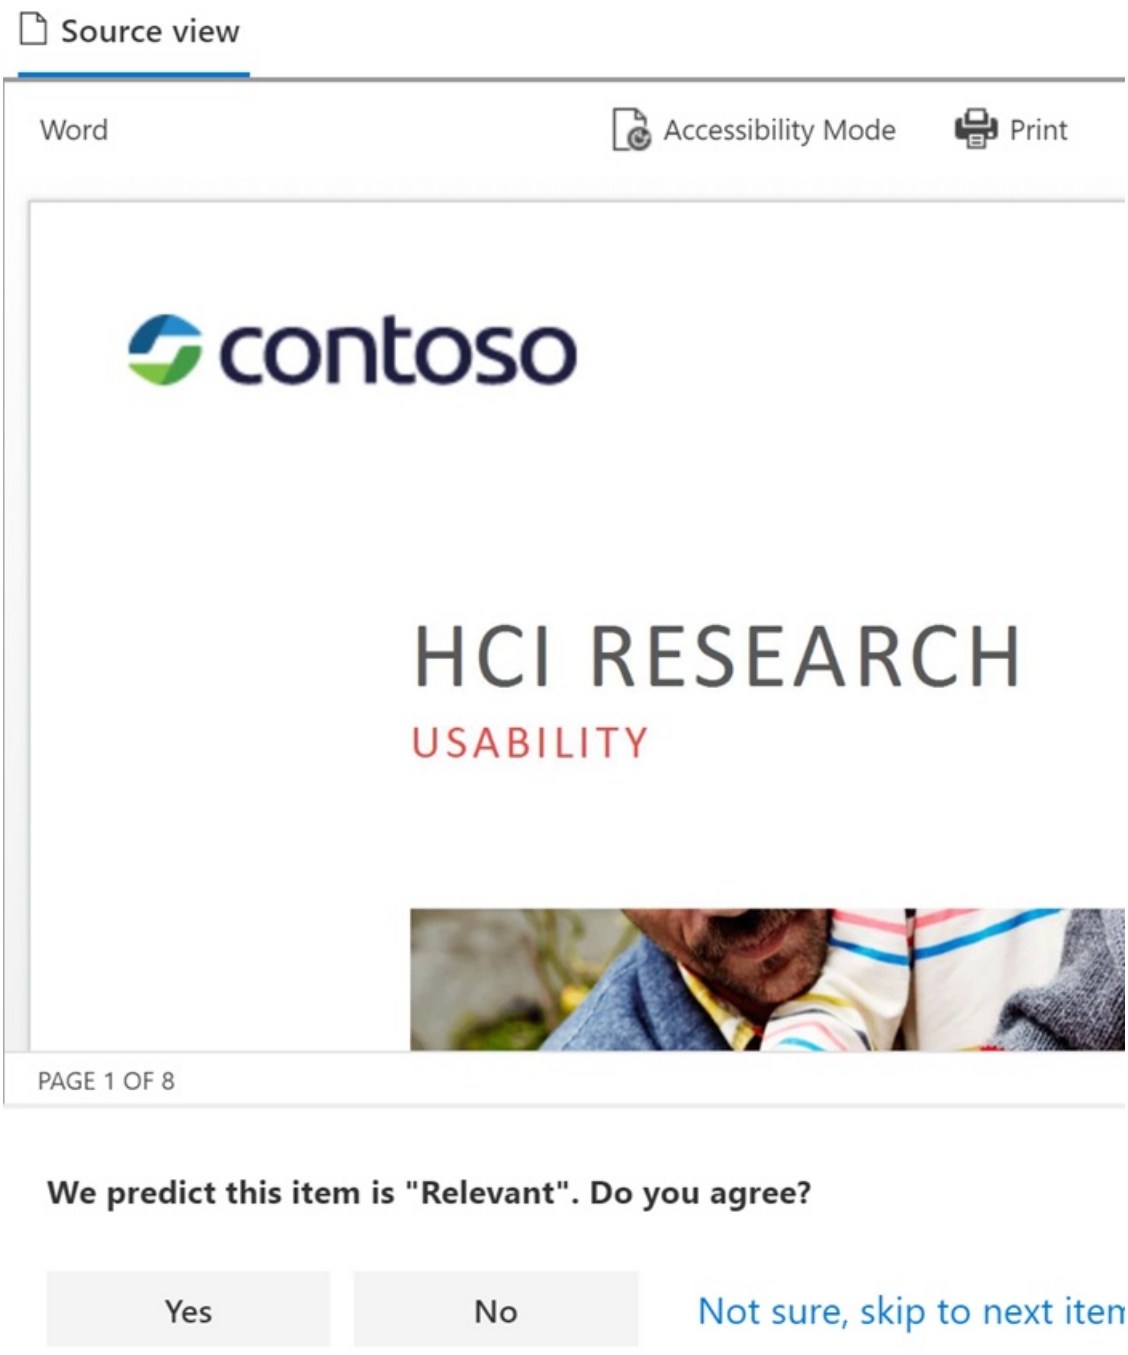

19. Review *at least* 200 items. Once the accuracy score has stabilized, the **publish** option will become available and the classifier status will say `Ready to use`.

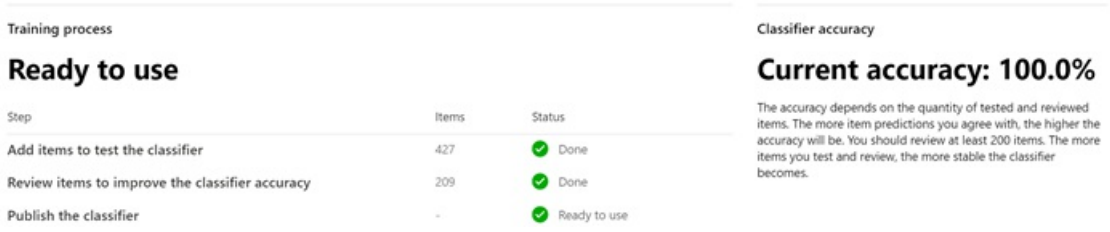

20. Publish the classifier.

21. Once published your classifier will be available as a condition in Office auto-labeling with sensitivity labels, auto-apply retention label policy based on a condition and in Communication compliance.

How to retrain a classifier in communications compliance

2/18/2021 • 3 minutes to read • Edit Online

A Microsoft 365 trainable classifier is a tool you can train to recognize various types of content by giving it samples to look at. Once trained, you can use it to identify item for application of Office sensitivity labels, communications compliance policies, and retention label policies.

This article shows you how to improve the performance of custom trainable classifiers and some pre-trained classifiers by providing them additional feedback.

To learn more about the different types of classifiers, see Learn about trainable classifiers.

Permissions

To access classifiers in the Microsoft 365 Compliance center:

- the Compliance admin role or Compliance Data Administrator is required to train a classifier

You'll need accounts with these permissions to use classifiers in these scenarios:

- Communication compliance policy scenario: Insider Risk Management Admin, Supervisory Review Administrator

Overall workflow

> **IMPORTANT**
>
> You provide feedback in the compliance solution that is using the classifier as a condition. **If you don't have a communications compliance policy that uses a classifier as a condition, stop here.**

As you use your classifiers, you may want to increase the precision of the classifications that they're making. You do this by evaluating the quality of the classifications made for items it has identified as being a match or not a match. After you make 30 evaluations for a classifier it takes that feedback and automatically retrains itself.

To understand more about the overall workflow of retraining a classifier, see Process flow for retraining a classifier.

> **NOTE**
>
> A classifier must already be published and in use before it can be retrained.

How to retrain a classifier in communication compliance policies

1. Open the Communication compliance policy that uses a classifier as a condition and choose one of the identified items from the **Pending** list.
2. Choose the ellipsis and **Improve classification**.
3. In the **Detailed feedback** pane, if the item is a true positive, choose, **Match**. If the item is a false positive, that is it was incorrectly included in the category, choose **Not a match**.

4. If there is another classifier that would be more appropriate for the item, you can choose it from the **Suggest other trainable classifiers** list. This will trigger the other classifier to evaluate the item.

> **TIP**
> You can provide feedback on multiple items simultaneously by choosing them all and then choosing **Provide detailed feedback** in the command bar.

5. Choose **Send feedback** to send your evaluation of the `match`, `not a match` classifications and suggest other trainable classifiers. When you've provided 30 instances of feedback to a classifier, it will automatically retrain. Retraining can take from 1-4 hours. Classifiers can only be retrained twice per day.

> **IMPORTANT**
> This information goes to the classifier in your tenant, **it does not go back to Microsoft**.

6. Open the **Data classification** page in the **Microsoft 365 compliance center**.
7. Open **Trainable classifiers**.
8. The classifier that was used in your Communications compliance policy will appear under the **Re-training** heading.

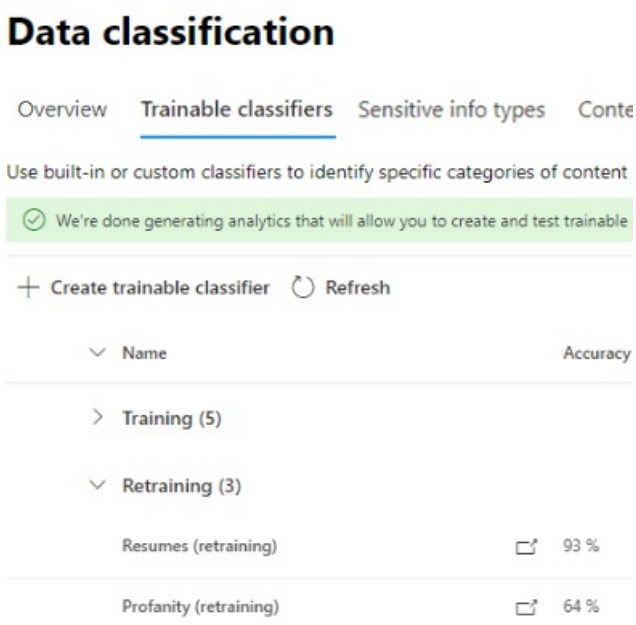

9. Once retraining completes, choose the classifier to open the retraining overview.

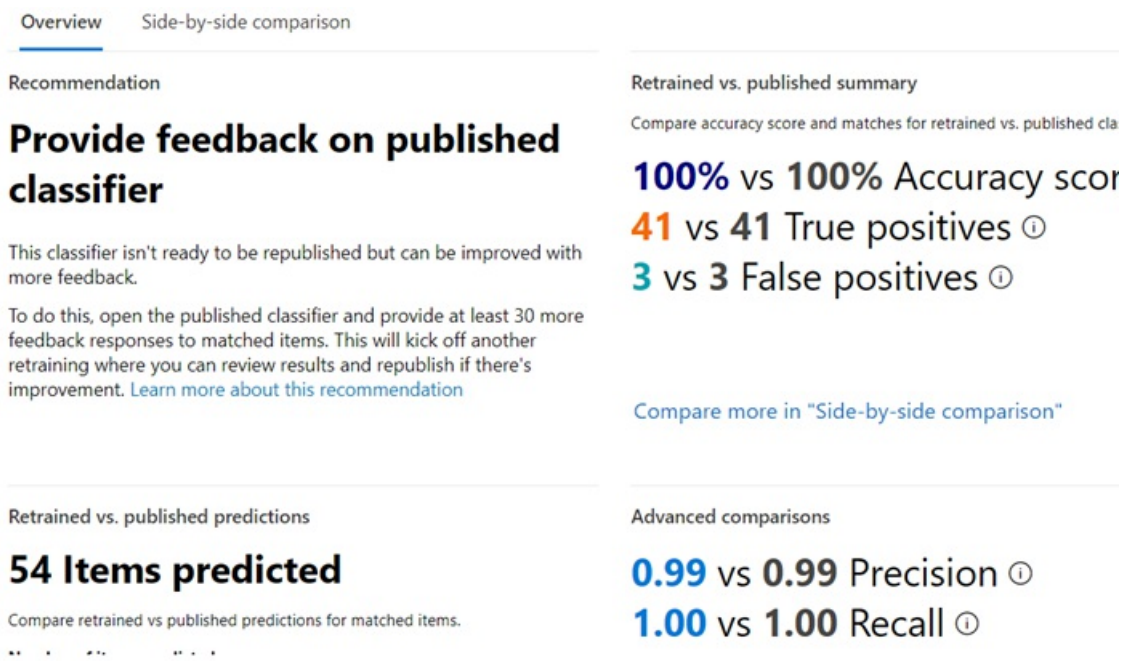

10. Review the recommended action, and the prediction comparisons of the retrained and currently published versions of the classifier.
11. If you satisfied with the results of the retraining, choose **Re-publish**.
12. If you are not satisfied with the results of the retraining, you can choose to provide additional feedback to the classifier in the Communications compliance interface and start another retraining cycle or do nothing in which case the currently published version of the classifier will continue to be used.

Details on republishing recommendations

Here is a little information on how we formulate the recommendation to re-publish a retrained classifier or suggest further retraining. This requires a little deeper understanding of how trainable classifiers work.

After a retrain, we evaluate the classifier's performance on both the items with feedback as well as any items originally used to train the classifier.

- For built-in models, items used to train the classifier are the items used by Microsoft to build the model.
- For custom models, items used in the original training the classifier are from the sites you had added for test and review.

We compare the performance numbers on both sets of items for the retrained and published classifier to provide a recommendation on whether there was improvement to republish.

See also

- Learn about trainable classifiers
- Default crawled file name extensions and parsed file types in SharePoint Server

How to retrain a classifier in content explorer

2/18/2021 • 4 minutes to read • Edit Online

A Microsoft 365 trainable classifier is a tool you can train to recognize various types of content by giving it samples to look at. Once trained, you can use it to identify item for application of Office sensitivity labels, communications compliance policies, and retention label policies.

This article shows you how to improve the performance of custom trainable classifiers and some pre-trained classifiers by providing them additional feedback.

To learn more about the different types of classifiers, see Learn about trainable classifiers.

Permissions

To access classifiers in the Microsoft 365 Compliance center:

- the Compliance admin role or Compliance Data Administrator is required to train a classifier

You'll need accounts with these permissions to use classifiers in these scenarios:

- Retention label policy scenario: Record Management and Retention Management roles

Overall workflow

> **IMPORTANT**
> You provide feedback in content explorer for auto-apply retention label policies to Exchange items and uses the classifier as a condition. **If you don't have a retention policy that auto-applies a retention label to Exchange items and uses a classifier as a condition, stop here.**

As you use your classifiers, you may want to increase the precision of the classifications that they're making. You do this by evaluating the quality of the classifications made for items it has identified as being a match or not a match. After you make 30 evaluations for a classifier it takes that feedback and automatically retrains itself.

To understand more about the overall workflow of retraining a classifier, see Process flow for retraining a classifier.

> **NOTE**
> A classifier must already be published and in use before it can be retrained.

How to retrain a classifier in content explorer

1. Sign in to Microsoft 365 compliance center with compliance admin or security admin role access and open **Microsoft 365 compliance center** > **Data classification** > **Content explorer**.
2. Under the **Filter on labels, info types, or categories** list, expand **Trainable classifiers**.

> **IMPORTANT**
> It can take up to eight days for aggregated items to appear under the trainable classifiers heading.

3. Choose the trainable classifier you used in you auto-apply retention label policy. This is the trainable classifier you will give feedback on.

> **NOTE**
> If an item has an entry in the **Retention label** column, it means that the item was classified as a `match`. If an item doesn't have an entry in the **Retention label** column, it means it was classified as a `close match`. You can improve the classifier precision the most by providing feedback on `close match` items.

4. Choose an item and open it.

> **TIP**
> You can provide feedback on multiple items simultaneously by choosing them all and then choosing **Improve classification** in the command bar.

5. Choose **Provide feedback**.
6. In the **Detailed feedback** pane, if the item is a true positive, choose, **Match**. If the item is a false positive, that is it was incorrectly included in the category, choose **Not a match**.
7. If there is another classifier that would be more appropriate for the item, you can choose it from the **Suggest other trainable classifiers** list. This will trigger the other classifier to evaluate the item.
8. Choose **Send feedback** to send your evaluation of the `match`, `not a match` classifications and suggest other trainable classifiers. When you've provided 30 instances of feedback to a classifier, it will automatically retrain. Retraining can take from one to four hours. Classifiers can only be retrained twice per day.

> **IMPORTANT**
> This information goes to the classifier in your tenant, **it does not go back to Microsoft**.

9. Open **Trainable classifiers**.
10. The classifier that was used in your Communications compliance policy will appear under the **Re-training** heading.

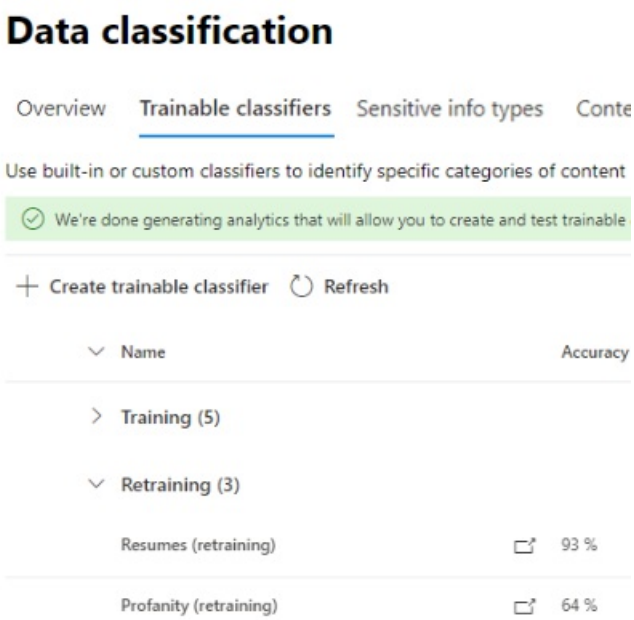

11. Once retraining completes, choose the classifier to open the retraining overview.

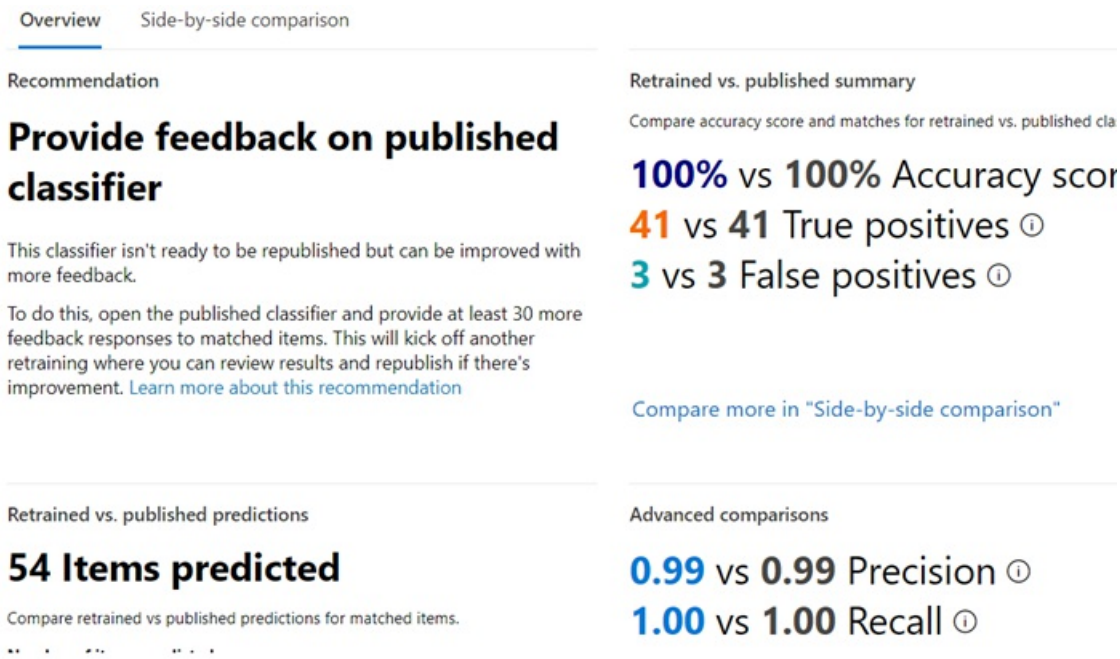

12. Review the recommended action, and the prediction comparisons of the retrained and currently published versions of the classifier.
13. If you satisfied with the results of the retraining, choose **Re-publish**.
14. If you are not satisfied with the results of the retraining, you can choose to provide additional feedback to the classifier in the Communications compliance interface and start another retraining cycle or do nothing in which case the currently published version of the classifier will continue to be used.

Details on republishing recommendations

Here is a little information on how we formulate the recommendation to re-publish a retrained classifier or suggest further retraining. This requires a little deeper understanding of how trainable classifiers work.

After a retrain, we evaluate the classifier's performance on both the items with feedback as well as any items originally used to train the classifier.

- For built-in models, items used to train the classifier are the items used by Microsoft to build the model.
- For custom models, items used in the original training the classifier are from the sites you had added for test and review.

We compare the performance numbers on both sets of items for the retrained and published classifier to provide a recommendation on whether there was improvement to republish.

See also

- Learn about trainable classifiers
- Default crawled file name extensions and parsed file types in SharePoint Server

Know your data - data classification overview

2/18/2021 • 5 minutes to read • Edit Online

As a Microsoft 365 administrator or compliance administrator, you can evaluate and then tag content in your organization in order to control where it goes, protect it no matter where it is and to ensure that it is preserved and deleted according to your organizations needs. You do this through the application of sensitivity labels, retention labels, and sensitive information type classification. There are various ways to do the discovery, evaluation and tagging, but the end result is that you may have very large number of documents and emails that are tagged and classified with one or both of these labels. After you apply your retention labels and sensitivity labels, you'll want to see how the labels are being used across your tenant and what is being done with those items. The data classification page provides visibility into that body of content, specifically:

- the number items that have been classified as a sensitive information type and what those classifications are
- the top applied sensitivity labels in both Microsoft 365 and Azure Information Protection
- the top applied retention labels
- a summary of activities that users are taking on your sensitive content
- the locations of your sensitive and retained data

You also manage these features on the data classification page:

- trainable classifiers
- sensitive information types

You can find data classification in the **Microsoft 365 compliance center** or **Microsoft 365 security center** > **Classification** > **Data Classification**.

Take a video tour of our data classification features.

Data classification will scan your sensitive content and labeled content before you create any policies. This is called **zero change management**. This lets you see the impact that all the retention and sensitivity labels are having in your environment and empower you to start assessing your protection and governance policy needs.

Prerequisites

A number of different subscriptions support Endpoint DLP. To see licensing options for Endpoint DLP see Information Protection licensing for guidance.

Permissions

In order to get access to the data classification page, an account must be assigned membership in any one of these roles or role groups.

Microsoft 365 role groups

- Global administrator
- Compliance administrator
- Security administrator
- Compliance data administrator

Sensitive information types used most in your content

Microsoft 365 comes with many definitions of sensitive information types, such as an item containing a social security number or a credit card number. For more information on sensitive information types, see Sensitive information type entity definitions.

The sensitive information type card shows the top sensitive information types that have been found and labeled across your organization.

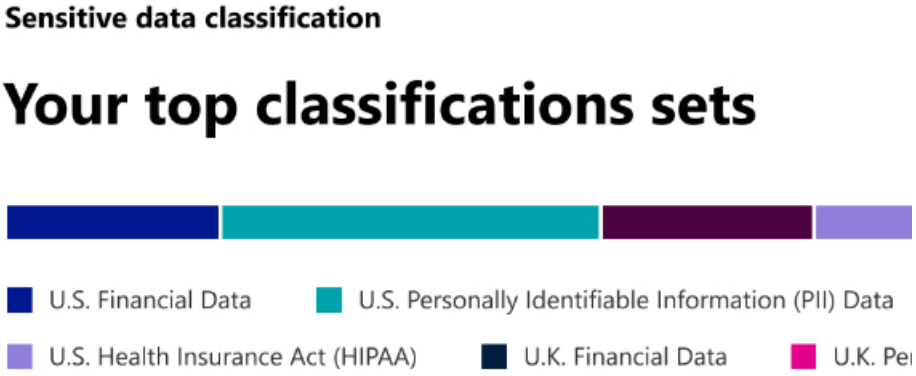

Explore data by industry regulators in the content explorer

To find out how many items are in any given classification category, hover over the bar for the category.

NOTE

If the card displays the message "No data found with sensitive information". It means that there are no items in your organization that have been classified as being a sensitive information type or no items that have been crawled. To get started with labels, see:

- Get started with sensitivity labels
- Get started with retention policies and retention labels
- Sensitive information type entity definitions

Top sensitivity labels applied to content

When you apply a sensitivity label to an item either through Microsoft 365 or Azure Information Protection (AIP), two things happen:

- a tag that indicates the value of the item to your org is embedded in the document and will follow it everywhere it goes
- the presence of the tag enables various protective behaviors, such as mandatory watermarking or encryption. With end point protection enabled you can even prevent an item from leaving your organizational control.

For more information on sensitivity labels, see: Learn about sensitivity labels

Sensitivity labels must be enabled for files that are in SharePoint and OneDrive in order for the corresponding data to surface in the data classification page. For more information, see Enable sensitivity labels for Office files in SharePoint and OneDrive.

The sensitivity label card shows the number of items (email or document) by sensitivity level.

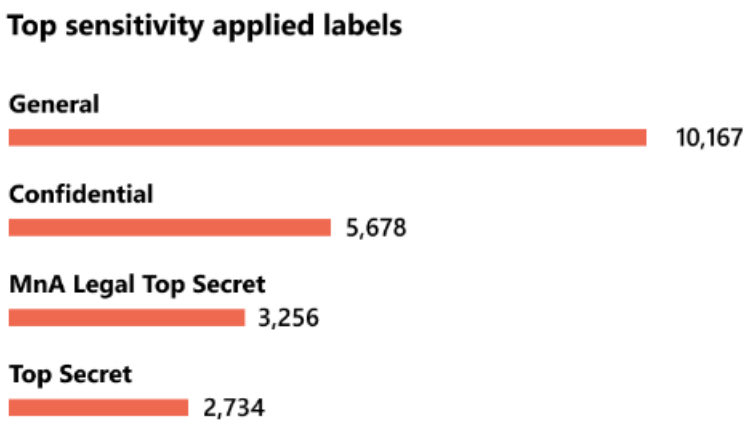

NOTE

If you haven't created or published any sensitivity labels or no content has had a sensitivity label applied, this card will display the message "No sensitivity labels detected". To get started with sensitivity labels, see:

- Get started with sensitivity labels or for AIP Configure the Azure information protection policy

Top retention labels applied to content

Retention labels are used to manage the retention and disposition of content in your organization. When applied, they can be used to control how an item will be kept before deletion, whether it should be reviewed prior to deletion, when its retention period expires, and whether it should be marked as a record. For more information, see Learn about retention policies and retention labels.

The top applied retention labels card shows you how many items have a given retention label.

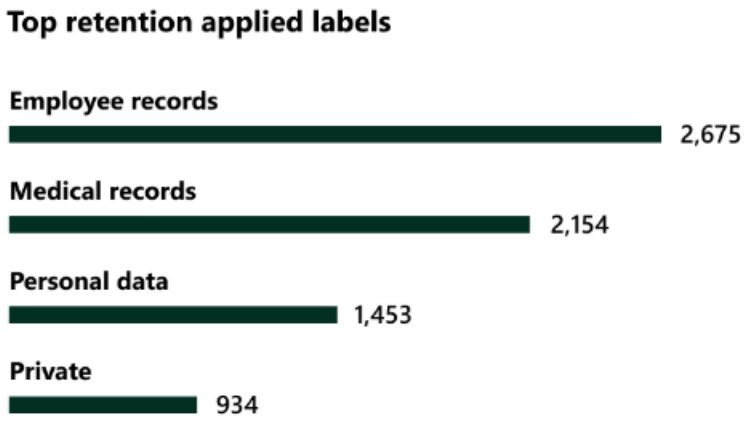

NOTE

If this card displays the message, "No retention labels detected", it means you haven't created or published any retention labels or no content has had a retention label applied. To get started with retention labels, see:

- Get started with retention policies and retention labels

Top activities detected

This card provides a quick summary of the most common actions that users are taking on the sensitivity labeled items. You can use the Activity explorer to drill deep down on eight different activities that Microsoft 365 tracks on labeled content and content that is located on Windows 10 endpoints.

> **NOTE**
>
> If this card displays the message, "No activity detected" it means that there's been no activity on the files or that user and admin auditing isn't turned on. To turn the audit logs on , see:
>
> - Search the audit log in security & compliance center

Sensitivity and retention labeled data by location

The point of the data classification reporting is to provide visibility into the number of items that have which label as well as their location. These cards let you know how many labeled items the are in Exchange, SharePoint, and OneDrive etc.

> **NOTE**
>
> If this card displays the message, "No locations detected, it means you haven't created or published any sensitivity labels or no content has had a retention label applied. To get started with sensitivity labels, see:
>
> - Sensitivity labels

See also

- View label activity
- View labeled content
- Learn about sensitivity labels
- Learn about retention policies and retention labels
- Sensitive information type entity definitions
- Learn about trainable classifiers (preview)

Get started with content explorer

2/18/2021 • 4 minutes to read • Edit Online

The data classification content explorer allows you to natively view the items that were summarized on the overview page.

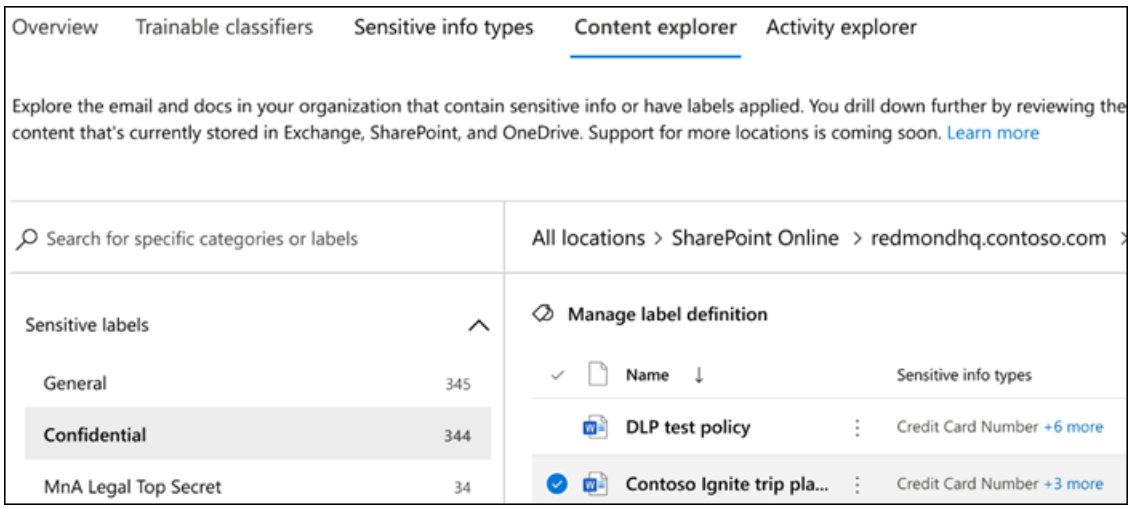

Prerequisites

Every account that accesses and uses data classification must have a license assigned to it from one of these subscriptions:

- Microsoft 365 (E5)
- Office 365 (E5)
- Advanced Compliance (E5) add-on
- Advanced Threat Intelligence (E5) add-on
- Microsoft 365 E5/A5 Info Protection & Governance
- Microsoft 365 E5/A5 Compliance

Permissions

In order to get access to the content explorer tab, an account must be assigned membership in any one of these roles or role groups.

Microsoft 365 role groups

- Global administrator
- Compliance administrator
- Security administrator
- Compliance data administrator

> **IMPORTANT**
> Membership in these role groups does not allow you to view the list of items in content explorer or to view the contents of the items in content explorer.

Required permissions to access items in content explorer

Access to content explorer is highly restricted because it lets you read the contents of scanned files.

> **IMPORTANT**
> These permissions supercede permissions that are locally assigned to the items, which allows viewing of the content.

There are two roles that grant access to content explorer and it is granted using the Microsoft Security & Compliance Center:

- **Content Explorer List viewer**: Membership in this role group allows you to see each item and its location in list view. The `data classification list viewer` role has been pre-assigned to this role group.
- **Content Explorer Content viewer**: Membership in this role group allows you to view the contents of each item in the list. The `data classification content viewer` role has been pre-assigned to this role group.

The account you use to access content explorer must be in one or both of the role groups. These are independent role groups and are not cumulative. For example, if you want to grant an account the ability to view the items and their locations only, grant Content Explorer List viewer rights. If you want that same account to also be able to view the contents of the items in the list, grant Content Explorer Content viewer rights as well.

You can also assign either or both of the roles to a custom role group to tailor access to content explorer.

A Global admin, Compliance admin, or Data admin can assign the necessary Content Explorer List Viewer, and Content Explorer Content Viewer role group membership.

Content explorer

Content explorer shows a current snapshot of the items that have a sensitivity label, a retention label or have been classified as a sensitive information type in your organization.

Sensitive information types

A DLP policy can help protect sensitive information, which is defined as a **sensitive information type**. Microsoft 365 includes definitions for many common sensitive information types from across many different regions that are ready for you to use. For example, a credit card number, bank account numbers, national ID numbers, and Windows Live ID service numbers.

> **NOTE**
> Content explorer doesn't currently scan for sensitive information types in Exchange Online.

Sensitivity labels

A sensitivity label is simply a tag that indicates the value of the item to your organization. It can be applied manually, or automatically. Once applied it gets embedded in the document and will follow it everywhere it goes. A sensitivity label enables various protective behaviors, such as mandatory watermarking or encryption.

Sensitivity labels must be enabled for files that are in SharePoint and OneDrive in order for the corresponding data to surface in the data classification page. For more information, see Enable sensitivity labels for Office files in SharePoint and OneDrive.

Retention labels

A retention label allows you to define how long a labeled item is kept and the steps to be taken prior to deleting it. They are applied manually or automatically via policies. They can play a role in helping your organization stay in compliance with legal and regulatory requirements.

How to use content explorer

1. Open **Microsoft 365 compliance center** > **Data classification** > **Content explorer**.
2. If you know the name of the label, or the sensitive information type, you can type that into the filter box.
3. Alternately, you can browse for the item by expanding the label type and selecting the label from the list.
4. Select a location under **All locations** and drill down the folder structure to the item.
5. Double-click to open the item natively in content explorer.

Export

The **export** control will create a .csv file that contains a listing of whatever is showing in the **All locations** pane.

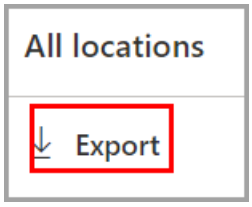

Search

When you drill down into a location, such as an Exchange folder, or a SharePoint or OneDrive site, the **search** tool appears.

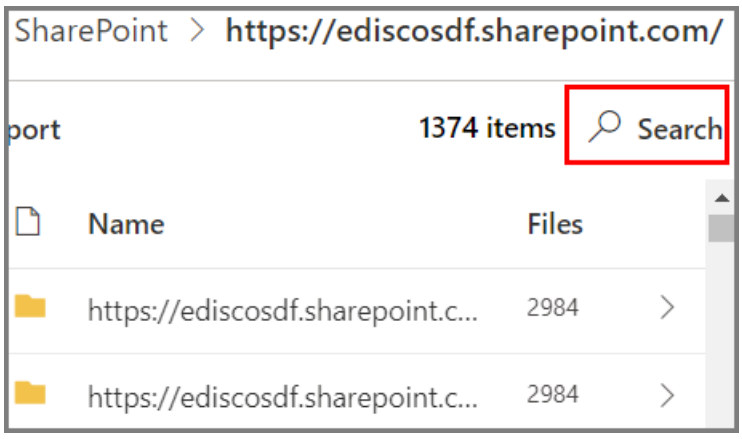

The scope of the search tool is what is displaying in the **All locations** pane and what you can search on varies depending on the selected location.

When **Exchange** is the selected location, you can search on the full email address of the mailbox, for example `user@domainname.com`.

When either **SharePoint** or **OneDrive** are selected location, the search tool will appear when you drill down to site names, folders and files.

> **NOTE**
>
> **OneDrive** We have listened to your valuable feedback on OneDrive integration during our preview program. Based on that feedback, the OneDrive functionality will remain in preview till all fixes are in place. Depending on your tenant, some customers may not see OneDrive as a location. We appreciate your continued support on this.

You can search on:

VALUE	EXAMPLE
full site name	`https://contoso.onmicrosoft.com/sites/sitename`
root folder name - gets all subfolders	`/sites`

VALUE	EXAMPLE
file name	`RES_Resume_1234.txt`
text at the beginning of file name	`RES`
text after an underscore character (_) in file name	`Resume` or `1234`
file extension	`txt`

See also

- Learn about sensitivity labels
- Learn about retention policies and retention labels
- Sensitive information type entity definitions.md
- Overview of data loss prevention

Get started with activity explorer

2/18/2021 • 2 minutes to read • Edit Online

The data classification overview and content explorer tabs give you visibility into what content has been discovered and labeled, and where that content is. Activity explorer rounds out this suite of functionality by allowing you to monitor what's being done with your labeled content. Activity explorer provides a historical view.

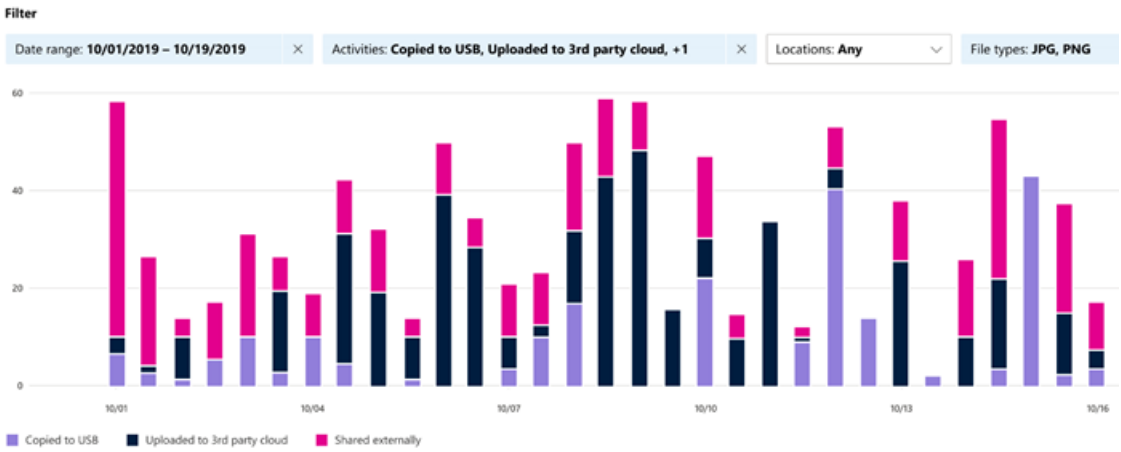

There are over 30 different filters available for use, some are:

- date range
- activity type
- location
- user
- sensitivity label
- retention label
- file path
- DLP policy

Prerequisites

Every account that accesses and uses data classification must have a license assigned to it from one of these subscriptions:

- Microsoft 365 (E5)
- Office 365 (E5)
- Advanced Compliance (E5) add-on
- Advanced Threat Intelligence (E5) add-on
- Microsoft 365 E5/A5 Info Protection & Governance
- Microsoft 365 E5/A5 Compliance

Permissions

In order to get access to the activity explorer tab, an account must be assigned membership in any one of these roles or role groups.

Microsoft 365 role groups

- Global administrator

- Compliance administrator
- Security administrator
- Compliance data administrator

Activity type

Microsoft 365 monitors and reports on types of activities across SharePoint Online, and OneDrive like:

- label applied
- label changed (upgraded, downgraded, or removed)
- auto-labeling simulation

The value of understanding what actions are being taken with your sensitive labeled content is that you can see if the controls that you have already put into place, such as data loss prevention policies are effective or not. If not, or if you discover something unexpected, such as a large number of items that are labeled `highly confidential` and are downgraded `general`, you can manage your various policies and take new actions to restrict the undesired behavior.

> **NOTE**
> Activity explorer doesn't currently monitor retention activities for Exchange Online.

See also

- Learn about sensitivity labels
- Learn about retention policies and retention labels
- Sensitive information type entity definitions

Data classification release notes

2/18/2021 • 2 minutes to read • Edit Online

Exchange mailbox count

You will notice a small tool tip appear when you drill into Exchange mailboxes. This is to call out the fact that the aggregate count displayed for sensitive information type, sensitivity label and retention label may not exactly match the number of items that you will find inside the mailbox. This is because the drill-down into the folder fetches the live view of content, which is classified, while the aggregated count is calculated.

Rendering of encrypted documents

SharePoint, Exchange, and OneDrive files that are encrypted don't render in the content explorer. This is a sensitive issue that requires a balance between the need to see file contents in content explorer and the need to keep the contents encrypted. With the permissions granted by **Content Explorer List Viewer**, and **Content Explorer Content Viewer** role groups, you will see a list view of the files, the file metadata, and a link you can use to access the content via the web client.

Supported characters in retention label names in SharePoint search

SharePoint search doesn't support retention label names with `-`, or `_` in them. For example, `Label-MIP` and `Label_MIP` aren't supported. SharePoint search does support those characters in sensitivity label names and sensitive information type names.

OneDrive remains in preview

Thanks for your valuable feedback on OneDrive integration during our preview program. As we work through the specifics, you may run into inconsistent data / flows. We'll continue to showcase OneDrive in preview until all fixes are in place. We appreciate your continued support.

See also

- Get started with data classification (preview)
- View label activity (preview)
- View labeled content (preview)
- Learn about sensitivity labels
- Learn about retention policies and retention labels
- Sensitive information type entity definitions

Learn about sensitivity labels

2/18/2021 • 17 minutes to read • Edit Online

Microsoft 365 licensing guidance for security & compliance.

To get their work done, people in your organization collaborate with others both inside and outside the organization. This means that content no longer stays behind a firewall—it can roam everywhere, across devices, apps, and services. And when it roams, you want it to do so in a secure, protected way that meets your organization's business and compliance policies.

Sensitivity labels from the Microsoft Information Protection solution let you classify and protect your organization's data, while making sure that user productivity and their ability to collaborate isn't hindered.

Example showing available sensitivity labels in Excel, from the **Home** tab on the Ribbon. In this example, the applied label displays on the status bar:

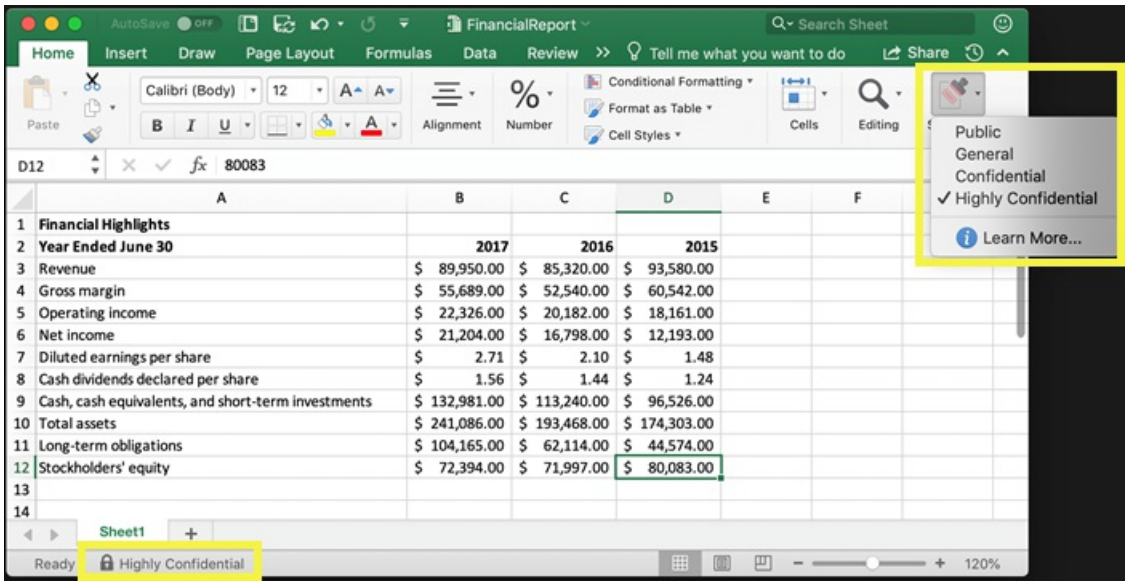

To apply sensitivity labels, users must be signed in with their Microsoft 365 work or school account.

> **NOTE**
>
> Sensitivity labels are newly supported For US Government tenants (GCC and GCC-H). For more information, see the release notes for Microsoft 365 Apps for enterprise, Version 2101: January 26.
>
> For the Azure Information Protection unified labeling client and scanner, see Azure Information Protection Premium Government Service Description.

You can use sensitivity labels to:

- **Provide protection settings that include encryption and content markings.** For example, apply a "Confidential" label to a document or email, and that label encrypts the content and applies a "Confidential" watermark. Content markings include headers and footers as well as watermarks, and encryption can also restrict what actions authorized people can take on the content.
- **Protect content in Office apps across different platforms and devices.** Supported by Word, Excel, PowerPoint, and Outlook on the Office desktop apps and Office on the web. Supported on Windows, macOS, iOS, and Android.

- **Protect content in third-party apps and services** by using Microsoft Cloud App Security. With Cloud App Security, you can detect, classify, label, and protect content in third-party apps and services, such as SalesForce, Box, or DropBox, even if the third-party app or service does not read or support sensitivity labels.
- **Protect containers** that include Teams, Microsoft 365 Groups, and SharePoint sites. For example, set privacy settings, external user access and external sharing, and access from unmanaged devices.
- **Extend sensitivity labels to Power BI**: When you turn on this capability, you can apply and view labels in Power BI, and protect data when it's saved outside the service.
- **Extend sensitivity labels to assets in Azure Purview**: When you turn on this capability, currently in preview, you can apply your sensitivity labels to assets such as SQL columns, files in Azure Blob Storage, and more.
- **Extend sensitivity labels to third-party apps and services.** Using the Microsoft Information Protection SDK, third-party apps can read sensitivity labels and apply protection settings.
- **Classify content without using any protection settings.** You can also simply assign a label as a result of classifying the content. This provides users with a visual mapping of classification to your organization's label names, and can use the labels to generate usage reports and see activity data for your sensitive content. Based on this information, you can always choose to apply protection settings later.

In all these cases, sensitivity labels in Microsoft 365 can help you take the right actions on the right content. With sensitivity labels, you can classify data across your organization, and enforce protection settings based on that classification.

For more information about these and other scenarios that are supported by sensitivity labels, see Common scenarios for sensitivity labels. New features are being developed all the time that support sensitivity labels, so you might also find it useful to reference the Microsoft 365 roadmap.

What a sensitivity label is

When you assign a sensitivity label to content, it's like a stamp that's applied and is:

- **Customizable.** Specific to your organization and business needs, you can create categories for different levels of sensitive content in your organization. For example, Personal, Public, General, Confidential, and Highly Confidential.
- **Clear text.** Because a label is stored in clear text in the metadata for files and emails, third-party apps and services can read it and then apply their own protective actions, if required.
- **Persistent.** Because the label is stored in metadata for files and emails, the label roams with the content, no matter where it's saved or stored. The unique label identification becomes the basis for applying and enforcing the policies that you configure.

When viewed by users, a sensitivity label appears like a tag on apps that they use and can be easily integrated into their existing workflows.

Each item that supports sensitivity labels can have a single sensitivity label applied to it. Documents and emails can have both a sensitivity label and a retention label applied to them.

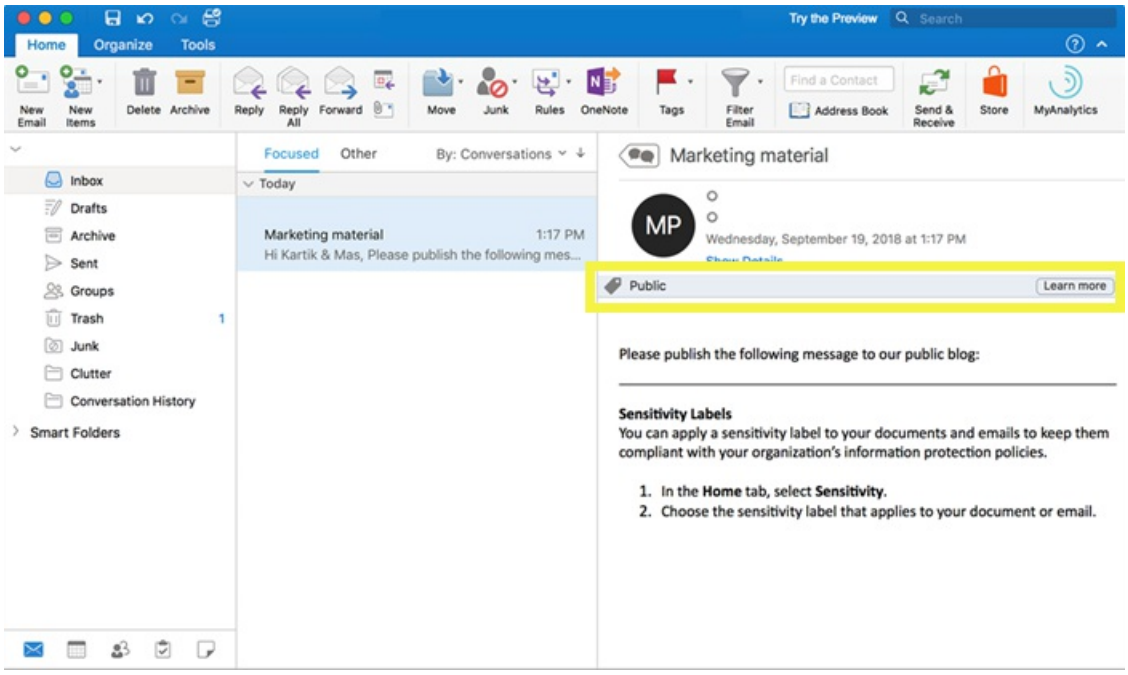

What sensitivity labels can do

After a sensitivity label is applied to an email or document, any configured protection settings for that label are enforced on the content. You can configure a sensitivity label to:

- **Encrypt** emails and documents to prevent unauthorized people from accessing this data. You can additionally choose which users or group have permissions to perform which actions and for how long. For example, you can choose to allow all users in your organization to modify a document while a specific group in another organization can only view it. Alternatively, instead of administrator-defined permissions, you can allow your users to assign permissions to the content when they apply the label.

  For more information about the **Encryption** settings when you create or edit a sensitivity label, see Restrict access to content by using encryption in sensitivity labels.

- **Mark the content** when you use Office apps, by adding watermarks, headers, or footers to email or documents that have the label applied. Watermarks can be applied to documents but not email. Example header and watermark:

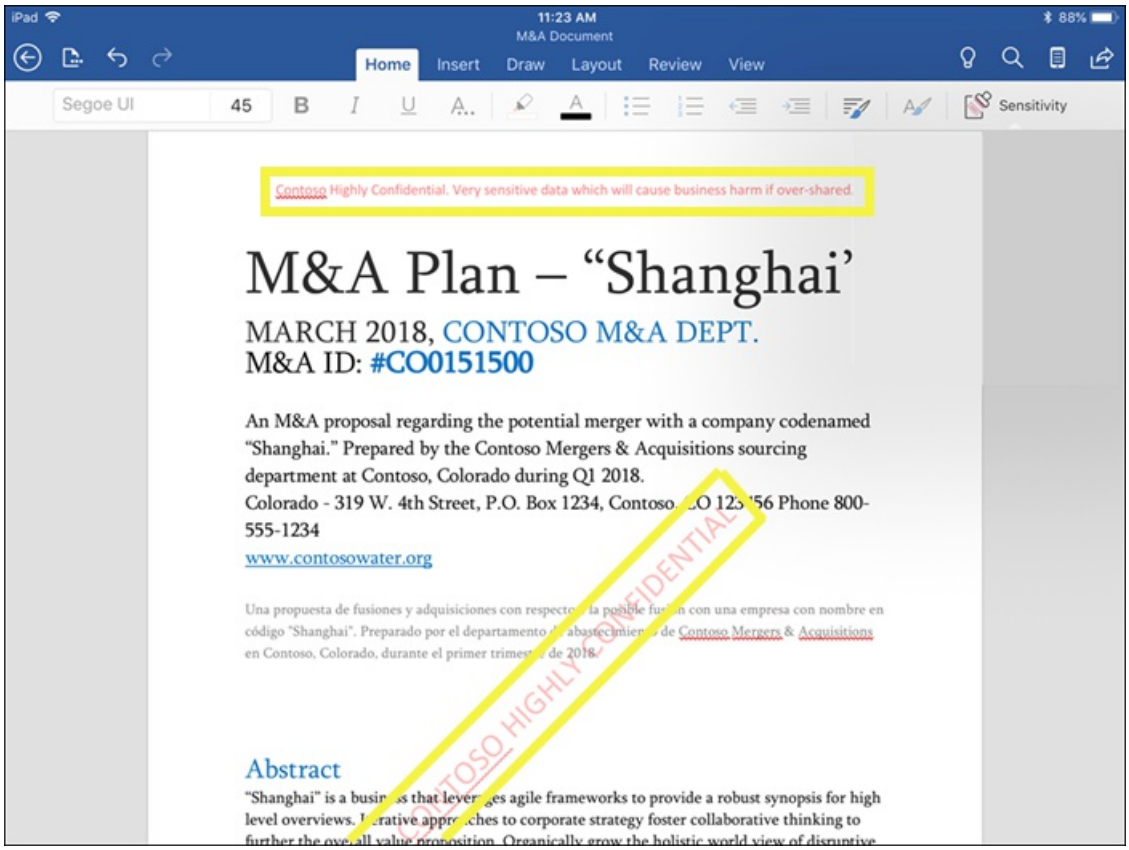

Need to check when content markings are applied? See When Office apps apply content marking and encryption.

Some, but not all apps support dynamic markings by using variables. For example, insert the label name or document name into the header, footer, or watermark. For more information, see Dynamic markings with variables.

String lengths: Watermarks are limited to 255 characters. Headers and footers are limited to 1024 characters, except in Excel. Excel has a total limit of 255 characters for headers and footers but this limit includes characters that aren't visible, such as formatting codes. If that limit is reached, the string you enter is not displayed in Excel.

- **Protect content in containers such as sites and groups** when you enable the capability to use sensitivity labels with Microsoft Teams, Microsoft 365 groups, and SharePoint sites.

  You can't configure protection settings for groups and sites until you enable this capability. This label configuration doesn't result in documents or emails being automatically labeled but instead, the label settings protect content by controlling access to the container where content can be stored. These settings include privacy settings, external user access and external sharing, and access from unmanaged devices.

- **Apply the label automatically to files and emails, or recommend a label.** Choose how to identify sensitive information that you want labeled, and the label can be applied automatically, or you can prompt users to apply the label that you recommend. If you recommend a label, the prompt displays whatever text you choose. For example:

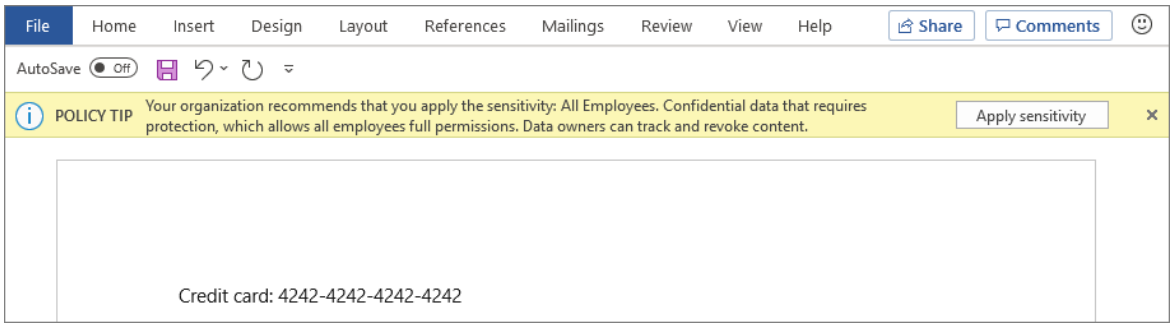

For more information about the **Auto-labeling for files and emails** settings when you create or edit a sensitivity label, see Apply a sensitivity label to content automatically for Office apps, and Automatically label your data in Azure Purview.

Label scopes

When you create a sensitivity label, you're asked to configure the label's scope which determines two things:

- Which label settings you can configure for that label
- Where the label will be visible to users

This scope configuration lets you have sensitivity labels that are just for documents and emails and can't be selected for containers. And similarly, sensitivity labels that are just for containers and can't be selected for documents and emails. New, and currently in preview, you can also select the scope for Azure Purview assets:

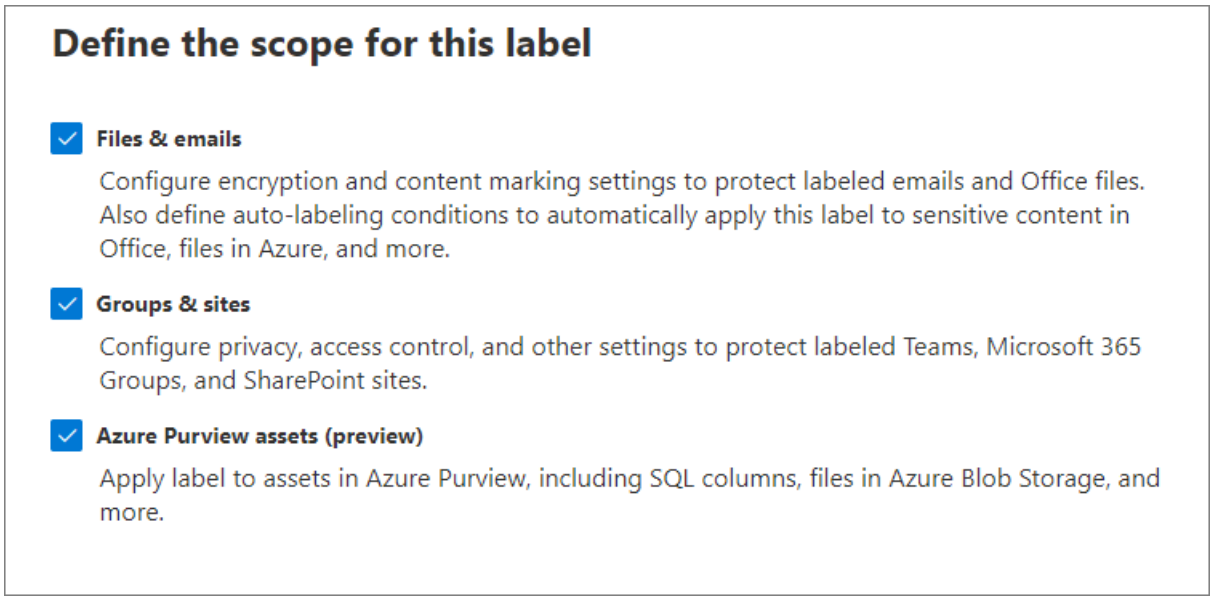

By default, the **Files & emails** scope is always selected. The other scopes are selected by default when the features are enabled for your tenant:

- **Groups & sites**: Enable sensitivity labels for containers and synchronize labels
- **Azure Purview assets (preview)**: Automatically label your content in Azure Purview

If you change the defaults so not all scopes are selected, you see the first page of the configuration settings for scopes you haven't selected, but you can't configure the settings. For example, if the scope for files and emails is not selected, you can't select the options on the next page:

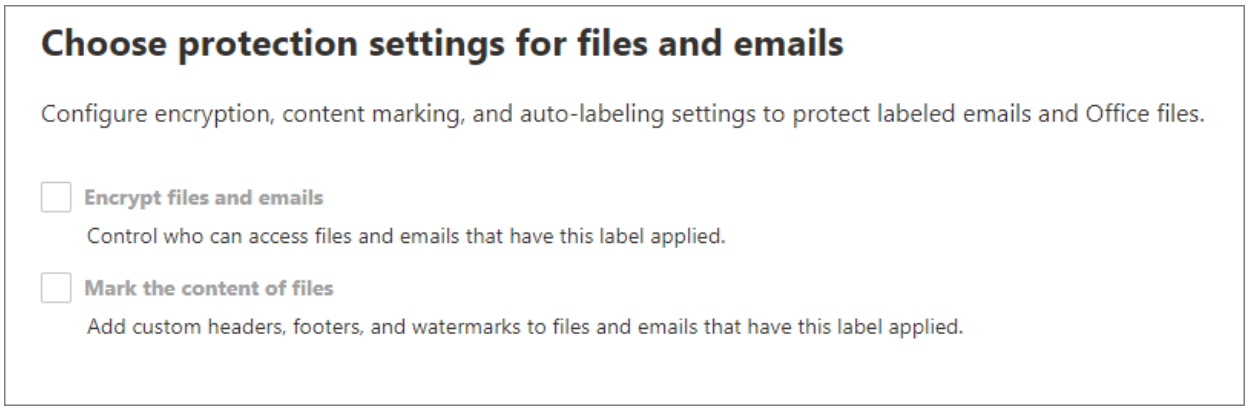

For these pages that have unavailable options, select **Next** to continue. Or, select **Back** to change the label's scope.

Label priority (order matters)

When you create your sensitivity labels in your admin center, they appear in a list on the **Sensitivity** tab on the **Labels** page. In this list, the order of the labels is important because it reflects their priority. You want your most restrictive sensitivity label, such as Highly Confidential, to appear at the **bottom** of the list, and your least restrictive sensitivity label, such as Public, to appear at the **top**.

You can apply just one sensitivity label to an item such as a document, email, or container. If you set an option that requires your users to provide a justification for changing a label to a lower classification, the order of this list identifies the lower classifications. However, this option does not apply to sublabels.

The ordering of sublabels is used with automatic labeling, though. When you configure labels to be applied automatically or as a recommendation, multiple matches can result for more than one label. To determine the label to apply or recommend, the label ordering is used: The last sensitive label is selected, and then if applicable, the last sublabel.

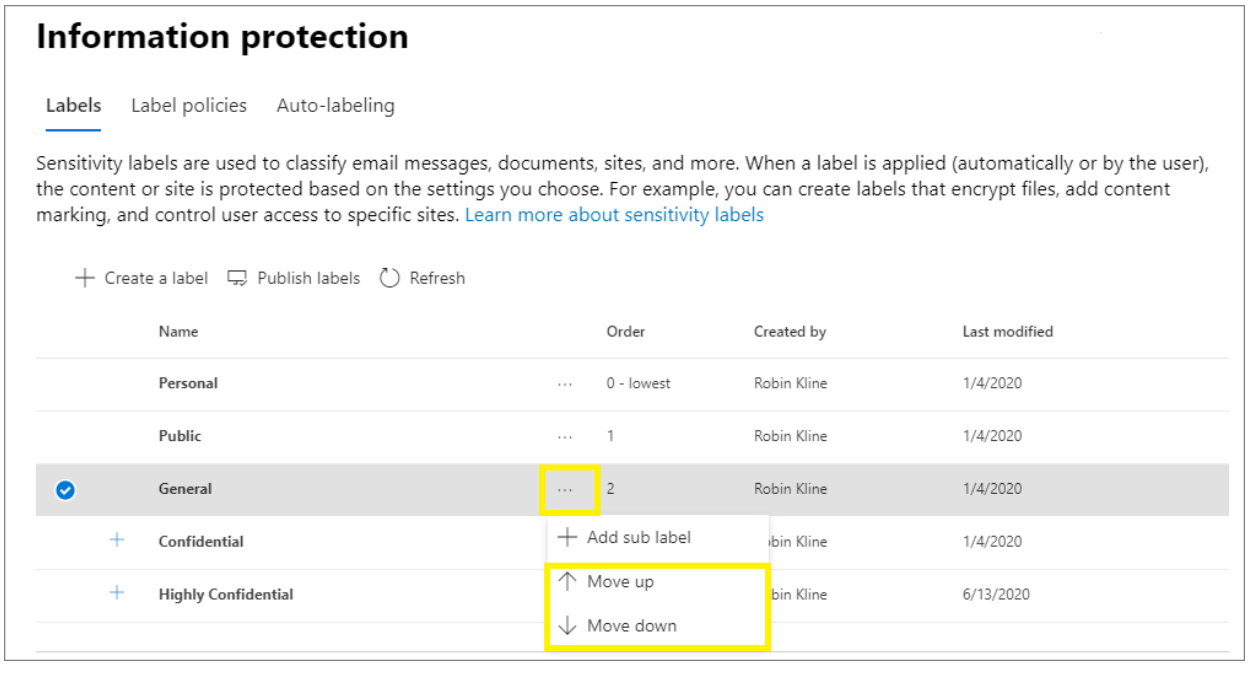

Sublabels (grouping labels)

With sublabels, you can group one or more labels below a parent label that a user sees in an Office app. For example, under Confidential, your organization might use several different labels for specific types of that classification. In this example, the parent label Confidential is simply a text label with no protection settings, and because it has sublabels, it can't be applied to content. Instead, users must choose Confidential to view the sublabels, and then they can choose a sublabel to apply to content.

Sublabels are simply a way to present labels to users in logical groups. Sublabels don't inherit any settings from

their parent label. When you publish a sublabel for a user, that user can then apply that sublabel to content but can't apply just the parent label.

Don't choose a parent label as the default label, or configure a parent label to be automatically applied (or recommended). If you do, the parent label won't be applied to content.

Example of how sublabels display for users:

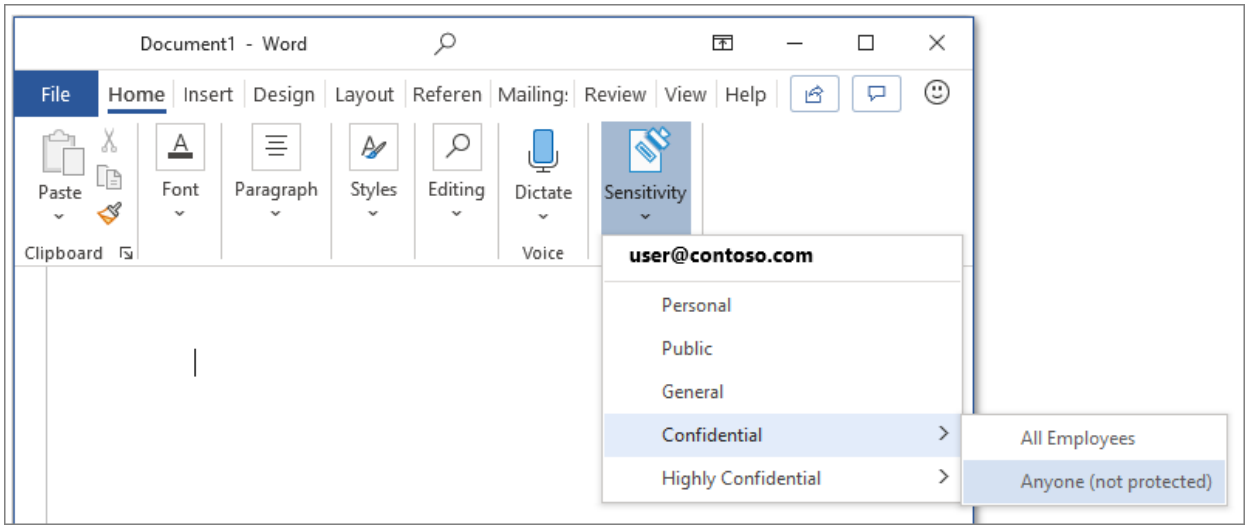

Editing or deleting a sensitivity label

If you delete a sensitivity label from your admin center, the label is not automatically removed from content, and any protection settings continue to be enforced on content that had that label applied.

If you edit a sensitivity label, the version of the label that was applied to content is what's enforced on that content.

What label policies can do

After you create your sensitivity labels, you need to publish them, to make them available to people and services in your organization. The sensitivity labels can then be applied to Office documents and emails, and other items that support sensitivity labels.

Unlike retention labels, which are published to locations such as all Exchange mailboxes, sensitivity labels are published to users or groups. Apps that support sensitivity labels can then display them to those users and groups as applied labels, or as labels that they can apply.

When you configure a label policy, you can:

- **Choose which users and groups see the labels.** Labels can be published to any specific user or email-enabled security group, distribution group, or Microsoft 365 group (which can have dynamic membership) in Azure AD.
- **Apply a default label** to all new documents and email created by the users and groups included in the label policy, and the same or different default label to containers (if you've enabled sensitivity labels for Microsoft Teams, Microsoft 365 groups, and SharePoint sites). Users can always change the default label if it's not the right label for their document or email.

  Consider using a default label to set a base level of protection settings that you want applied to all your content. However, without user training and other controls, this setting can also result in inaccurate labeling. It's usually not a good idea to select a label that applies encryption as a default label to documents. For example, many organizations need to send and share documents with external users who might not have apps that support the encryption or they might not use an account that can be authorized. For more information about this scenario, see Sharing encrypted documents with external

users.

- **Require a justification for changing a label.** If a user tries to remove a label or replace it with a label that has a lower-order number, you can require the user provides a justification to perform this action. For example, a user opens a document labeled Confidential (order number 3) and replaces that label with one named Public (order number 1). Currently, the justification reason is used only by the Azure Information Protection unified labeling client, which sends this information to Azure Information Protection analytics.

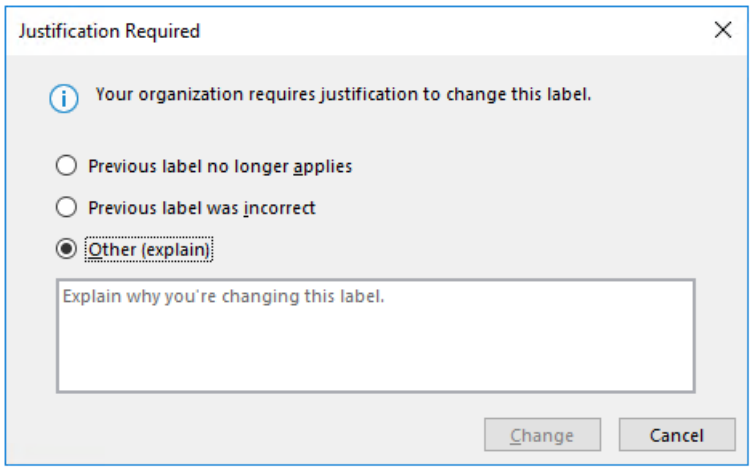

- **Require users to apply a label** with one option for email and documents, and another for containers. Also known as mandatory labeling, these options ensure a label must be applied before users can save documents and send emails, and create new groups or sites.

  For documents and emails, a label can be assigned manually by the user, automatically as a result of a condition that you configure, or be assigned by default (the default label option previously described). An example prompt shown in Outlook when a user is required to assign a label:

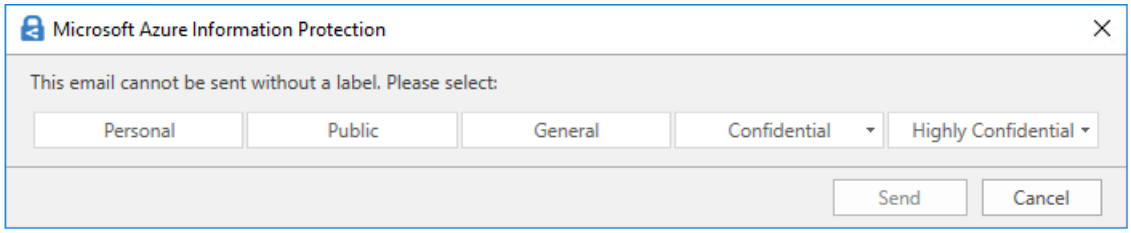

> **NOTE**
> Mandatory labeling for documents and emails isn't available for all apps or all platforms. For more information, see Require users to apply a label to their email and documents.

  For containers, a label must be assigned at the time the group or site is created.

  Consider using this option to help increase your labeling coverage. However, without user training, these settings can result in inaccurate labeling. In addition, unless you also set a corresponding default label, mandatory labeling can frustrate your users with the frequent prompts.

- **Provide help link to a custom help page.** If your users aren't sure what your sensitivity labels mean or how they should be used, you can provide a Learn More URL that appears at the bottom of the **Sensitivity label** menu in the Office apps:

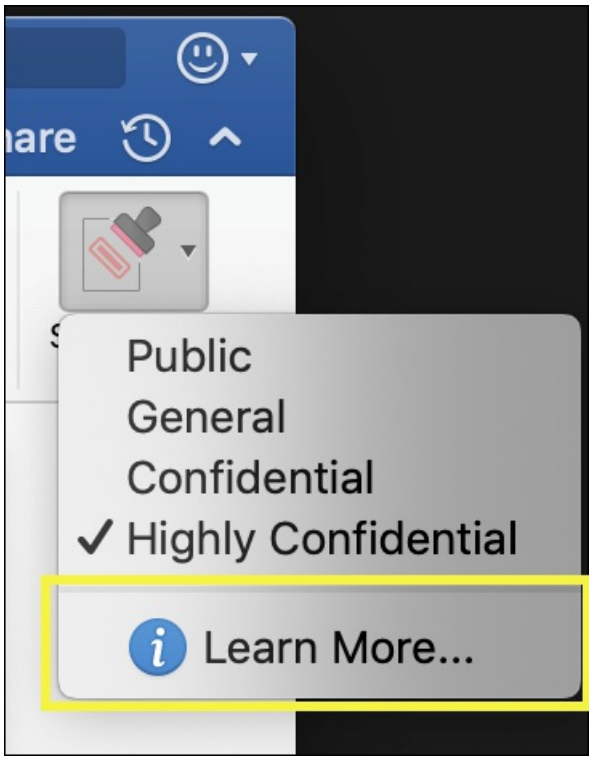

After you create a label policy that assigns new sensitivity labels to users and groups, users start to see those labels in their Office apps. Allow up to 24 hours for the latest changes to replicate throughout your organization.

There is no limit to the number of sensitivity labels that you can create and publish, with one exception: If the label applies encryption, there is a maximum of 500 labels that you can create. However, as a best practice to lower admin overheads and reduce complexity for your users, try to keep the number of labels to a minimum. Real-world deployments have proved effectiveness to be noticeably reduced when users have more than five main labels or more than five sublabels per main label.

Label policy priority (order matters)

You make your sensitivity labels available to users by publishing them in a sensitivity label policy that appears in a list on the **Sensitivity policies** tab on the **Label policies** page. Just like sensitivity labels (see Label priority (order matters)), the order of the sensitivity label policies is important because it reflects their priority. The label policy with lowest priority is shown at the **top**, and the label policy with the highest priority is shown at the **bottom**.

A label policy consists of:

- A set of labels.
- The users and groups that will be assigned the policy with labels.
- The scope of the policy and policy settings for that scope (such as default label for files and emails).

You can include a user in multiple label policies, and the user will see all the sensitivity labels from those policies. However, a user gets the policy settings from only the label policy with the highest priority.

If you're not seeing the label or label policy setting that you expect for a user or group, check the order of the sensitivity label policies. To reorder the label policies, select a sensitivity label policy > choose the ellipsis on the right > **Move down** or **Move up**.

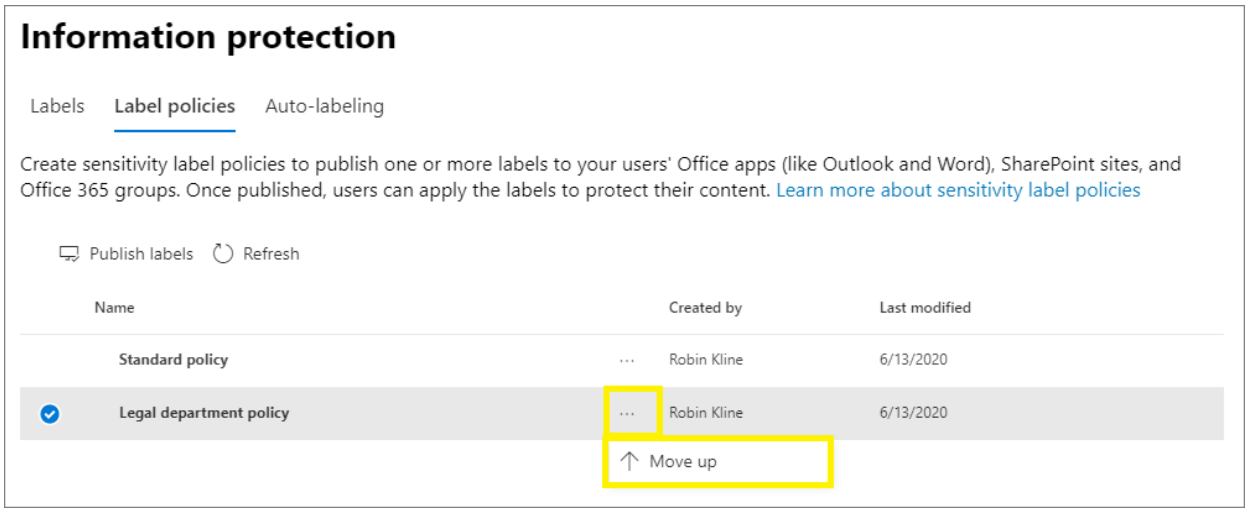

If you use retention labels in addition to sensitivity labels, it's important to remember that priority matters for sensitivity label policies, but not for retention labels.

Sensitivity labels and Azure Information Protection

If you have deployed labels with Azure Information Protection, use the following sections for guidance before you start to use sensitivity labels.

Azure Information Protection labels

> **NOTE**
> Label management for Azure Information Protection labels in the Azure portal is being deprecated **March 31, 2021**. Learn more from the official deprecation notice.

If you are using Azure Information Protection labels because your tenant isn't yet on the unified labeling platform, we recommend that you avoid creating sensitivity labels until you activate unified labeling. In this scenario, the labels you see in the Azure portal are Azure Information Protection labels rather than sensitivity labels. These labels can be used by the Azure Information Protection client (classic) on Windows computers, but can't be used by devices running macOS, iOS, or Android. To resolve this, migrate these labels to sensitivity labels.

The metadata applied by both sets of labels are compatible, so you don't need to relabel documents and emails when the migration is complete.

Azure Information Protection clients

When you use sensitivity labels in Microsoft 365 Apps for enterprise apps on Windows computers, you have a choice of using an Azure Information Protection client, or use labeling that's built into Office.

By default, built-in labeling is turned off in these apps when the Azure Information Protection client is installed. For more information, including how to change this default behavior, see Office built-in labeling client and the Azure Information Protection client.

Even when you use built-in labeling in Office apps, you can also use the Azure Information Protection unified labeling client with sensitivity labels for the following:

- A scanner to discover sensitive information that's stored on-premises, and then optionally, label that content
- Right-click options in File Explorer for users to apply labels to all file types
- A viewer to display encrypted files for text, images, or PDF documents

- A PowerShell module to discover sensitive information in files on premises, and apply or remove labels and encryption from these files.

If you are new to Azure Information Protection, or if you are an existing Azure Information Protection customer that has just migrated your labels, see Choose which labeling client to use for Windows computers from the Azure Information Protection documentation.

Sensitivity labels and Microsoft Cloud App Security

By using Cloud App Security (CAS), you can discover, classify, label, and protect content in third-party services and apps, such as SalesForce, Box, or Dropbox.

Cloud App Security works with both Azure Information Protection labels and sensitivity labels:

- If the labeling admin centers have one or more sensitivity labels published to at least one user: Sensitivity labels are used.
- If the labeling admin centers don't have sensitivity labels published: Azure Information Protection labels are used.

For instructions to use Cloud App Security with these labels, see Azure Information Protection integration.

Sensitivity labels and the Microsoft Information Protection SDK

Because a sensitivity label is stored as clear text in the metadata of a document, third-party apps and services can read from and write to this labeling metadata to supplement your labeling deployment. Additionally, software developers can use the Microsoft Information Protection SDK to fully support labeling and encryption capabilities across multiple platforms. To learn more, see the General Availability announcement on the Tech Community blog.

You can also learn about partner solutions that are integrated with Microsoft Information Protection.

Deployment guidance

For deployment planning and guidance that includes licensing information, permissions, deployment strategy, a list of supported scenarios, and end-user documentation, see Get started with sensitivity labels.

Get started with sensitivity labels

2/18/2021 • 6 minutes to read • Edit Online

Microsoft 365 licensing guidance for security & compliance.

For information about what sensitivity labels are and how they can help you protect your organization's data, see Learn about sensitivity labels.

If you have Azure Information Protection, determine whether you need to migrate labels to the unified labeling platform, and which labeling client to use:

- How can I determine if my tenant is on the unified labeling platform?
- Choose which labeling client to use for Windows computers

When you're ready to start protecting your organization's data by using sensitivity labels:

1. **Create the labels.** Create and name your sensitivity labels according to your organization's classification taxonomy for different sensitivity levels of content. Use common names or terms that make sense to your users. If you don't already have an established taxonomy, consider starting with label names such as Personal, Public, General, Confidential, and Highly Confidential. You can then use sublabels to group similar labels by category. When you create a label, use the tooltip text to help users select the appropriate label.

   For more extensive guidance for defining a classification taxonomy, download the white paper, "Data Classification & Sensitivity Label Taxonomy" from the Service Trust Portal.

2. **Define what each label can do.** Configure the protection settings you want associated with each label. For example, you might want lower sensitivity content (such as a "General" label) to have just a header or footer applied, while higher sensitivity content (such as a "Confidential" label) should have a watermark and encryption.

3. **Publish the labels.** After your sensitivity labels are configured, publish them by using a label policy. Decide which users and groups should have the labels and what policy settings to use. A single label is reusable—you define it once, and then you can include it in several label policies assigned to different users. So for example, you could pilot your sensitivity labels by assigning a label policy to just a few users. Then when you're ready to roll out the labels across your organization, you can create a new label policy for your labels and this time, specify all users.

The basic flow for deploying and applying sensitivity labels:

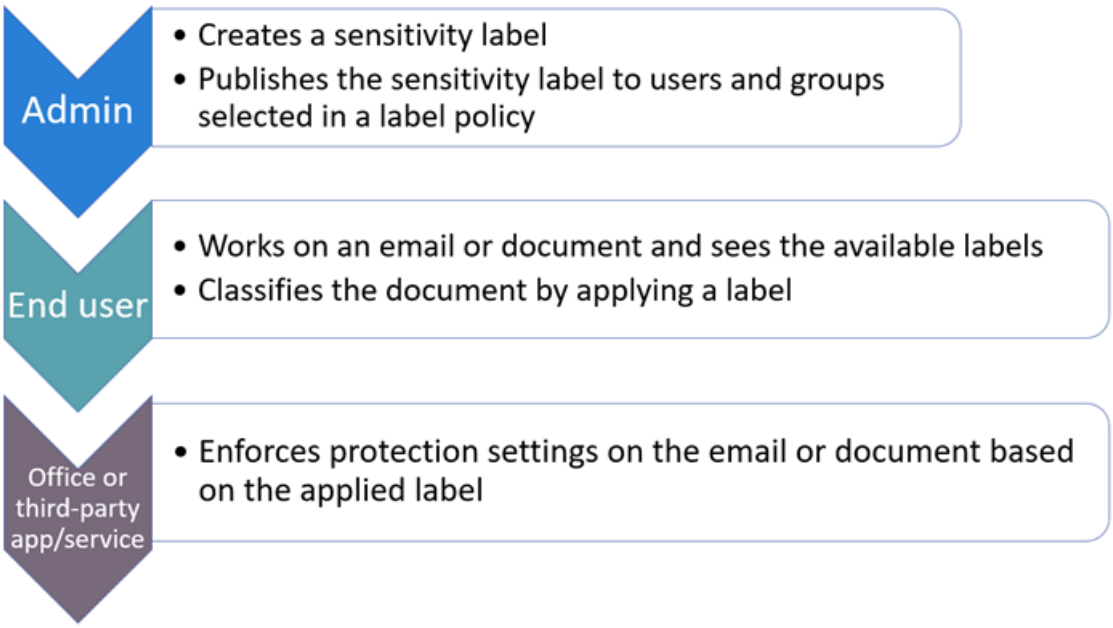

Subscription and licensing requirements for sensitivity labels

A number of different subscriptions support sensitivity labels and the licensing requirements for users depend on the features you use.

To see the options for licensing your users to benefit from Microsoft 365 compliance features, see the Microsoft 365 licensing guidance for security & compliance. For sensitivity labels, see the Information Protection section and related PDF or Excel download.

Permissions required to create and manage sensitivity labels

Members of your compliance team who will create sensitivity labels need permissions to the Microsoft 365 compliance center, Microsoft 365 security center, or the Security & Compliance Center.

By default, global administrators for your tenant have access to these admin centers and can give compliance officers and other people access, without giving them all of the permissions of a tenant admin. For this delegated limited admin access, add users to the **Compliance Data Administrator**, **Compliance Administrator**, or **Security Administrator** role group.

Alternatively to using the default roles, you can create a new role group and add either **Sensitivity Label Administrator** or **Organization Configuration** roles to this group. For a read-only role, use **Sensitivity Label Reader**.

For instructions to add users to the default roles or create your own role groups, see Give users access to the Office 365 Security & Compliance Center.

These permissions are required only to create and configure sensitivity labels and their label policies. They are not required to apply the labels in apps or services. If additional permissions are needed for specific configurations that relate to sensitivity labels, those permissions will be listed in their respective documentation instructions.

Deployment strategy for sensitivity labels

A successful strategy to deploy sensitivity labels for an organization is to create a working virtual team that identifies and manages the business and technical requirements, proof of concept testing, internal checkpoints and approvals, and final deployment for the production environment.

Using the table in the next section, we recommend identifying your top one or two scenarios that map to your

most impactful business requirements. After these scenarios are deployed, return to the list to identify the next one or two priorities for deployment.

You'll find additional general deployment guidance in the downloadable Data Loss Prevention and Microsoft Information Protection Deployment Acceleration Guide. For more information, see the blog post, Microsoft 365 Information Protection and Compliance Deployment Acceleration Guides.

Common scenarios for sensitivity labels

All scenarios require you to Create and configure sensitivity labels and their policies.

I WANT TO ...	DOCUMENTATION
Manage sensitivity labels for Office apps so that content is labeled as it's created—includes support for manual labeling on all platforms	Manage sensitivity labels in Office apps
Enable users to label and protect files from Windows computers using Office apps, File Explorer, and PowerShell	Azure Information Protection unified labeling client for Windows
Encrypt documents and emails with sensitivity labels and restrict who can access that content and how it can be used	Restrict access to content by using sensitivity labels to apply encryption
Enable sensitivity labels for Office on the web, with support for coauthoring, eDiscovery, data loss prevention, search—even when documents are encrypted	Enable sensitivity labels for Office files in SharePoint and OneDrive
Automatically apply sensitivity labels to documents and emails	Apply a sensitivity label to content automatically
Use sensitivity labels to protect content in Teams and SharePoint	Use sensitivity labels with Microsoft Teams, Microsoft 365 groups, and SharePoint sites
Prevent or warn users about sharing files or emails with a specific sensitivity label	Use sensitivity labels as conditions in DLP policies (preview)
Discover, label, and protect files stored in data stores that are on premises	Deploying the Azure Information Protection scanner to automatically classify and protect files
Discover, label, and protect files stored in data stores that are in the cloud	Discover, classify, label, and protect regulated and sensitive data stored in the cloud
Apply and view labels in Power BI, and protect data when it's saved outside the service	Sensitivity labels in Power BI
Monitor and understand how sensitivity labels are being used in my organization	Know your data - data classification overview Get started with data classification
Extend sensitivity labels to third-party apps and services	Microsoft Information Protection SDK
Extend sensitivity labels across content in Azure Blob Storage, Azure files, Azure Data Lake Storage Gen1, and Azure Data Lake Storage Gen12	Automatically label your content in Azure Purview

End-user documentation for sensitivity labels

The most effective end-user documentation will be customized guidance and instructions you provide for the label names and configurations you choose. For built-in labeling, you can use the label policy setting **Provide users with a link to a custom help page** to specify an internal link for this documentation. Users can then easily access it by selecting **Learn More** from the **Sensitivity** button on the Office ribbon for Word, PowerPoint, Excel, and Outlook.

To help you write your customized documentation, see the following blog post for a download package that you can use to train users and drive adoption: End User Training for Sensitivity Labels in M365 – How to Accelerate Your Adoption.

You can also use the following resources for basic instructions:

- Apply sensitivity labels to your files and email in Office
    - Known issues with sensitivity labels in Office
- Automatically apply or recommend sensitivity labels to your files and emails in Office
    - Known issues with automatically applying or recommending sensitivity labels
- Azure Information Protection unified labeling user guide

If your sensitivity labels apply encryption for PDF documents, these documents can be opened with Microsoft Edge on Windows or Mac. For more information, and alternative readers, see Which PDF readers are supported for protected PDFs?

Create and configure sensitivity labels and their policies

2/18/2021 • 11 minutes to read • Edit Online

> *Microsoft 365 licensing guidance for security & compliance*.

All Microsoft Information Protection solutions (sometimes abbreviated to MIP) are implemented by using sensitivity labels. To create and publish these labels, go to your labeling admin center, such as the Microsoft 365 compliance center. You can also use the Microsoft 365 security center, or the Security & Compliance Center.

First, create and configure the sensitivity labels that you want to make available for apps and other services. For example, the labels you want users to see and apply from Office apps.

Then, create one or more label policies that contain the labels and policy settings that you configure. It's the label policy that publishes the labels and settings for your chosen users and locations.

Before you begin

The global admin for your organization has full permissions to create and manage all aspects of sensitivity labels. If you aren't signing in as a global admin, see Permissions required to create and manage sensitivity labels.

Create and configure sensitivity labels

1. In your labeling admin center, navigate to sensitivity labels:
   - Microsoft 365 compliance center:
     - **Solutions** > **Information protection**

     If you don't immediately see this option, first select **Show all**.
   - Microsoft 365 security center:
     - **Classification** > **Sensitivity labels**
   - Security & Compliance Center:
     - **Classification** > **Sensitivity labels**
2. On the **Labels** page, select **+ Create a label** to start the New sensitivity label wizard.

   For example, from the Microsoft 365 compliance center:

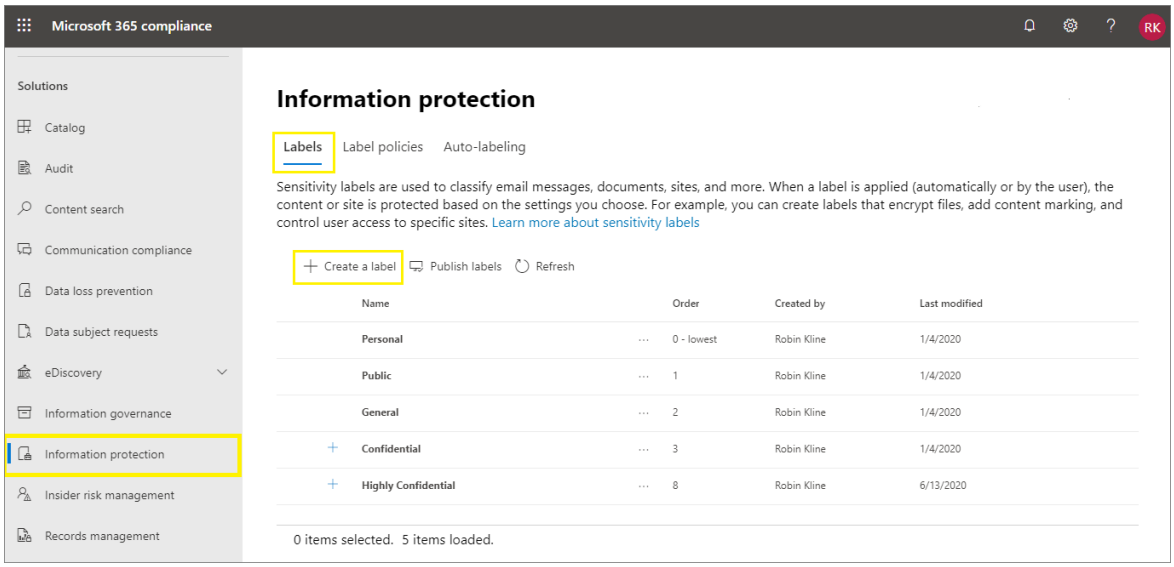

> **NOTE**
>
> By default, tenants don't have any labels and you must create them. The labels in the example picture show default labels that were migrated from Azure Information Protection.

3. On the **Define the scope for this label** page, the options selected determine the label's scope for the settings that you can configure and where they will be visible when they are published:

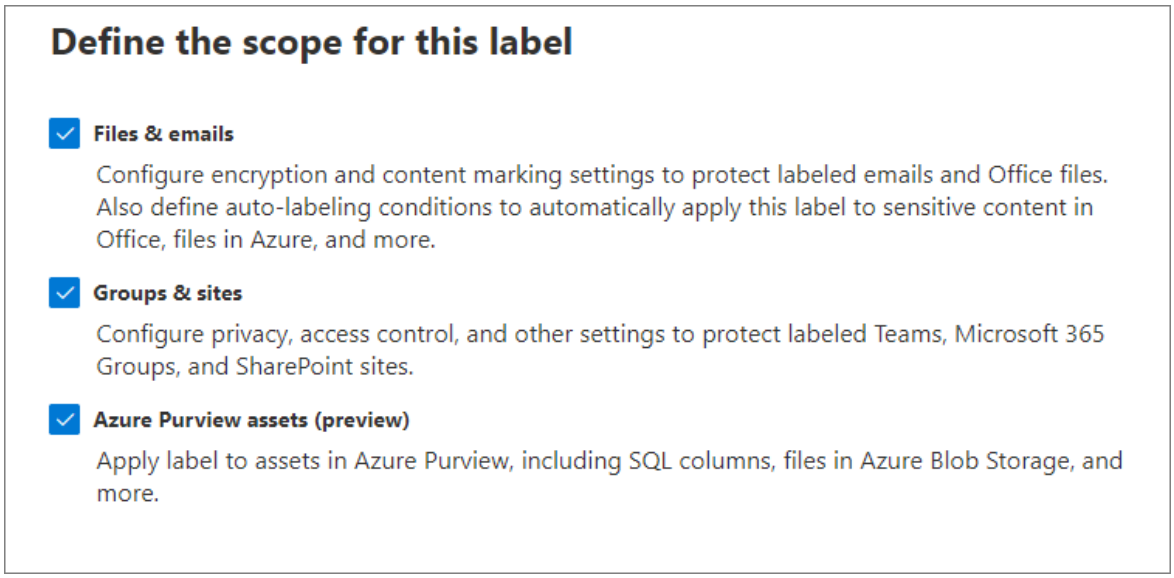

   - If **Files & emails** is selected, you can configure settings in this wizard that apply to apps that support sensitivity labels, such as Office Word and Outlook. If this option isn't selected, the wizard displays the first page of these settings but you can't configure them and the labels won't be available for users to select in these apps.
   - If **Groups & sites** is selected, you can configure settings in this wizard that apply to Microsoft 365 groups, and sites for Teams and SharePoint. If this option isn't selected, the wizard displays the first page of these settings but you can't configure them and the labels won't be available for users to select for groups and site.

   For information about the **Azure Purview assets (preview)** scope, see Automatically label your content in Azure Purview.

4. Follow the prompts in the wizard for the label settings.

   For more information about the label settings, see What sensitivity labels can do from the overview

information and use the help in the wizard for individual settings.

5. Repeat these steps to create more labels. However, if you want to create a sublabel, first select the parent label and select ... for **More actions**, and then select **Add sub label**.

6. When you have created all the labels you need, review their order and if necessary, move them up or down. To change the order of a label, select ... for **More actions**, and then select **Move up** or **Move down**. For more information, see Label priority (order matters) from the overview information.

To edit an existing label, select it, and then select the **Edit label** button:

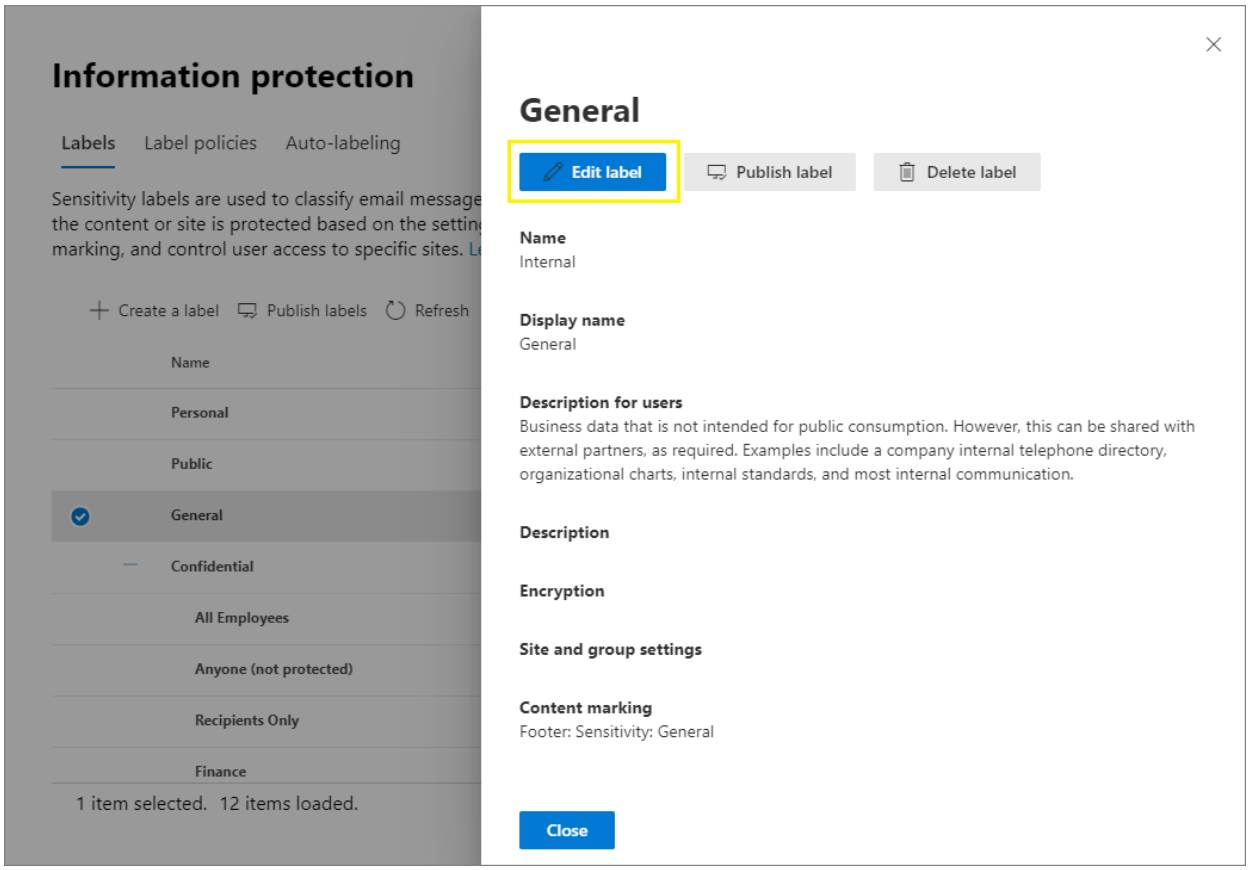

This button starts the **Edit sensitivity label** wizard, which lets you change all the label settings in step 4.

Don't delete a label unless you understand the impact for users. For more information, see the Removing and deleting labels section.

> **NOTE**
>
> If you edit a label that's already published by using a label policy, no extra steps are needed when you finish the wizard. For example, you don't need to add it to a new label policy for the changes to become available to the same users. However, allow up to 24 hours for the changes to replicate to all apps and services.

Until you publish your labels, they won't be available to select in apps or for services. To publish the labels, they must be added to a label policy.

> **IMPORTANT**
>
> On this **Labels** tab, do not select the **Publish labels** tab (or the **Publish label** button when you edit a label) unless you need to create a new label policy. You need multiple label policies only if users need different labels or different policy settings. Aim to have as few label policies as possible—it's not uncommon to have just one label policy for the organization.

Additional label settings with Security & Compliance Center PowerShell

Additional label settings are available with the Set-Label cmdlet from Security & Compliance Center PowerShell.

For example:

- Use the *LocaleSettings* parameter for multinational deployments so that users see the label name and tooltip in their local language. The following section has an example configuration that specifies the label name and tooltip text for French, Italian, and German.
- For the Azure Information Protection unified labeling client only, specify advanced settings that include setting a label color, and applying a custom property when a label is applied. For the full list, see Available advanced settings for labels from this client's admin guide.

Example configuration to configure a sensitivity label for different languages

The following example shows the PowerShell configuration for a label named "Public" with placeholder text for the tooltip. In this example, the label name and tooltip text are configured for French, Italian, and German.

As a result of this configuration, users who have Office apps that use those display languages see their label names and tooltips in the same language. Similarly, if you have the Azure Information Protection unified labeling client installed to label files from File Explorer, users who have those language versions of Windows see their label names and tooltips in their local language when they use the right-click actions for labeling.

For the languages that you need to support, use the Office language identifiers (also known as language tags), and specify your own translation for the label name and tooltip.

Before you run the commands in PowerShell, you must first connect to Security & Compliance Center PowerShell.

```
$Languages = @("fr-fr","it-it","de-de")
$DisplayNames=@("Publique","Publico","Oeffentlich")
$Tooltips = @("Texte Français","Testo italiano","Deutscher text")
$label = "Public"
$DisplayNameLocaleSettings = [PSCustomObject]@{LocaleKey='DisplayName';
Settings=@(
@{key=$Languages[0];Value=$DisplayNames[0];}
@{key=$Languages[1];Value=$DisplayNames[1];}
@{key=$Languages[2];Value=$DisplayNames[2];})}
$TooltipLocaleSettings = [PSCustomObject]@{LocaleKey='Tooltip';
Settings=@(
@{key=$Languages[0];Value=$Tooltips[0];}
@{key=$Languages[1];Value=$Tooltips[1];}
@{key=$Languages[2];Value=$Tooltips[2];})}
Set-Label -Identity $Label -LocaleSettings (ConvertTo-Json $DisplayNameLocaleSettings -Depth 3 -Compress),
(ConvertTo-Json $TooltipLocaleSettings -Depth 3 -Compress)
```

Publish sensitivity labels by creating a label policy

1. In your labeling admin center, navigate to sensitivity labels:
   - Microsoft 365 compliance center:
     - **Solutions** > **Information protection**

     If you don't immediately see this option, first select **Show all**.
   - Microsoft 365 security center:
     - **Classification** > **Sensitivity labels**
   - Security & Compliance Center:
     - **Classification** > **Sensitivity labels**

2. Select the **Label policies** tab, and then **Publish labels** to start the Create policy wizard:

   For example, from the Microsoft 365 compliance center:

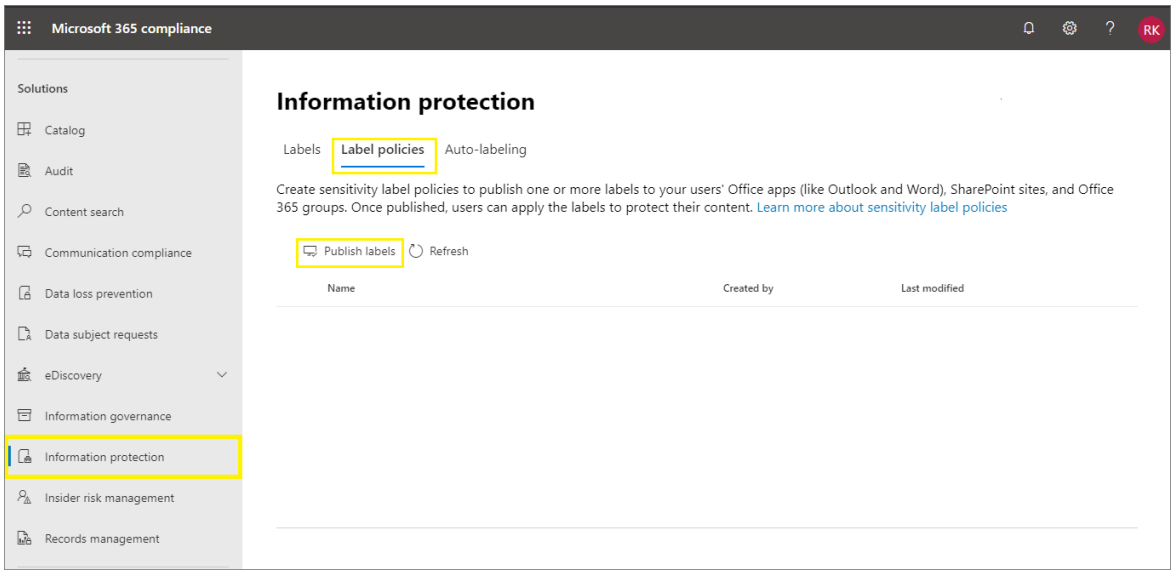

> **NOTE**
>
> By default, tenants don't have any label policies and you must create them.

3. In the wizard, select **Choose sensitivity labels to publish**. Select the labels that you want to make available in apps and to services, and then select **Add**.

> **IMPORTANT**
>
> If you select a sublabel, make sure you also select its parent label.

4. Review the selected labels and to make any changes, select **Edit**. Otherwise, select **Next**.

5. Follow the prompts to configure the policy settings.

   The policy settings that you see match the scope of the labels that you selected. For example, if you selected labels that have just the **Files & emails** scope, you don't see the policy settings **Apply this label by default to groups and sites** and **Require users to apply a label to their groups and sites**.

   For more information about these settings, see What label policies can do from the overview information and use the help in the wizard for individual settings.

   For labels configured for **Azure Purview assets (preview)**: These labels don't have any associated policy settings.

6. Repeat these steps if you need different policy settings for different users or scopes. For example, you want additional labels for a group of users, or a different default label for a subset of users. Or, if you have configured labels to have different scopes.

7. If you create more than one label policy that might result in a conflict for a user, review the policy order and if necessary, move them up or down. To change the order of a label policy, select **...** for **More actions**, and then select **Move up** or **Move down**. For more information, see Label policy priority (order matters) from the overview information.

Completing the wizard automatically publishes the label policy. To make changes to a published policy, simply

edit it. There is no specific publish or republish action for you to select.

To edit an existing label policy, select it, and then select the **Edit Policy** button:

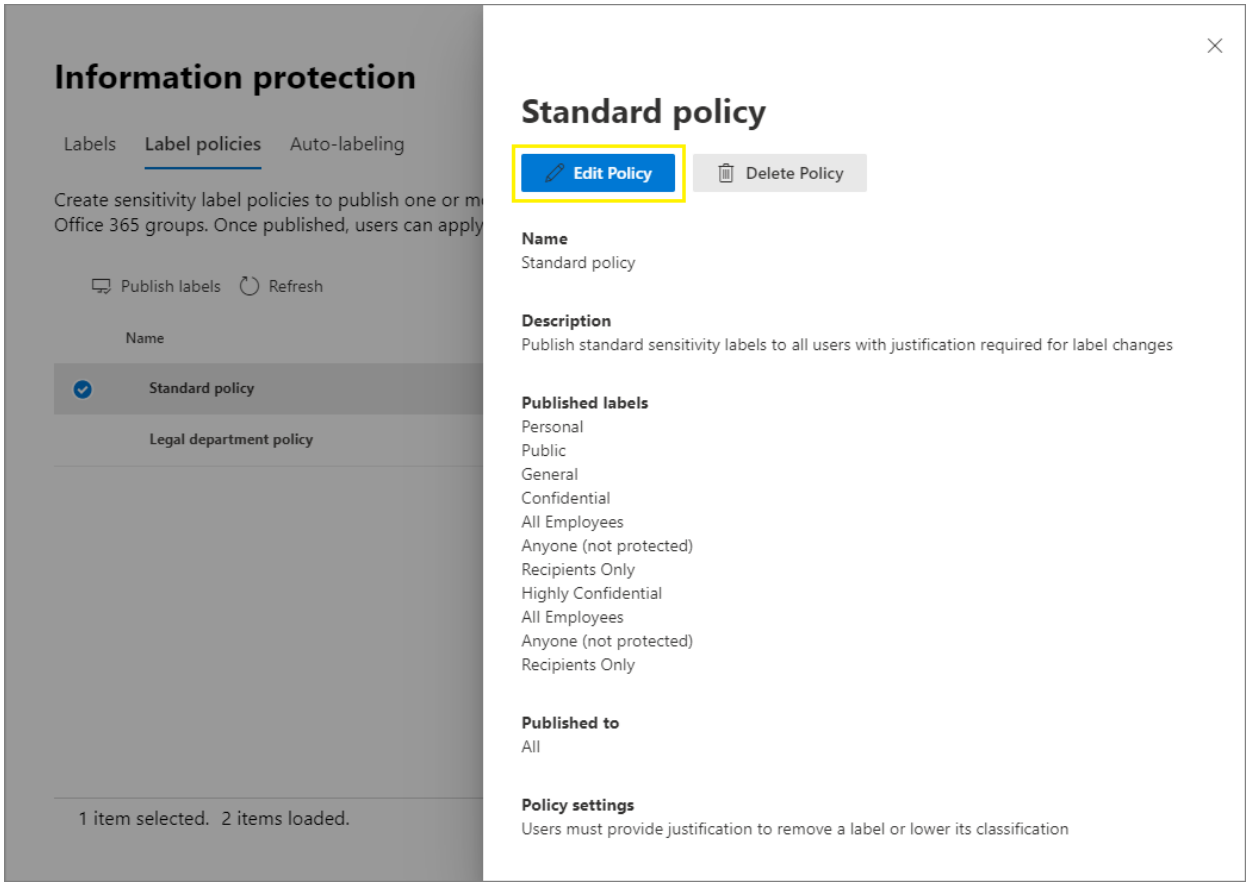

This button starts the **Create policy** wizard, which lets you edit which labels are included and the label settings. When you complete the wizard, any changes are automatically replicated to the selected users and services.

When you use built-in labeling for Office apps on Windows, macOS, iOS, and Android, users see new labels within four hours, and within one hour for Office on the web. However, allow up to 24 hours for changes to replicate to all apps and services.

Additional label policy settings with Security & Compliance Center PowerShell

Additional label policy settings are available with the Set-LabelPolicy cmdlet from Security & Compliance Center PowerShell.

For the Azure Information Protection unified labeling client only, you can specify advanced settings that include setting a different default label for Outlook, and implement pop-up messages in Outlook that warn, justify, or block emails being sent. For the full list, see Available advanced settings for label policies from this client's admin guide.

Use PowerShell for sensitivity labels and their policies

You can now use Security & Compliance Center PowerShell to create and configure all the settings you see in your labeling admin center. This means that in addition to using PowerShell for settings that aren't available in the labeling admin centers, you can now fully script the creation and maintenance of sensitivity labels and sensitivity label policies.

See the following documentation for supported parameters and values:

- New-Label
- New-LabelPolicy
- Set-Label

- Set-LabelPolicy

You can also use Remove-Label and Remove-LabelPolicy if you need to script the deletion of sensitivity labels or sensitivity label policies. However, before you delete sensitivity labels, make sure you read the following section.

Removing and deleting labels

In a production environment, it's unlikely that you will need to remove sensitivity labels from a label policy, or delete sensitivity labels. It's more likely that you might need to do one or either of these actions during an initial testing phase. Make sure you understand what happens when you do either of these actions.

Removing a label from a label policy is less risky than deleting it, and you can always add it back to a label policy later if needed:

- When you remove a label from a label policy so that the label is no longer published to the originally specified users, the next time the label policy is refreshed, users no longer see that label to select in their Office app. However, if the label has been applied to documents or emails, the label isn't removed from that content. Any encryption that was applied by the label remains and the underlying protection template remains published.
- For labels that are removed but have previously been applied to content, users who are using built-in labeling for Word, Excel, and PowerPoint, still see the applied label name on the status bar. Similarly, labels that are removed that were applied to SharePoint sites still display the label name in the **Sensitivity** column.

In comparison, when you delete a label:

- If the label applied encryption, the underlying protection template is archived so that previously protected content can still be opened. Because of this archived protection template, you won't be able to create a new label with the same name. Although it's possible to delete a protection template by using PowerShell, don't do this unless you're sure you don't need to open content that was encrypted with the archived template.
- For desktop apps: The label information in the metadata remains, but because a label ID to name mapping is no longer possible, users don't see the applied label name displayed (for example, on the status bar) so users will assume the content isn't labeled. If the label applied encryption, the encryption remains and when the content is opened, uses still see the name and description of the now archived protection template.
- For Office on the web: Users don't see the label name on status bar or in the **Sensitivity** column. The label information in the metadata remains only if the label didn't apply encryption. If the label applied encryption, and you've enabled sensitivity labels for SharePoint and OneDrive, the label information in the metadata is removed and the encryption is removed.

When you remove a sensitivity label from a label policy, or delete a sensitivity label, these changes can take up to one hour to replicate to all users and services.

Next steps

To configure and use your sensitivity labels for specific scenarios, use the following articles:

- Restrict access to content by using encryption in sensitivity labels
- Apply a sensitivity label to content automatically
- Use sensitivity labels with teams, groups, and sites
- Enable sensitivity labels for Office files in SharePoint and OneDrive

To monitor how your labels are being used, see Get started with data classification.

Restrict access to content by using sensitivity labels to apply encryption

2/18/2021 • 21 minutes to read • Edit Online

Microsoft 365 licensing guidance for security & compliance.

When you create a sensitivity label, you can restrict access to content that the label will be applied to. For example, with the encryption settings for a sensitivity label, you can protect content so that:

- Only users within your organization can open a confidential document or email.
- Only users in the marketing department can edit and print the promotion announcement document or email, while all other users in your organization can only read it.
- Users cannot forward an email or copy information from it that contains news about an internal reorganization.
- The current price list that is sent to business partners cannot be opened after a specified date.

When a document or email is encrypted, access to the content is restricted, so that it:

- Can be decrypted only by users authorized by the label's encryption settings.
- Remains encrypted no matter where it resides, inside or outside your organization, even if the file's renamed.
- Is encrypted both at rest (for example, in a OneDrive account) and in transit (for example, email as it traverses the internet).

Finally, as an admin, when you configure a sensitivity label to apply encryption, you can choose either to:

- **Assign permissions now**, so that you determine exactly which users get which permissions to content with that label.
- **Let users assign permissions** when they apply the label to content. This way, you can allow people in your organization some flexibility that they might need to collaborate and get their work done.

The encryption settings are available when you create a sensitivity label in the Microsoft 365 compliance center, Microsoft 365 security center, or the Security & Compliance Center.

Understand how the encryption works

Encryption uses the Azure Rights Management service (Azure RMS) from Azure Information Protection. This protection solution uses encryption, identity, and authorization policies. To learn more, see What is Azure Rights Management? from the Azure Information Protection documentation.

When you use this encryption solution, the **super user** feature ensures that authorized people and services can always read and inspect the data that has been encrypted for your organization. If necessary, the encryption can then be removed or changed. For more information, see Configuring super users for Azure Information Protection and discovery services or data recovery.

How to configure a label for encryption

1. Follow the general instructions to create or edit a sensitivity label and make sure **Files & emails** is selected for the label's scope:

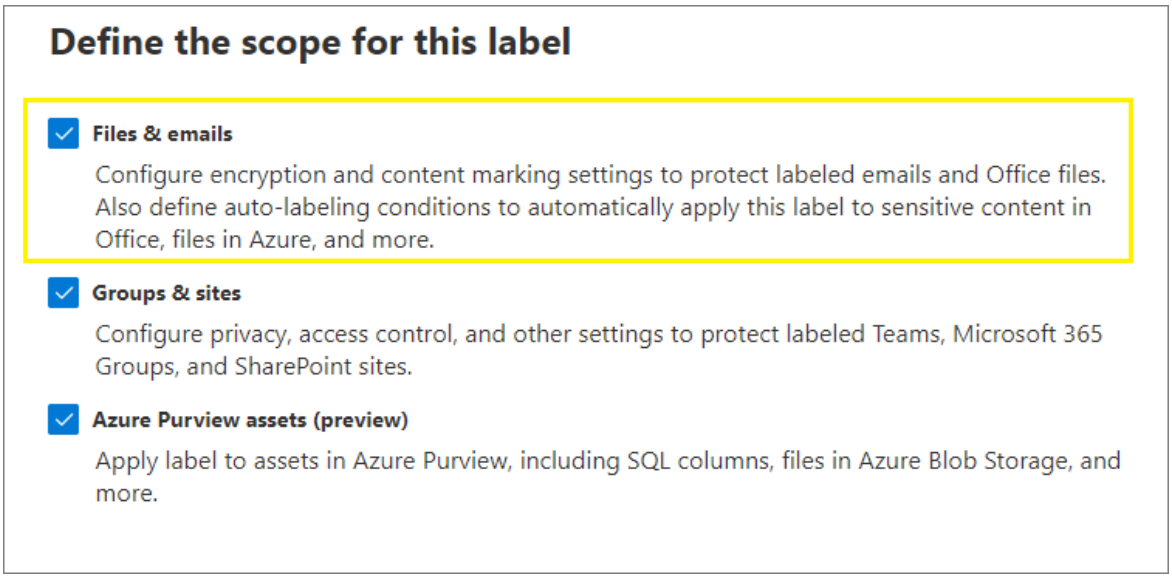

2. Then, on the **Choose protection settings for files and emails** page, make sure you select **Encrypt files and emails**

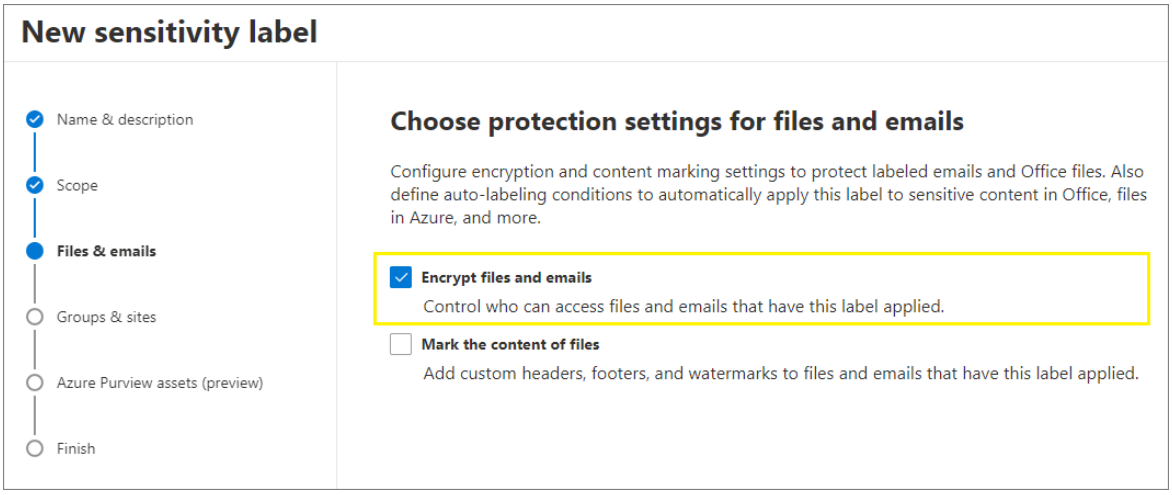

3. On the **Encryption** page of the wizard, select one of the following options:

   - **Remove encryption if the file is encrypted**: For more information about this scenario, see the What happens to existing encryption when a label's applied section. It's important to understand that this setting can result in a sensitivity label that users might not be able to apply when they don't have sufficient permissions.
   - **Configure encryption settings**: Turns on encryption and makes the encryption settings visible:

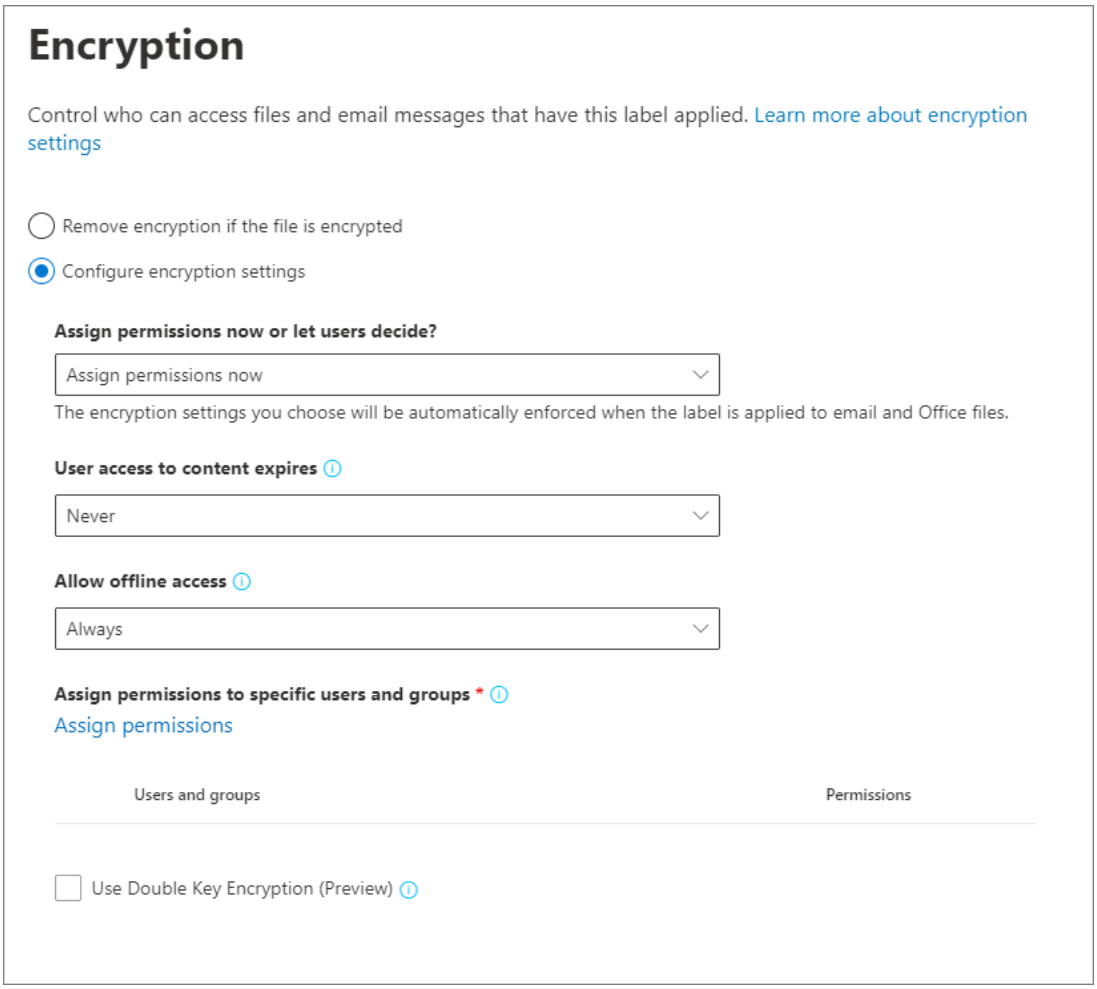

Instructions for these settings are in the following Configure encryption settings section.

What happens to existing encryption when a label's applied

If a sensitivity label is applied to unencrypted content, the outcome of the encryption options you can select is self-explanatory. For example, if you didn't select **Encrypt files and emails**, the content remains unencrypted.

However, the content might be already encrypted. For example, another user might have applied:

- Their own permissions, which include user-defined permissions when prompted by a label, custom permissions by the Azure Information Protection client, and the **Restricted Access** document protection from within an Office app.
- An Azure Rights Management protection template that encrypts the content independently from a label. This category includes mail flow rules that apply encryption by using rights protection.
- A label that applies encryption with permissions assigned by the administrator.

The following table identifies what happens to existing encryption when a sensitivity label is applied to that content:

	ENCRYPTION: NOT SELECTED	ENCRYPTION: CONFIGURED	ENCRYPTION: REMOVE
Permissions specified by a user	Original encryption is preserved	New label encryption is applied	Original encryption is removed
Protection template	Original encryption is preserved	New label encryption is applied	Original encryption is removed
Label with administator-defined permissions	Original encryption is removed	New label encryption is applied	Original encryption is removed

Note that in the cases where the new label encryption is applied or the original encryption is removed, this happens only if the user applying the label has a usage right or role that supports this action:

- The usage right Export or Full Control.
- The role of Rights Management issuer or Rights Management owner, or super user.

If the user doesn't have one of these rights or roles, the label can't be applied and so the original encryption is preserved. The user sees the following message: **You don't have permission to make this change to the sensitivity label. Please contact the content owner.**

For example, the person who applies Do Not Forward to an email message can relabel the thread to replace the encryption or remove it, because they are the Rights Management owner for the email. But with the exception of super users, recipients of this email can't relabel it because they don't have the required usage rights.

Email attachments for encrypted email messages

When an email message is encrypted by any method, any unencrypted Office documents that are attached to the email automatically inherit the same encryption settings.

Documents that are already encrypted and then added as attachments always preserve their original encryption.

Configure encryption settings

When you select **Configure encryption settings** on the **Encryption** page of the wizard to create or edit a sensitivity label, choose one of the following options:

- **Assign permissions now**, so that you can determine exactly which users get which permissions to content that has the label applied. For more information, see the next section Assign permissions now.
- **Let users assign permissions** when your users apply the label to content. With this option, you can allow people in your organization some flexibility that they might need to collaborate and get their work done. For more information, see the Let users assign permissions section on this page.

For example, if you have a sensitivity label named **Highly Confidential** that will be applied to your most sensitive content, you might want to decide now who gets what type of permissions to that content.

Alternatively, if you have a sensitivity label named **Business Contracts**, and your organization's workflow requires that your people collaborate on this content with different people on an ad hoc basis, you might want to allow your users to decide who gets permissions when they assign the label. This flexibility both helps your users' productivity and reduces the requests for your admins to update or create new sensitivity labels to address specific scenarios.

Choosing whether to assign permissions now or let users assign permissions:

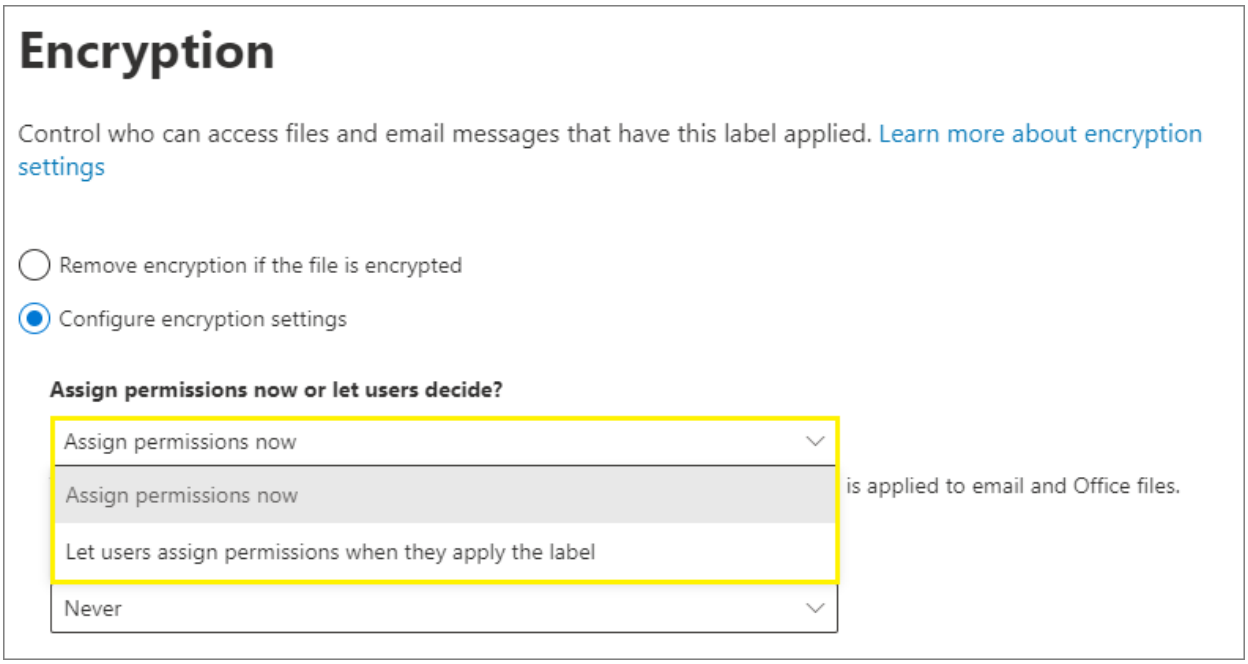

Assign permissions now

Use the following options to control who can access email or documents to which this label is applied. You can:

- **Allow access to labeled content to expire**, either on a specific date or after a specific number of days after the label is applied. After this time, users won't be able to open the labeled item. If you specify a date, it is effective midnight on that date in your current time zone. (Note that some email clients might not enforce expiration and show emails past their expiration date, due to their caching mechanisms.)
- **Allow offline access** never, always, or for a specific number of days after the label is applied. If you restrict offline access to never or a number of days, when that threshold is reached, users must be reauthenticated and their access is logged. For more information, see the next section on the Rights Management use license.

Settings for access control for encrypted content:

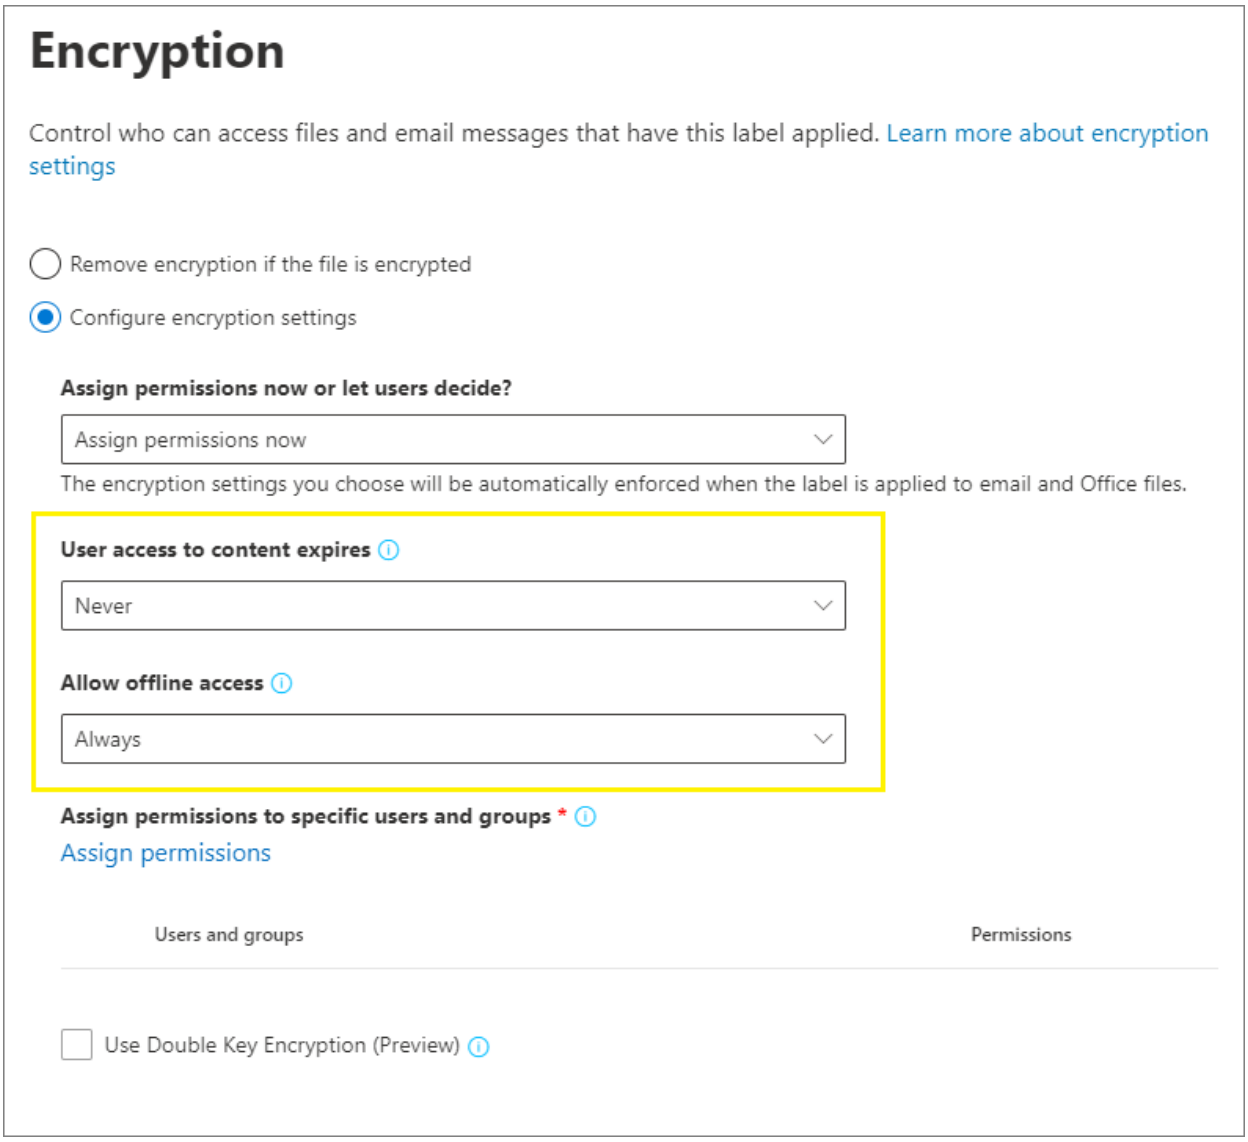

Rights Management use license for offline access

When a user opens a document or email that's been protected by encryption from the Azure Rights Management service, an Azure Rights Management use license for that content is granted to the user. This use license is a certificate that contains the user's usage rights for the document or email, and the encryption key that was used to encrypt the content. The use license also contains an expiration date if this has been set, and how long the use license is valid.

If no expiration date has been set, the default use license validity period for a tenant is 30 days. For the duration of the use license, the user is not reauthenticated or reauthorized for the content. This process lets the user continue to open the protected document or email without an internet connection. When the use license validity period expires, the next time the user accesses the protected document or email, the user must be reauthenticated and reauthorized.

In addition to reauthentication, the encryption settings and user group membership is reevaluated. This means that users could experience different access results for the same document or email if there are changes in the encryption settings or group membership from when they last accessed the content.

To learn how to change the default 30-day setting, see Rights Management use license.

Assign permissions to specific users or groups

You can grant permissions to specific people so that only they can interact with the labeled content:

1. First, add users or groups that will be assigned permissions to the labeled content.
2. Then, choose which permissions those users should have for the labeled content.

Assigning permissions:

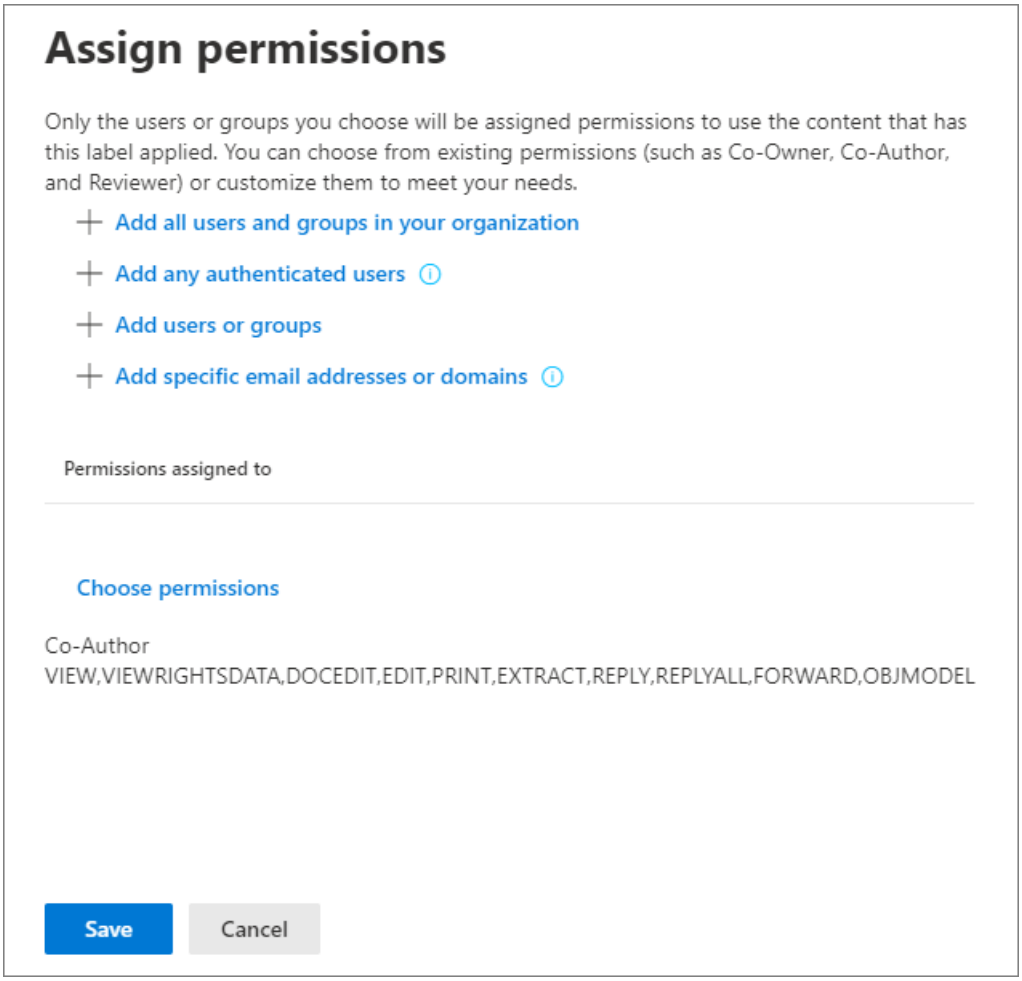

Add users or groups

When you assign permissions, you can choose:

- Everyone in your organization (all tenant members). This setting excludes guest accounts.
- Any authenticated users. Make sure you understand the requirements and limitations of this setting before selecting it.
- Any specific user or email-enabled security group, distribution group, or Microsoft 365 group (formerly Office 365 group) in Azure AD. The Microsoft 365 group can have static or dynamic membership. Note that you can't use a dynamic distribution group from Exchange because this group type isn't synchronized to Azure AD, and you can't use a security group that isn't email-enabled.
- Any email address or domain. Use this option to specify all users in another organization who uses Azure AD, by entering any domain name from that organization. You can also use this option for social providers, by entering their domain name such as **gmail.com**, **hotmail.com**, or **outlook.com**.

  > **NOTE**
  > If you specify a domain from an organization that uses Azure AD, you can't restrict access to that specific domain. Instead, all verified domains in Azure AD are automatically included for the tenant that owns the domain name you specify.

When you choose all users and groups in your organization or browse the directory, the users or groups must have an email address.

As a best practice, use groups rather than users. This strategy keeps your configuration simpler.

Requirements and limitations for "Add any authenticated users"

This setting doesn't restrict who can access the content that the label encrypts, while still encrypting the content and providing you with options to restrict how the content can be used (permissions), and accessed (expiry and offline access). However, the application opening the encrypted content must be able to support the authentication being used. For this reason, federated social providers such as Google, and onetime passcode authentication work for email only, and only when you use Exchange Online. Microsoft accounts can be used with Office 365 apps and the Azure Information Protection viewer.

> **NOTE**
>
> Consider using this setting with SharePoint and OneDrive integration with Azure AD B2B when sensitivity labels are enabled for Office files in SharePoint and OneDrive.

Some typical scenarios for any authenticated users setting:

- You don't mind who views the content, but you want to restrict how it is used. For example, you don't want the content to be edited, copied, or printed.
- You don't need to restrict who accesses the content, but you want to be able to confirm who opens it.
- You have a requirement that the content must be encrypted at rest and in transit, but it doesn't require access controls.

Choose permissions

When you choose which permissions to allow for those users or groups, you can select either:

- A predefined permissions level with a preset group of rights, such as Co-Author or Reviewer.
- Custom permissions, where you choose one or more usage rights.

For more information to help you select the appropriate permissions, see Usage rights and descriptions.

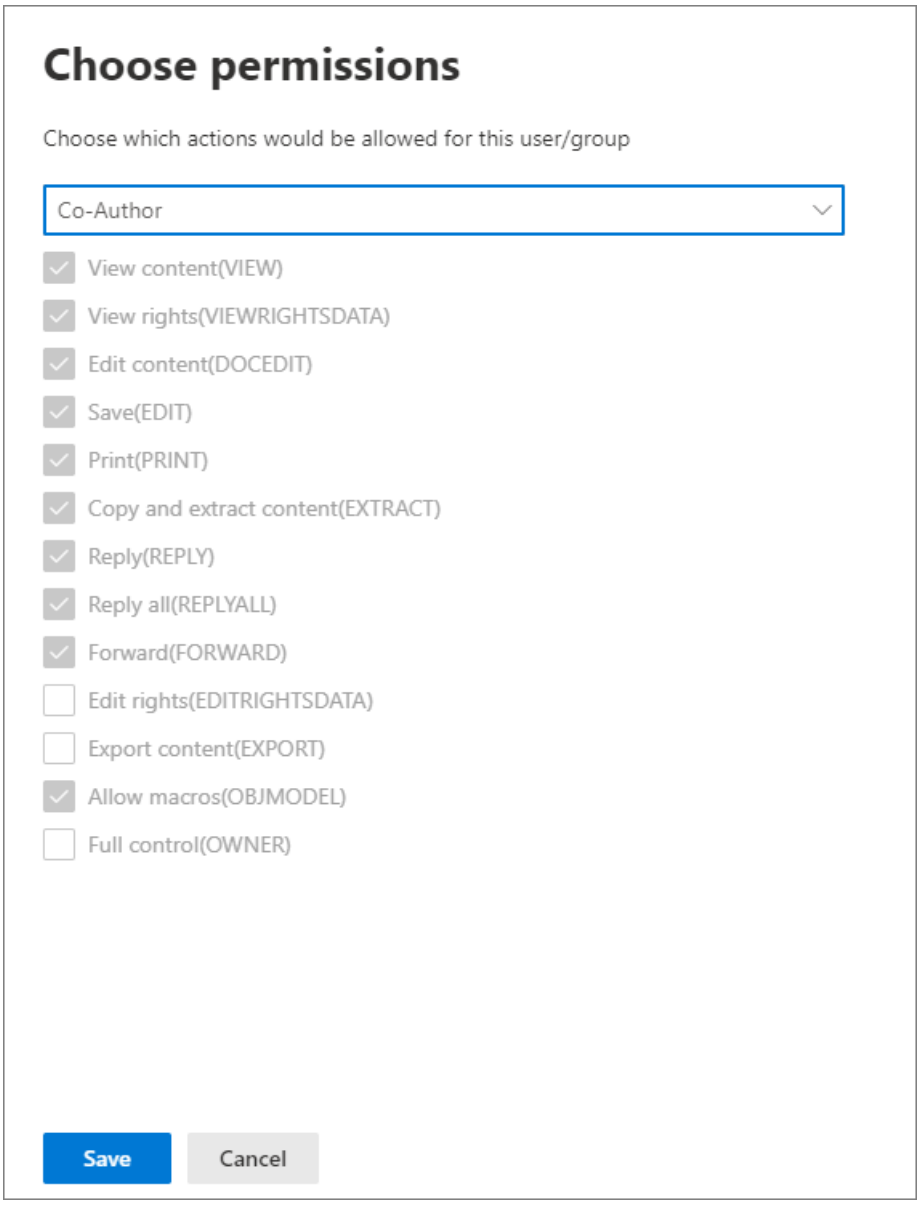

Note that the same label can grant different permissions to different users. For example, a single label can assign some users as Reviewer and a different user as Co-author, as shown in the following screenshot.

To do this, add users or groups, assign them permissions, and save those settings. Then repeat these steps, adding users and assigning them permissions, saving the settings each time. You can repeat this configuration as often as necessary, to define different permissions for different users.

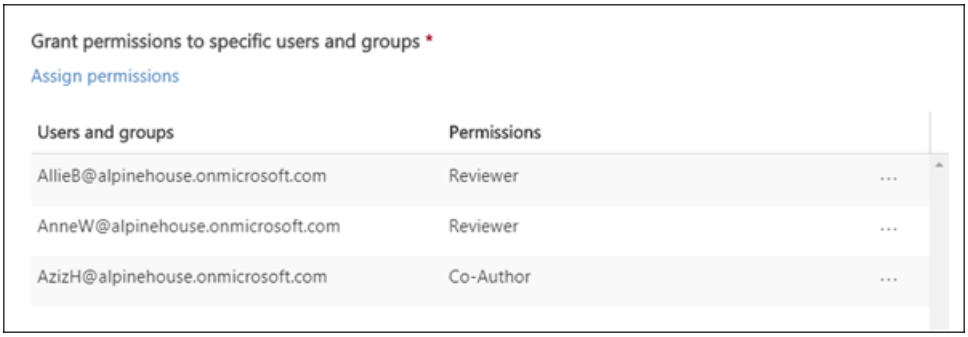

Rights Management issuer (user applying the sensitivity label) always has Full Control

Encryption for a sensitivity label uses the Azure Rights Management service from Azure Information Protection. When a user applies a sensitivity label to protect a document or email by using encryption, that user becomes the Rights Management issuer for that content.

The Rights Management issuer is always granted Full Control permissions for the document or email, and in addition:

- If the encryption settings include an expiration date, the Rights Management issuer can still open and edit the document or email after that date.
- The Rights Management issuer can always access the document or email offline.
- The Rights Management issuer can still open a document after it is revoked.

For more information, see Rights Management issuer and Rights Management owner.

Double Key Encryption

> **NOTE**
> This feature is currently supported only by the Azure Information Protection unified labeling client.

Select this option only after you have configured the Double Key Encryption service and you need to use this double key encryption for files that will have this label applied.

For more information, prerequisites, and configuration instructions, see Double Key Encryption (DKE).

Let users assign permissions

You can use these options to let users assign permissions when they manually apply a sensitivity label to content:

- In Outlook, a user can select restrictions equivalent to the Do Not Forward option for their chosen recipients.
- In Word, PowerPoint, and Excel, a user is prompted to select their own permissions for specific users, groups, or organizations.

  > **NOTE**
  > This option for Word, PowerPoint, and Excel is supported by the Azure Information Protection unified labeling client. For apps that use built-in labeling, check which apps support it.
  >
  > If this option is selected but isn't supported for a user's app, either that label doesn't display to the user, or the label displays for consistency, but it can't be applied with an explanation message to users.

When the options are supported, use the following table to identify when users see the sensitivity label:

SETTING	LABEL VISIBLE IN OUTLOOK	LABEL VISIBLE IN WORD, EXCEL, POWERPOINT
In Outlook, enforce restrictions equivalent to the Do Not Forward option	Yes	No
In Word, PowerPoint, and Excel, prompt users to specify permissions	No	Yes

When both settings are selected, the label is therefore visible in both Outlook and in Word, Excel, and PowerPoint.

A sensitivity label that lets users assign permissions can be applied to content only manually by users; it can't be auto-applied or used as a recommended label.

Configuring the user-assigned permissions:

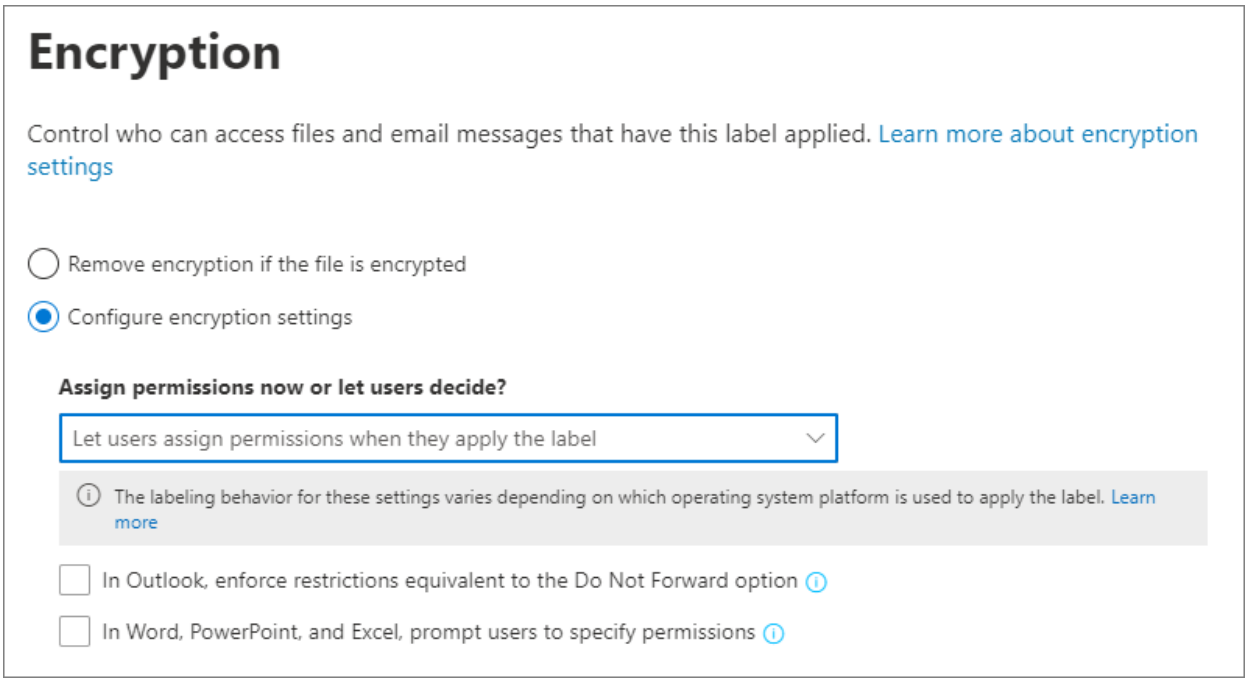

Outlook restrictions

In Outlook, when a user applies a sensitivity label that lets them assign permissions to a message, the restrictions are the same as the Do Not Forward option. The user will see the label name and description at the top of the message, which indicates the content's being protected. Unlike Word, PowerPoint, and Excel (see the next section), users aren't prompted to select specific permissions.

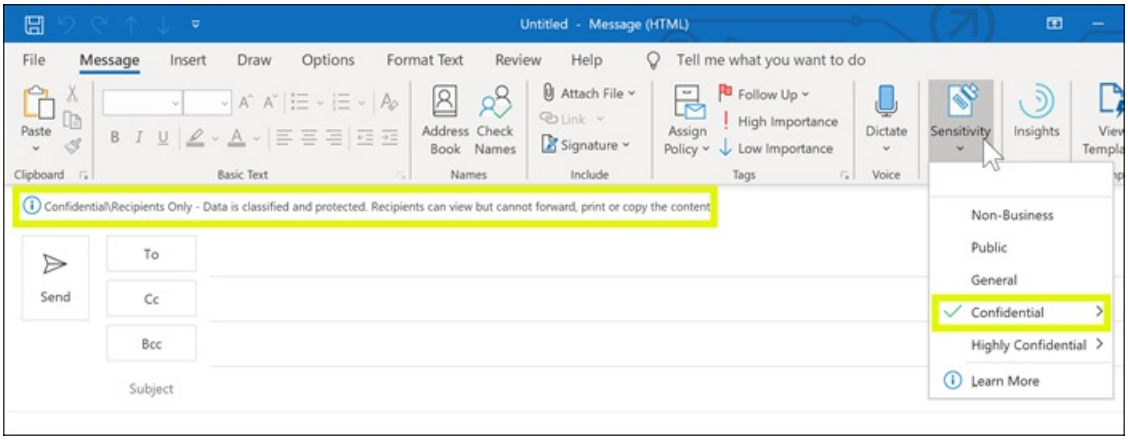

When the Do Not Forward option is applied to an email, the email is encrypted and recipients must be authenticated. Then, the recipients cannot forward it, print it, or copy from it. For example, in the Outlook client, the Forward button is not available, the Save As and Print menu options are not available, and you cannot add or change recipients in the To, Cc, or Bcc boxes.

Unencrypted Office documents that are attached to the email automatically inherit the same restrictions. The usage rights applied to these documents are Edit Content, Edit; Save; View, Open, Read; and Allow Macros. If the user wants different usage rights for an attachment, or the attachment is not an Office document that supports this inherited protection, the user needs to protect the file before attaching it to the email.

Word, PowerPoint, and Excel permissions

In Word, PowerPoint, and Excel, when a user applies a sensitivity label that lets them assign permissions to a document, they are prompted to specify their choice of users and permissions when the encryption is applied.

For example, with the Azure Information Protection unified labeling client, users can:

- Select a permission level, such as Viewer (which assigns View Only permission) or Co-Author (which assigns View, Edit, Copy, and Print permissions).

- Select users, groups, or organizations. This can include people both inside or outside your organizations.
- Set an expiration date, after which the selected users cannot access the content. For more information, see the above section Rights Management use license for offline access.

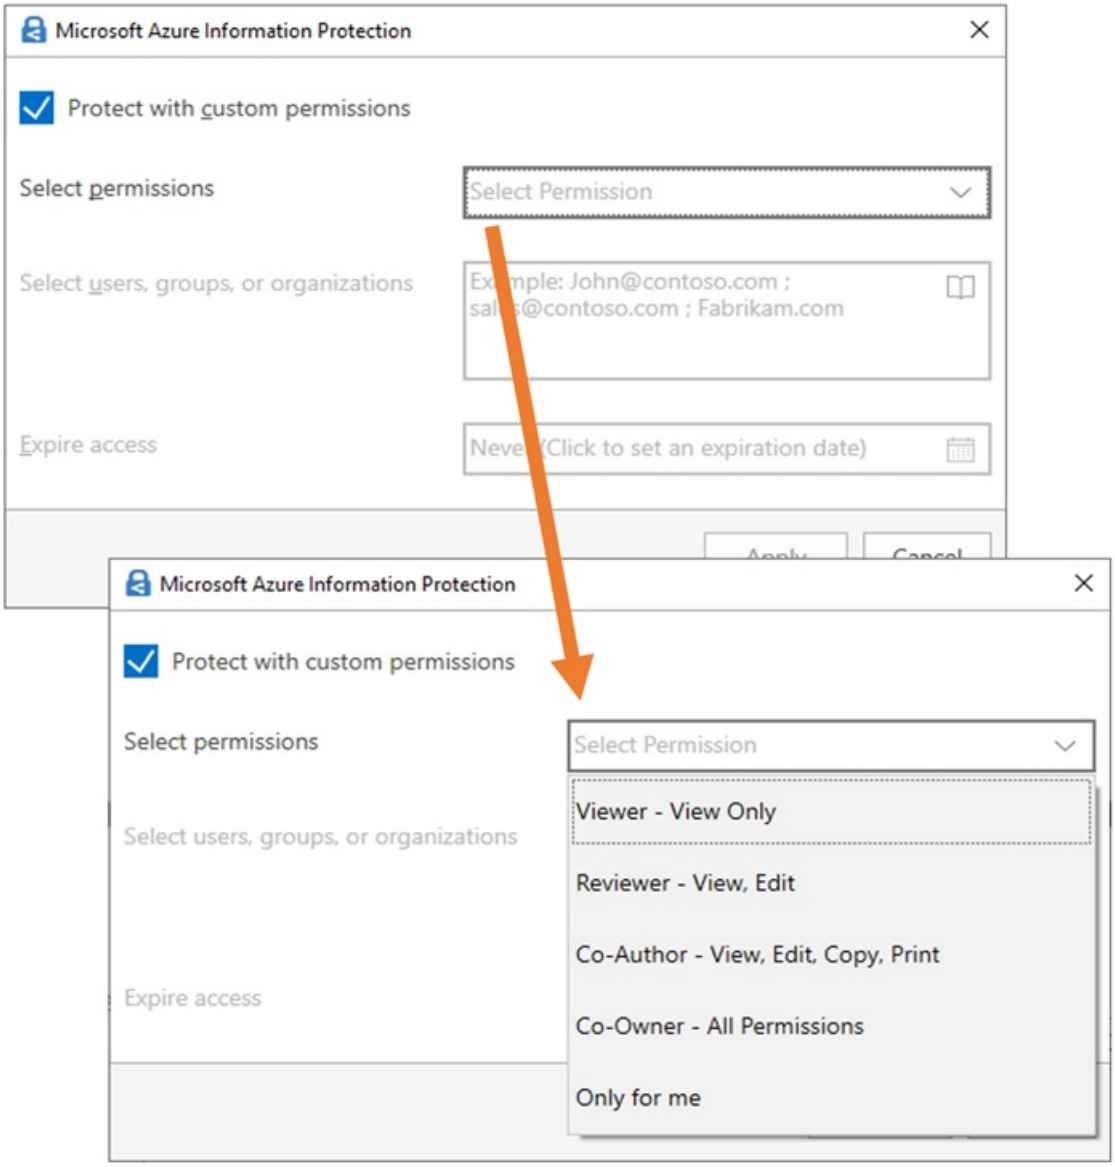

For built-in labeling, users see the same dialog box if they select the following:

- Windows: **File** tab > **Info** > **Protect Document** > **Restrict Access** > **Restricted Access**
- macOS: **Review** tab > **Protection** > **Permissions** > **Restricted Access**

Example configurations for the encryption settings

For each example that follows, do the configuration from the **Encryption** page of the wizard when **Configure encryption settings** is selected:

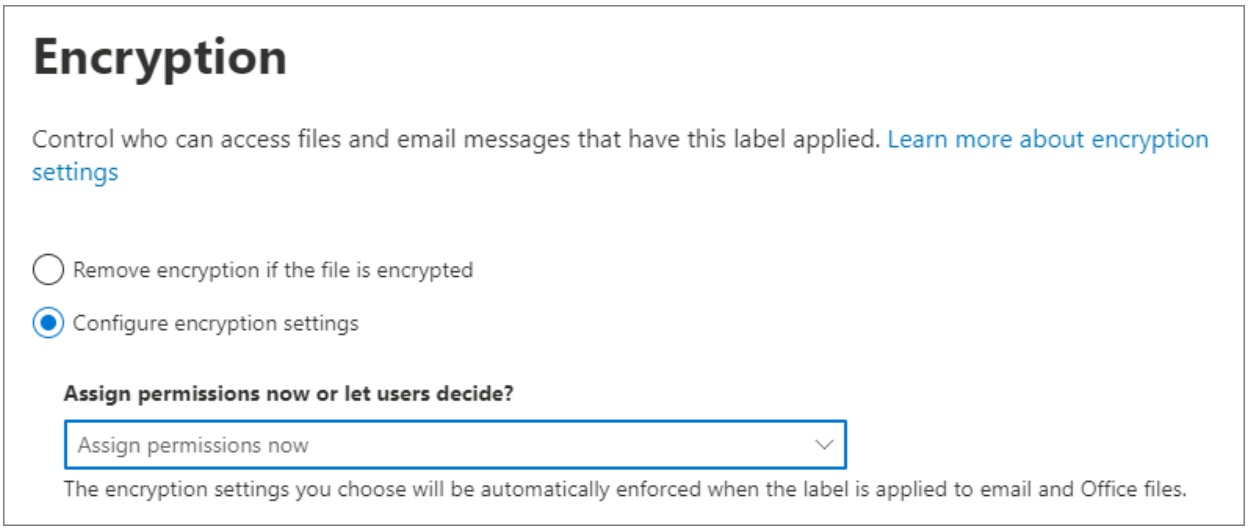

Example 1: Label that applies Do Not Forward to send an encrypted email to a Gmail account

This label displays only in Outlook and Outlook on the web, and you must use Exchange Online. Instruct users to select this label when they need to send an encrypted email to people using a Gmail account (or any other email account outside your organization).

Your users type the Gmail email address in the **To** box. Then, they select the label and the Do Not Forward option is automatically added to the email. The result is that recipients cannot forward the email, or print it, copy from it, or save the email outside their mailbox by using the **Save As** option.

1. On the **Encryption** page: For **Assign permissions now or let users decide?** select **Let users assign permissions when they apply the label**.
2. Select the checkbox: **In Outlook, enforce restrictions equivalent to the Do Not Forward option**.
3. If selected, clear the checkbox: **In Word, PowerPoint, and Excel, prompt users to specify permissions**.
4. Select **Next** and complete the wizard.

Example 2: Label that restricts read-only permission to all users in another organization

This label is suitable for sharing very sensitive documents as read-only, and the documents always require an internet connection to view them.

This label is not suitable for emails.

1. On the **Encryption** page: For **Assign permissions now or let users decide?** select **Assign permissions now**.
2. For **Allow offline access**, select **Never**.
3. Select **Assign permissions**.
4. On the **Assign permissions** pane, select **Add specific email addresses or domains**.
5. In the text box, enter the name of a domain from the other organization, for example, **fabrikam.com**. Then select **Add**.
6. Select **Choose permissions**.
7. On the **Choose permissions** pane, select the dropdown box, select **Viewer**, and then select **Save**.
8. Back on the **Assign Permissions** pane, select **Save**.
9. On the **Encryption** page, select **Next** and complete the wizard.

Example 3: Add external users to an existing label that encrypts content

The new users that you add will be able open documents and emails that have already been protected with this label. The permissions that you grant these users can be different from the permissions that the existing users have.

1. On the **Encryption** page: For **Assign permissions now or let users decide?** make sure **Assign permissions now** is selected.
2. Select **Assign permissions**.
3. On the **Assign permissions** pane, select **Add specific email addresses or domains**.
4. In the text box, enter the email address of the first user (or group) to add, and then select **Add**.
5. Select **Choose permissions**.
6. On the **Choose permissions** pane, select the permissions for this user (or group), and then select **Save**.
7. Back on the **Assign Permissions** pane, repeat steps 3 through 6 for each user (or group) that you want to add to this label. Then click **Save**.
8. On the **Encryption** page, select **Next** and complete the wizard.

Example 4: Label that encrypts content but doesn't restrict who can access it

This configuration has the advantage that you don't need to specify users, groups, or domains to encrypt an email or document. The content will still be encrypted and you can still specify usage rights, an expiry date, and offline access.

Use this configuration only when you do not need to restrict who can open the protected document or email. More information about this setting

1. On the **Encryption** page: For **Assign permissions now or let users decide?** make sure **Assign permissions now** is selected.
2. Configure settings for **User access to content expires** and **Allow offline access** as required.
3. Select **Assign permissions**.
4. On the **Assign permissions** pane, select **Add any authenticated users**.

   For **Users and groups**, you see **Authenticated users** automatically added. You can't change this value, only delete it, which cancels the **Add any authenticated users** selection.
5. Select **Choose permissions**.
6. On the **Choose permissions** pane, select the dropdown box, select the permissions you want, and then select **Save**.
7. Back on the **Assign Permissions** pane, select **Save**.
8. On the **Encryption** page, select **Next** and complete the wizard.

Considerations for encrypted content

Encrypting your most sensitive documents and emails helps to ensure that only authorized people can access this data. However, there are some considerations to take into account:

- If your organization hasn't enabled sensitivity labels for Office files in SharePoint and OneDrive:
  - Search, eDiscovery, and Delve will not work for encrypted files.
  - DLP policies work for the metadata of these encrypted files (including retention label information) but

not the content of these files (such as credit card numbers within files).
    - Users can't open encrypted files using Office on the web. When sensitivity labels for Office files in SharePoint and OneDrive are enabled, users can use Office on the web to open encrypted files, with some limitations that include encryption that has been applied with an on-premises key (known as "hold your own key", or HYOK), double key encryption, and encryption that has been applied independently from a sensitivity label.
- If you share encrypted documents with people outside your organization, you might need to create guest accounts and modify Conditional Access policies. For more information, see Sharing encrypted documents with external users.
- For multiple users to edit an encrypted file at the same time, they must all be using Office for the web. If this isn't the case, and the file is already open:
    - In Office apps (Windows, Mac, Android, and iOS), users see a **File In Use** message with the name of the person who has checked out the file. They can then view a read-only copy or save and edit a copy of the file, and receive notification when the file is available.
    - In Office for the web, users see an error message that they can't edit the document with other people. They can then select **Open in Reading View**.
- The AutoSave functionality in Office apps (Windows, Mac, Android, and iOS) is disabled for encrypted files. Users see a message that the file has restricted permissions that must be removed before AutoSave can be turned on.
- Encrypted files might take longer to open in Office apps (Windows, Mac, Android, and iOS).
- If a label that applies encryption is added by using an Office app when the document is checked out in SharePoint, and the user then discards the checkout, the document remains labeled and encrypted.
- The following actions for encrypted files aren't supported from Office apps (Windows, Mac, Android, and iOS), and users see an error message that something went wrong. However, SharePoint functionality can be used as an alternative:
    - View, restore, and save copies of previous versions. As an alternative, users can do these actions using Office on the web when you enable and configure versioning for a list or library.
    - Change the name or location of files. As an alternative, users can rename a file, folder, or link in a document library in SharePoint.

For the best collaboration experience for files that are encrypted by a sensitivity label, we recommend you use sensitivity labels for Office files in SharePoint and OneDrive and Office for the web.

Important prerequisites

Before you can use encryption, you might need to do some configuration tasks.

- Activate protection from Azure Information Protection

  For sensitivity labels to apply encryption, the protection service (Azure Rights Management) from Azure Information Protection must be activated for your tenant. In newer tenants, this is the default setting, but you might need to manually activate the service. For more information, see Activating the protection service from Azure Information Protection.

- Configure Exchange for Azure Information Protection

  Exchange does not have to be configured for Azure Information Protection before users can apply labels in Outlook to encrypt their emails. However, until Exchange is configured for Azure Information Protection, you do not get the full functionality of using Azure Rights Management protection with Exchange.

For example, users cannot view encrypted emails on mobile phones or with Outlook on the web, encrypted emails cannot be indexed for search, and you cannot configure Exchange Online DLP for Rights Management protection.

To ensure that Exchange can support these additional scenarios, see the following:

- For Exchange Online, see the instructions for Exchange Online: IRM Configuration.
- For Exchange on-premises, you must deploy the RMS connector and configure your Exchange servers.

Next steps

Need to share your labeled and encrypted documents with people outside your organization? See Sharing encrypted documents with external users.

Apply a sensitivity label to content automatically

2/18/2021 • 17 minutes to read • Edit Online

Microsoft 365 licensing guidance for security & compliance.

NOTE

For information about automatically applying a sensitivity label in Azure Purview (preview), see Automatically label your content in Azure Purview.

When you create a sensitivity label, you can automatically assign that label to files and emails when it matches conditions that you specify.

This ability to apply sensitivity labels to content automatically is important because:

- You don't need to train your users when to use each of your classifications.
- You don't need to rely on users to classify all content correctly.
- Users no longer need to know about your policies—they can instead focus on their work.

When content has been manually labeled, that label will never be replaced by automatic labeling. However, automatic labeling can replace a lower priority label that was automatically applied.

There are two different methods for automatically applying a sensitivity label to content in Microsoft 365:

- **Client-side labeling when users edit documents or compose (also reply or forward) emails**: Use a label that's configured for auto-labeling for files and emails (includes Word, Excel, PowerPoint, and Outlook).

  This method supports recommending a label to users, as well as automatically applying a label. But in both cases, the user decides whether to accept or reject the label, to help ensure the correct labeling of content. This client-side labeling has minimal delay for documents because the label can be applied even before the document is saved. However, not all client apps support auto-labeling. This capability is supported by the Azure Information Protection unified labeling client, and some versions of Office.

  For configuration instructions, see How to configure auto-labeling for Office apps on this page.

- **Service-side labeling when content is already saved (in SharePoint or OneDrive) or emailed (processed by Exchange Online)**: Use an auto-labeling policy.

  You might also hear this method referred to as auto-labeling for data at rest (documents in SharePoint and OneDrive) and data in transit (email that is sent or received by Exchange). For Exchange, it doesn't include emails at rest (mailboxes).

  Because this labeling is applied by services rather than by applications, you don't need to worry about what apps users have and what version. As a result, this capability is immediately available throughout your organization and suitable for labeling at scale. Auto-labeling policies don't support recommended labeling because the user doesn't interact with the labeling process. Instead, the administrator runs the policies in simulation mode to help ensure the correct labeling of content before actually applying the label.

  For configuration instructions, see How to configure auto-labeling policies for SharePoint, OneDrive, and

Exchange on this page.

Specific to auto-labeling for SharePoint and OneDrive:

- Office files for Word, PowerPoint, and Excel are supported. Open XML format is supported (such as .docx and .xlsx) but not Microsoft Office 97-2003 format (such as .doc and .xls).
  - These files can be auto-labeled when they are not part of an open session and whether they have been created, uploaded, or changed since you created auto-labeling policies, or they are existing files that have not been changed since you created your auto-labeling policies.
- Maximum of 25,000 automatically labeled files in your tenant per day.
- Maximum of 10 auto-labeling policies per tenant, each targeting up to 10 sites (SharePoint or OneDrive).
- Existing values for modified, modified by, and the date are not changed as a result of auto-labeling policies—for both simulation mode and when labels are applied.
- When the label applies encryption, the Rights Management issuer and Rights Management owner is the account that last modified the file.

Specific to auto-labeling for Exchange:

- Unlike manual labeling or auto-labeling with Office apps, Office attachments (Word, Excel, and PowerPoint files) and PDF attachments are also scanned for the conditions you specify in your auto-labeling policy. When there is a match, the email is labeled but not the attachment.
  - For these Office files, Open XML format is supported (such as .docx and .xlsx) but not Microsoft Office 97-2003 format (such as .doc and .xls).
- If you have Exchange mail flow rules or data loss prevention (DLP) policies that apply IRM encryption: When content is identified by these rules or policies and an auto-labeling policy, the label is applied. If that label applies encryption, the IRM settings from the Exchange mail flow rules or DLP policies are ignored. However, if that label doesn't apply encryption, the IRM settings from the mail flow rules or DLP policies are applied in addition to the label.
- Email that has IRM encryption with no label will be replaced by a label with any encryption settings when there is a match by using auto-labeling.
- Incoming email is labeled when there is a match with your auto-labeling conditions. However, if the label is configured for encryption, that encryption isn't applied.
- When the label applies encryption, the Rights Management issuer and Rights Management owner is the person who sends the email.

Compare auto-labeling for Office apps with auto-labeling policies

Use the following table to help you identify the differences in behavior for the two complementary automatic labeling methods:

FEATURE OR BEHAVIOR	LABEL SETTING: AUTO-LABELING FOR FILES AND EMAILS	POLICY: AUTO-LABELING
App dependency	Yes	No *
Restrict by location	No	Yes
Conditions: Trainable classifiers	Yes	No
Conditions: Sharing options and additional options for email	No	Yes

FEATURE OR BEHAVIOR	LABEL SETTING: AUTO-LABELING FOR FILES AND EMAILS	POLICY: AUTO-LABELING
Recommendations, policy tooltip, and user overrides	Yes	No
Simulation mode	No	Yes
Exchange attachments checked for conditions	No	Yes
Apply visual markings	Yes	Yes (email only)
Override IRM encryption applied without a label	Yes if the user has the minimum usage right of Export	Yes (email only)
Label incoming email	No	Yes (encryption not applied)

* Auto-labeling isn't currently available in all regions. If your tenant can't support this functionality, the Auto-labeling tab isn't visible in the admin labeling center.

How multiple conditions are evaluated when they apply to more than one label

The labels are ordered for evaluation according to their position that you specify in the policy: The label positioned first has the lowest position (least sensitive) and the label positioned last has the highest position (most sensitive). For more information on priority, see Label priority (order matters).

Don't configure a parent label to be applied automatically or recommended

Remember, you can't apply a parent label (a label with sublabels) to content. Make sure that you don't configure a parent label to be auto-applied or recommended in Office apps, and don't select a parent label for an auto-labeling policy. If you do, the parent label won't be applied to content.

To use automatic labeling with sublabels, make sure you publish both the parent label and the sublabel.

For more information on parent labels and sublabels, see Sublabels (grouping labels).

How to configure auto-labeling for Office apps

Automatic labeling in Office apps for Windows is supported by the Azure Information Protection unified labeling client. For built-in labeling in Office apps, this capability is in different stages of availability for different apps.

The auto-labeling settings for Office apps are available when you create or edit a sensitivity label. Make sure **Files & emails** is selected for the label's scope:

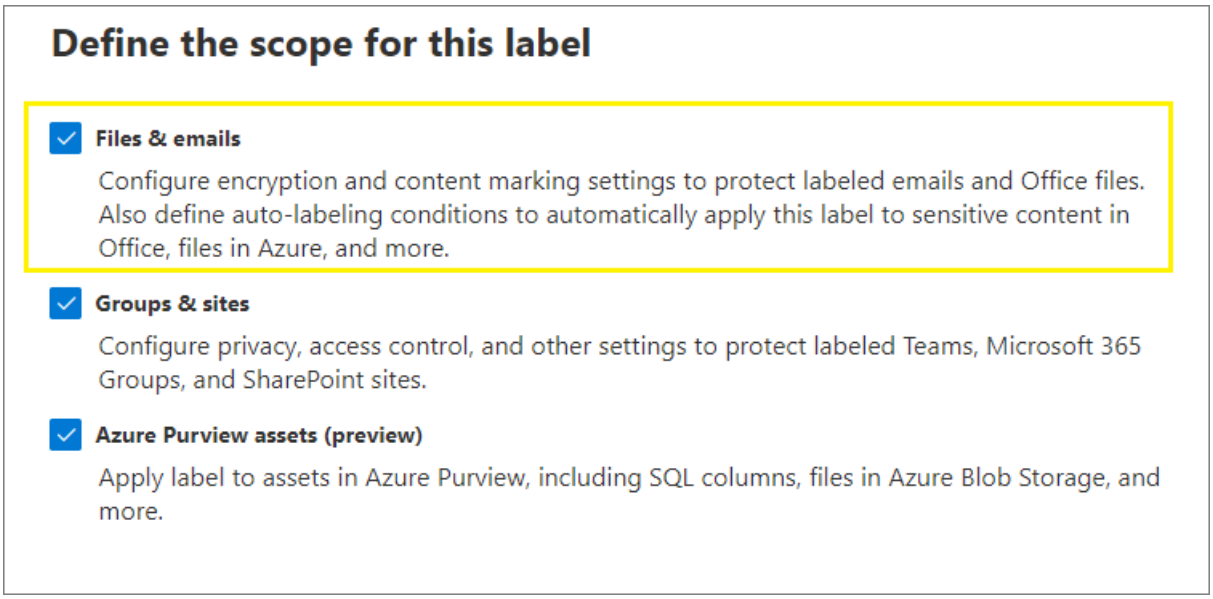

As you move through the wizard, you see the **Auto-labeling for files and emails** page where you can choose from a list of sensitive info types or trainable classifiers:

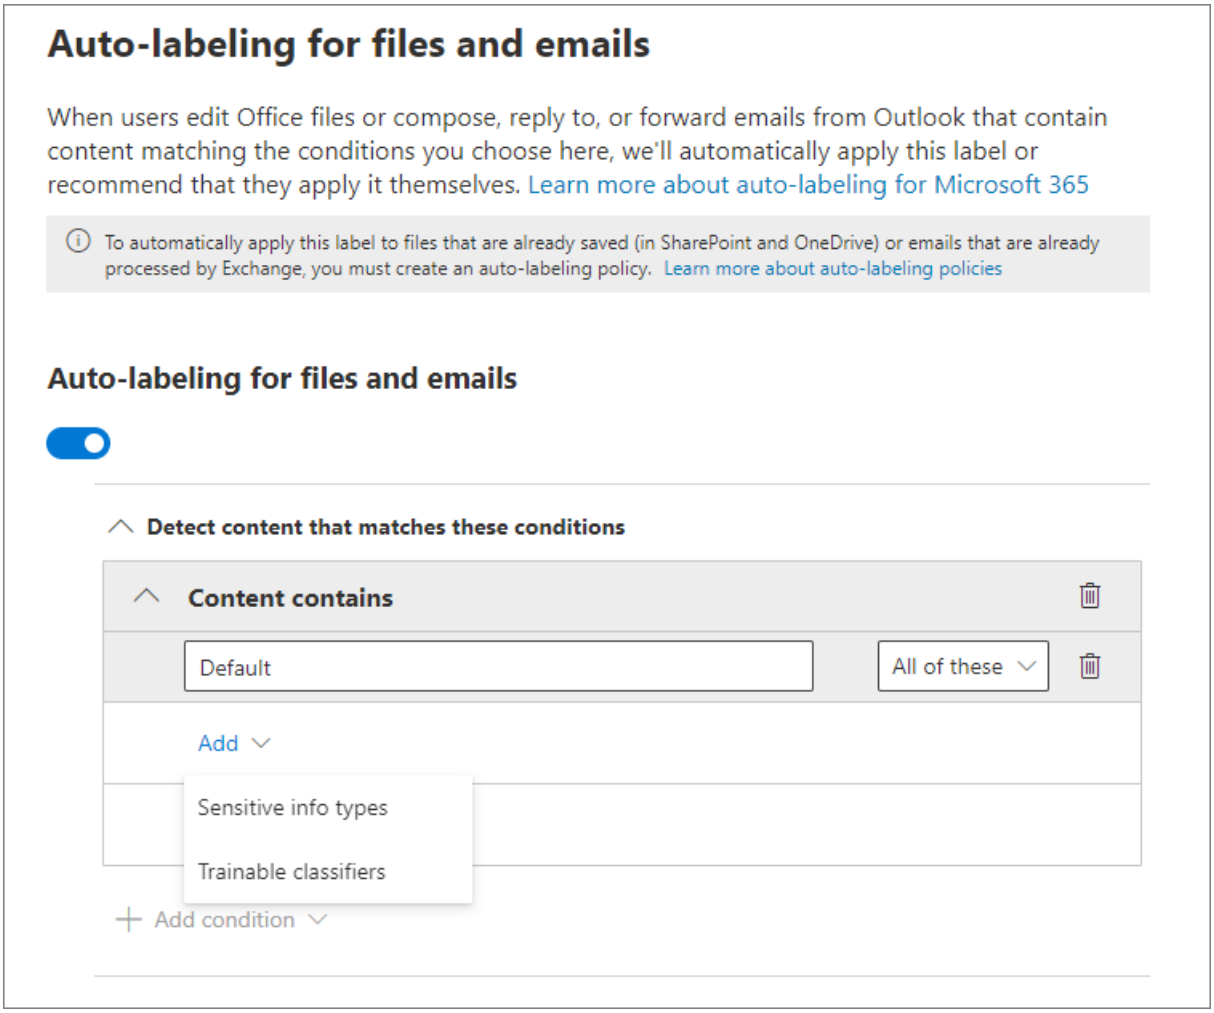

When this sensitivity label is automatically applied, the user sees a notification in their Office app. For example:

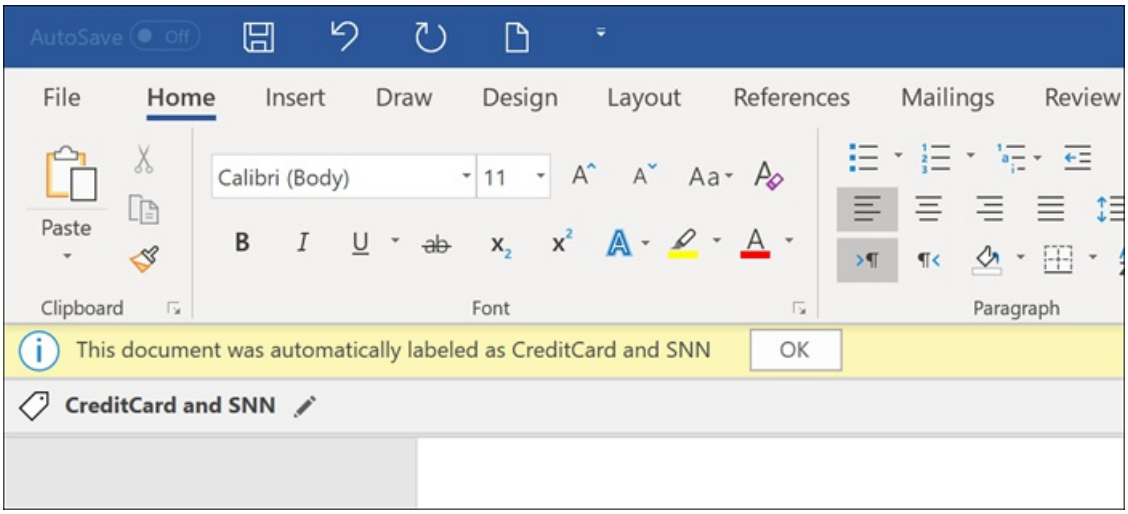

Configuring sensitive info types for a label

When you select the **Sensitive info types** option, you see the same list of sensitive information types as when you create a data loss prevention (DLP) policy. So you can, for example, automatically apply a Highly Confidential label to any content that contains customers' personal information, such as credit card numbers, social security numbers, or passport numbers:

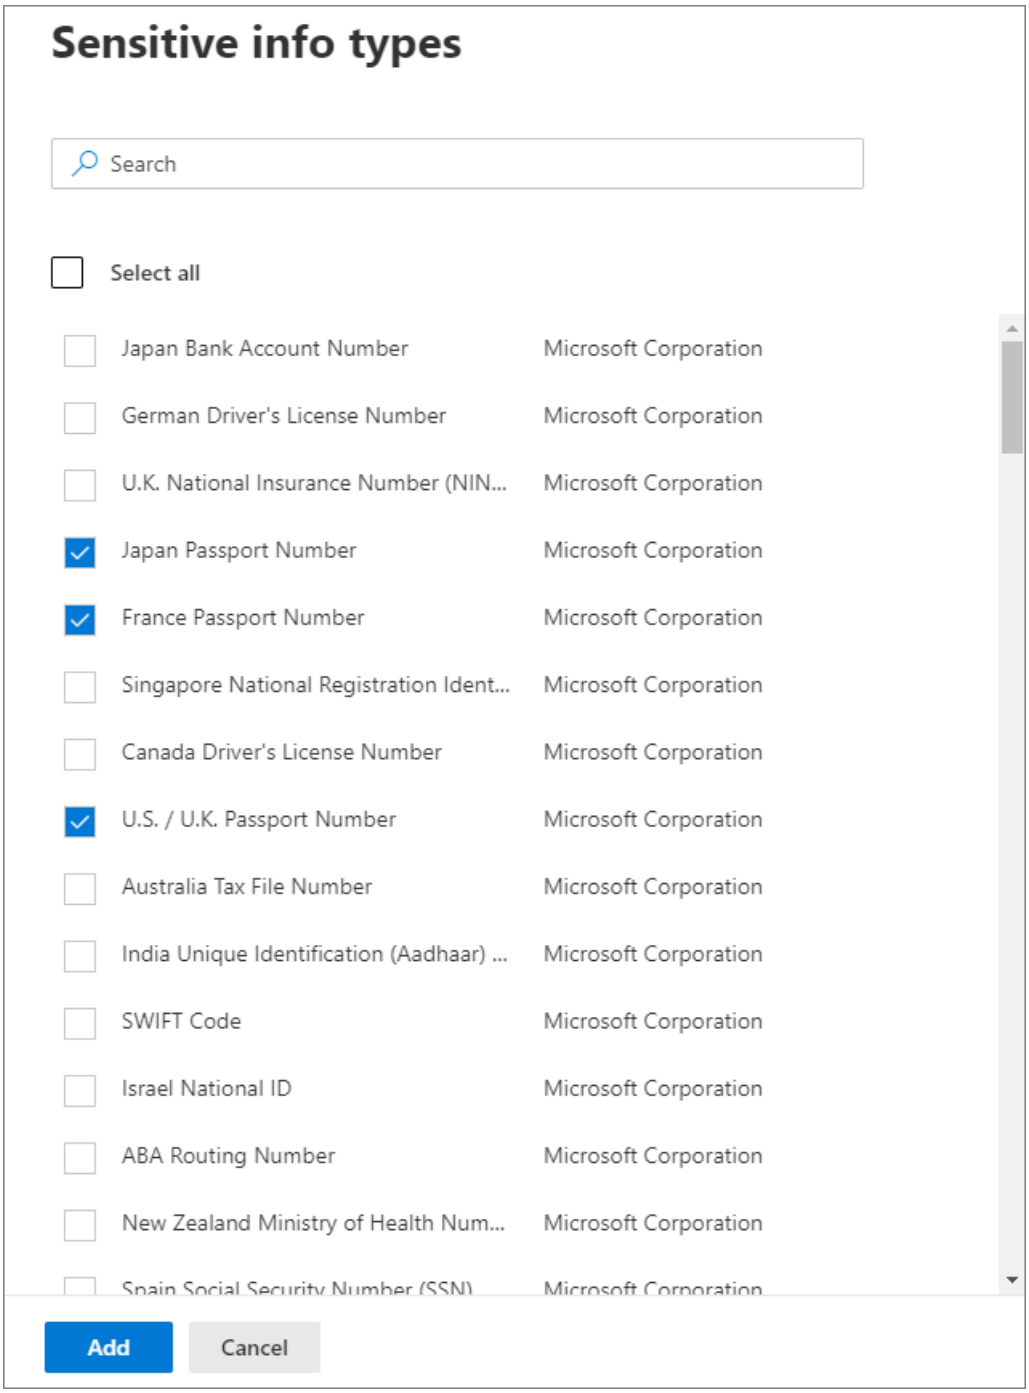

Similarly to when you configure DLP policies, you can then refine your condition by changing the instance count and match accuracy. For example:

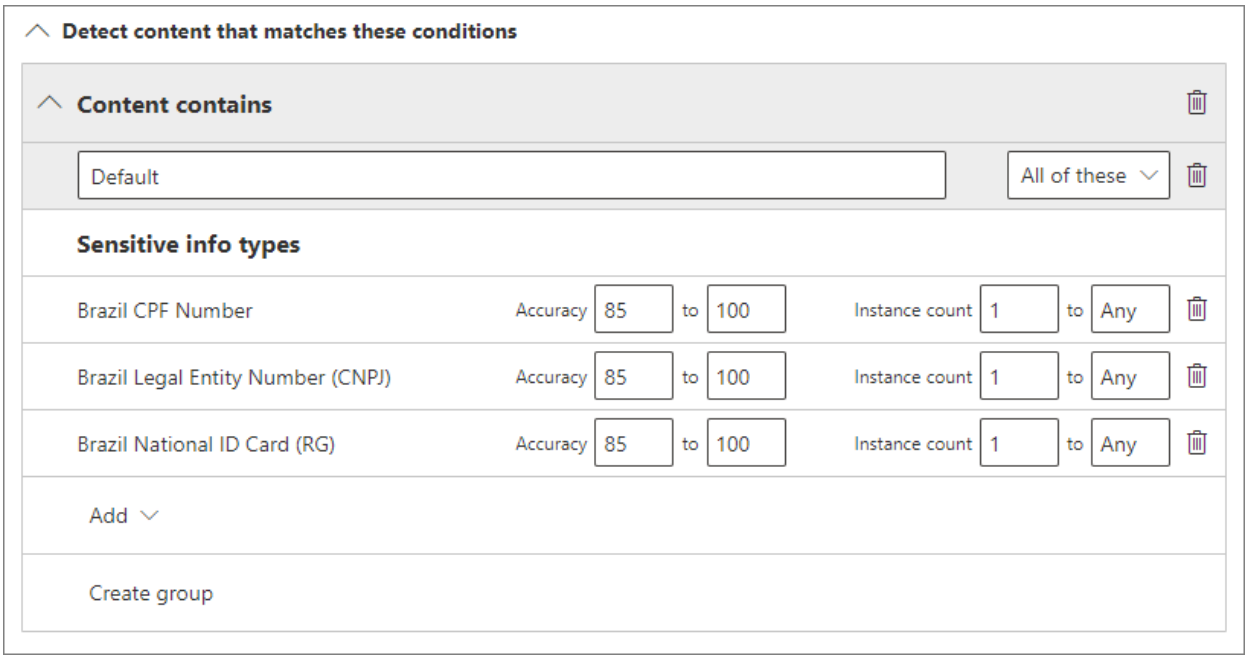

You can learn more about these configuration options from the DLP documentation: Tuning rules to make them easier or harder to match.

Also similarly to DLP policy configuration, you can choose whether a condition must detect all sensitive information types, or just one of them. And to make your conditions more flexible or complex, you can add groups and use logical operators between the groups.

Configuring trainable classifiers for a label

This option is currently in preview. If you use this option, make sure you have published in your tenant at least one other sensitivity label that's configured for auto-labeling and the sensitive info types option.

When you select the **Trainable classifiers** option, select one or more of the built-in trainable classifiers from Microsoft. If you've created your own custom trainable classifiers, these are also available to select:

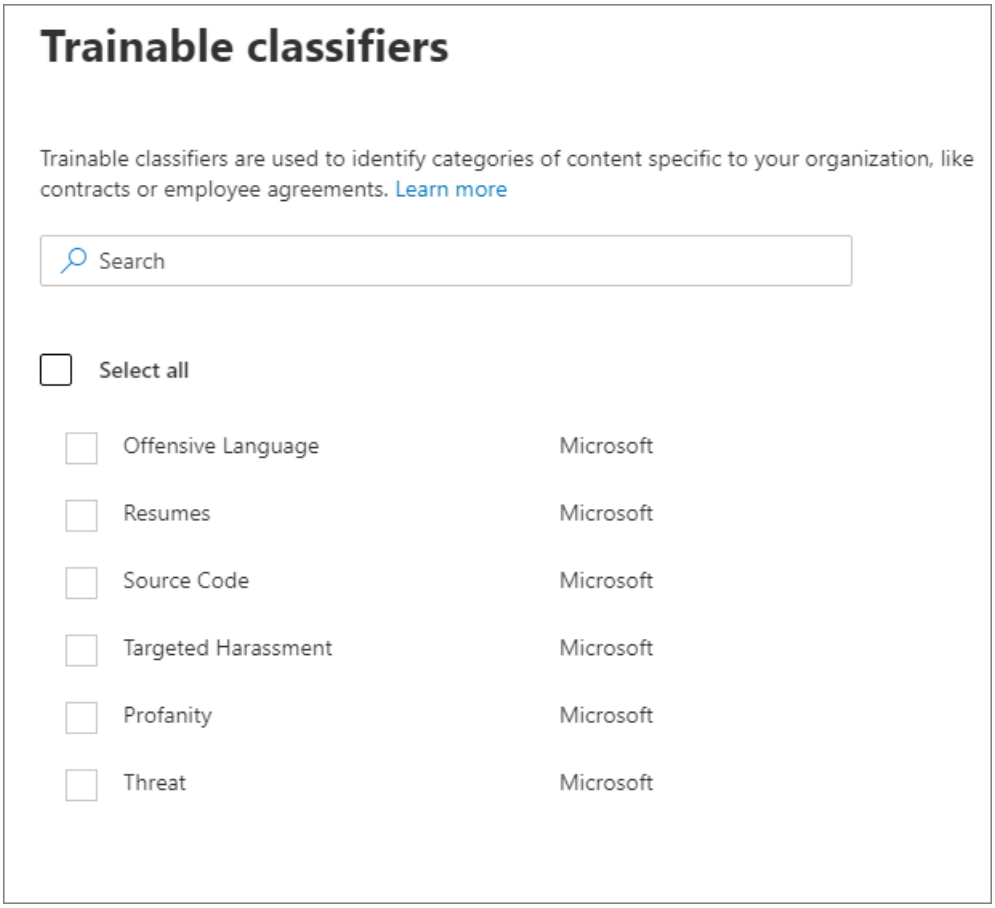

Caution

We are deprecating the **Offensive Language** built-in classifier because it has been producing a high number of false positives. Don't use this built-in classifier and if you are currently using it, you should move your business processes off it. We recommend using the **Targeted Harassment**, **Profanity**, and **Threat** built-in classifiers instead.

For more information about these classifiers, see Learn about trainable classifiers.

During the preview period for this option, the following apps support trainable classifiers for sensitivity labels:

- Microsoft 365 Apps for enterprise (formerly Office 365 ProPlus) for Windows, now rolling out to the Current Channel in version 2006 and later:
  - Word
  - Excel
  - PowerPoint
- Office for the web apps, when you have enabled sensitivity labels for Office files in SharePoint and OneDrive:
  - Word
  - Excel
  - PowerPoint
  - Outlook

Recommend that the user applies a sensitivity label

If you prefer, you can recommend to your users that they apply the label. With this option, your users can accept the classification and any associated protection, or dismiss the recommendation if the label isn't suitable for their content.

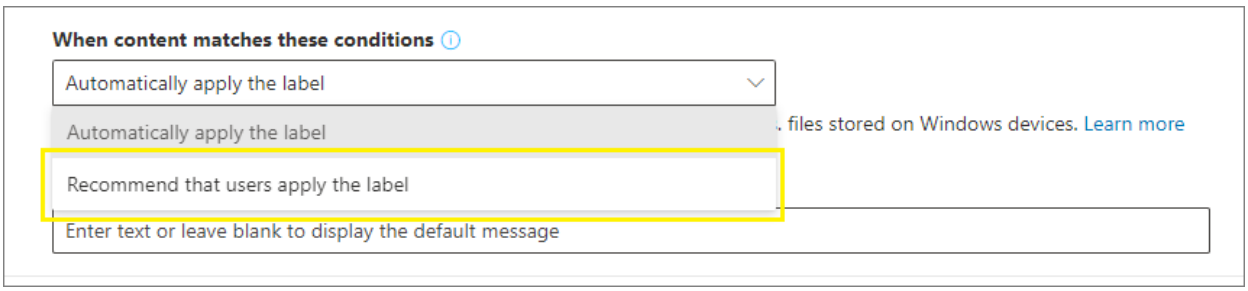

Here's an example of a prompt from the Azure Information Protection unified labeling client when you configure a condition to apply a label as a recommended action, with a custom policy tip. You can choose what text is displayed in the policy tip.

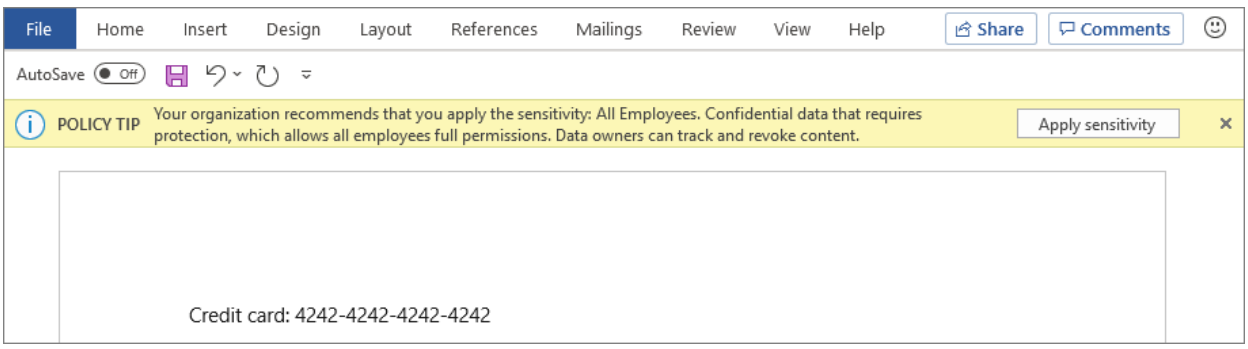

When automatic or recommended labels are applied

The implementation of automatic and recommended labeling in Office apps depend on whether you're using labeling that's built into Office, or the Azure Information Protection unified labeling client. In both cases, however:

- You can't use automatic labeling for documents and emails that were previously manually labeled, or previously automatically labeled with a higher sensitivity. Remember, you can only apply a single sensitivity label to a document or email (in addition to a single retention label).
- You can't use recommended labeling for documents or emails that were previously labeled with a higher sensitivity. When the content's already labeled with a higher sensitivity, the user won't see the prompt with the recommendation and policy tip.

Specific to built-in labeling:

- Not all Office apps support automatic (and recommended) labeling. For more information, see Support for sensitivity label capabilities in apps.
- For recommended labels in the desktop versions of Word, the sensitive content that triggered the recommendation is flagged so that users can review and remove the sensitive content instead of applying the recommended sensitivity label.
- For details about how these labels are applied in Office apps, example screenshots, and how sensitive information is detected, see Automatically apply or recommend sensitivity labels to your files and emails in Office.

Specific to the Azure Information Protection unified labeling client:

- Automatic and recommended labeling applies to Word, Excel, and PowerPoint when you save a document, and to Outlook when you send an email.
- For Outlook to support recommended labeling, you must first configure an advanced policy setting.
- Sensitive information can be detected in the body text in documents and emails, and to headers and footers—but not in the subject line or attachments of email.

How to configure auto-labeling policies for SharePoint, OneDrive, and Exchange

Make sure you're aware of the prerequisites before you configure auto-labeling policies.

Prerequisites for auto-labeling policies

- Simulation mode:
  - Auditing for Microsoft 365 must be turned on. If you need to turn on auditing or you're not sure whether auditing is already on, see Turn audit log search on or off.
  - To view file contents in the source view, you must have the **Content Explorer Content Viewer** role. Global admins don't have this role by default. If you don't have this permission, you don't see the preview pane when you select an item from the **Matched Items** tab.
- To auto-label files in SharePoint and OneDrive:
  - You have enabled sensitivity labels for Office files in SharePoint and OneDrive.
  - At the time the auto-labeling policy runs, the file mustn't be open by another process or user. A file that's checked out for editing falls into this category.
- If you plan to use custom sensitive information types rather than the built-in sensitivity types:
  - Custom sensitivity information types are evaluated for content that is added to SharePoint or OneDrive after the custom sensitivity information types are saved.
  - To test new custom sensitive information types, create them before you create your auto-labeling policy, and then create new documents with sample data for testing.
- One or more sensitivity labels created and published (to at least one user) that you can select for your

auto-labeling policies. For these labels:

- It doesn't matter if the auto-labeling in Office apps label setting is turned on or off, because that label setting supplements auto-labeling policies, as explained in the introduction.
- If the labels you want to use for auto-labeling are configured to use visual markings (headers, footers, watermarks), note that these are not applied to documents.
- If the labels apply encryption, they must be configured for the **Assign permissions now** setting.

Learn about simulation mode

Simulation mode is unique to auto-labeling policies and woven into the workflow. You can't automatically label documents and emails until your policy has run at least one simulation.

Workflow for an auto-labeling policy:

1. Create and configure an auto-labeling policy.
2. Run the policy in simulation mode, which can take 48 hours to complete.
3. Review the results, and if necessary, refine your policy. Rerun simulation mode and wait for it to complete again.
4. Repeat step 3 as needed.
5. Deploy in production.

The simulated deployment runs like the WhatIf parameter for PowerShell. You see results reported as if the auto-labeling policy had applied your selected label, using the rules that you defined. You can then refine your rules for accuracy if needed, and rerun the simulation. However, because auto-labeling for Exchange applies to emails that are sent and received, rather than emails stored in mailboxes, don't expect results for email in a simulation to be consistent unless you're able to send and receive the exact same email messages.

Simulation mode also lets you gradually increase the scope of your auto-labeling policy before deployment. For example, you might start with a single location, such as a SharePoint site, with a single document library. Then, with iterative changes, increase the scope to multiple sites, and then to another location, such as OneDrive.

Finally, you can use simulation mode to provide an approximation of the time needed to run your auto-labeling policy, to help you plan and schedule when to run it without simulation mode.

Creating an auto-labeling policy

1. In the Microsoft 365 compliance center, navigate to sensitivity labels:

   - **Solutions** > **Information protection**

   If you don't immediately see this option, first select **Show all**.

2. Select the **Auto-labeling** tab:

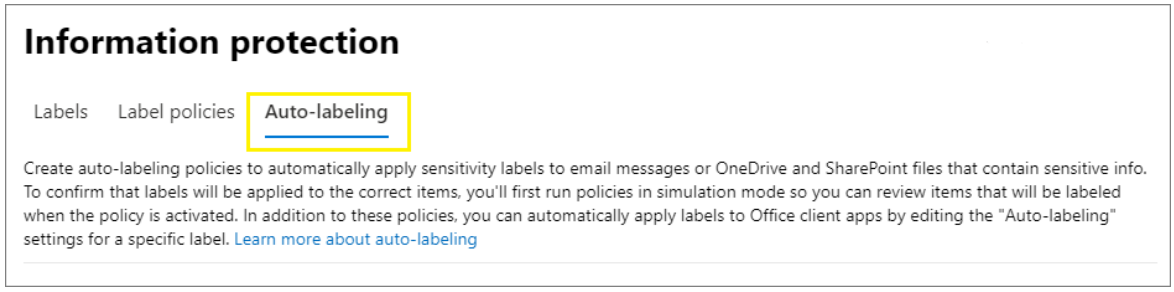

> **NOTE**
>
> If you don't see the **Auto-labeling** tab, this functionality isn't currently available in your region.

3. Select **+ Create auto-labeling policy**. This starts the New policy wizard:

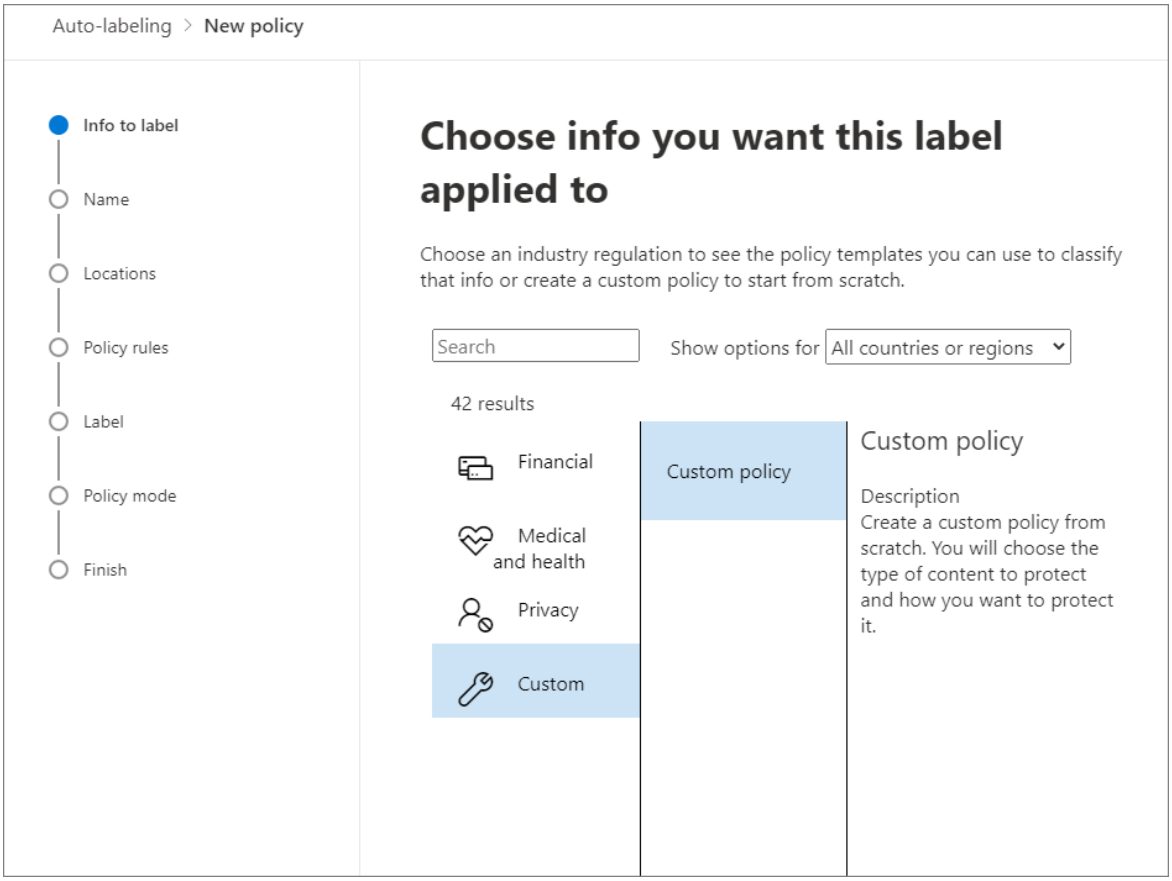

4. For the page **Choose info you want this label applied to**: Select one of the templates, such as **Financial** or **Privacy**. You can refine your search by using the **Show options for** dropdown. Or, select **Custom policy** if the templates don't meet your requirements. Select **Next**.

5. For the page **Name your auto-labeling policy**: Provide a unique name, and optionally a description to help identify the automatically applied label, locations, and conditions that identify the content to label.

6. For the page **Choose locations where you want to apply the label**: Select and specify locations for Exchange, SharePoint sites, and OneDrive. Then select **Next**.

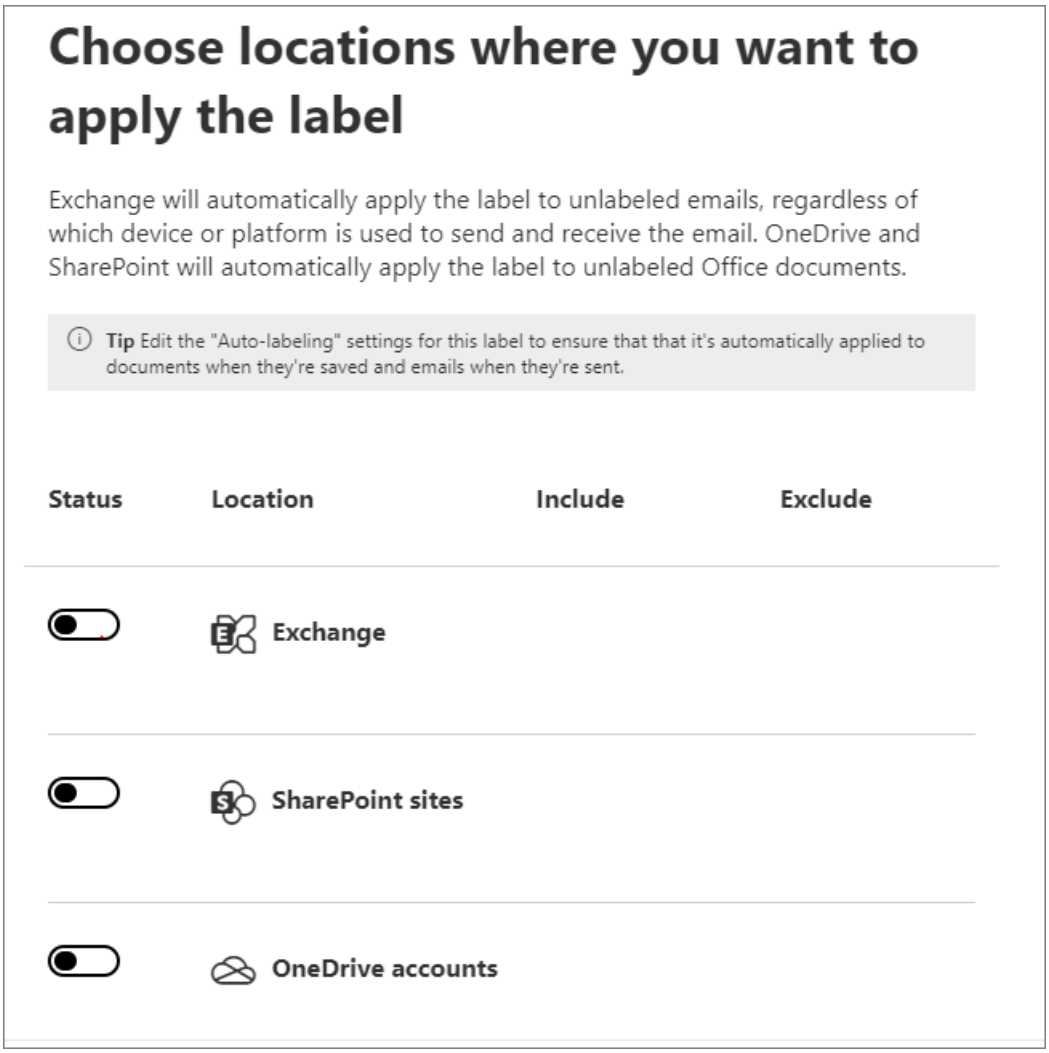

You must specify individual SharePoint sites and OneDrive accounts. For OneDrive, the URL for a user's OneDrive account is in the following format:

```
https://<tenant name>-my.sharepoint.com/personal/<user_name>_<tenant name>_com
```

For example, for a user in the contoso tenant that has a user name of "rsimone":

```
https://contoso-my.sharepoint.com/personal/rsimone_contoso_onmicrosoft_com
```

To verify the syntax for your tenant and identify URLs for users, see Get a list of all user OneDrive URLs in your organization.

7. For the **Set up common or advanced rules** page: Keep the default of **Common rules** to define rules that identify content to label across all your selected locations. If you need different rules per location, select **Advanced rules**. Then select **Next**.

   The rules use conditions that include sensitive information types and sharing options:

   - For sensitive information types, you can select both built-in and custom sensitive information types.
   - For the shared options, you can choose **only with people inside my organization** or **with people outside my organization**.

   If your only location is **Exchange**, or if you select **Advanced rules**, there are additional conditions that you can select:

   - Sender IP address is
   - Recipient domain is
   - Recipient is
   - Attachment's file extension is
   - Attachment is password protected

- Any email attachment's content could not be scanned
- Any email attachment's content didn't complete scanning

8. Depending on your previous choices, you'll now have an opportunity to create new rules by using conditions and exceptions.

   The configuration options for sensitive information types are the same as those you select for auto-labeling for Office apps. If you need more information, see Configuring sensitive info types for a label.

   When you have defined all the rules you need, and confirmed their status is on, select **Next** to move on to choosing a label to auto-apply.

9. For the **Choose a label to auto-apply** page: Select **+ Choose a label**, select a label from the **Choose a sensitivity label** pane, and then select **Next**.

10. For the **Decide if you want to test out the policy now or later** page: Select **Run policy in simulation mode** if you're ready to run the auto-labeling policy now, in simulation mode. Otherwise, select **Leave policy turned off**. Select **Next**:

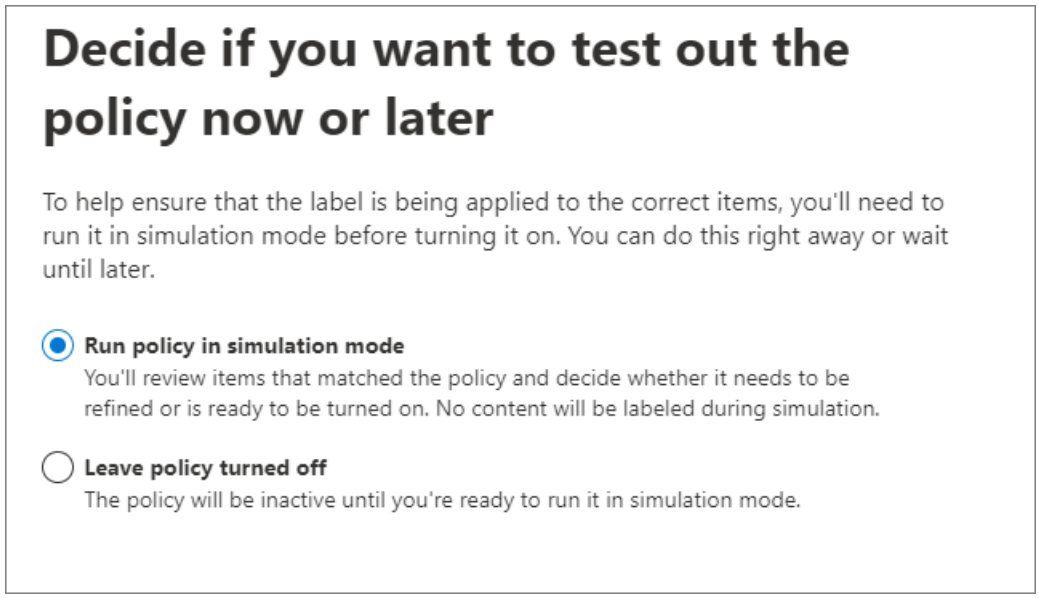

11. For the **Summary** page: Review the configuration of your auto-labeling policy and make any changes that needed, and complete the wizard.

Now on the **Information protection** > **Auto-labeling** page, you see your auto-labeling policy in the **Simulation** or **Off** section, depending on whether you chose to run it in simulation mode or not. Select your policy to see the details of the configuration and status (for example, **Policy simulation is still running**). For policies in simulation mode, select the **Matched items** tab to see which emails or documents matched the rules that you specified.

You can modify your policy directly from this interface:

- For a policy in the **Off** section, select the **Edit policy** button.
- For policy in the **Simulation** section, select the **Edit policy** option at the top of the page, from either tab:

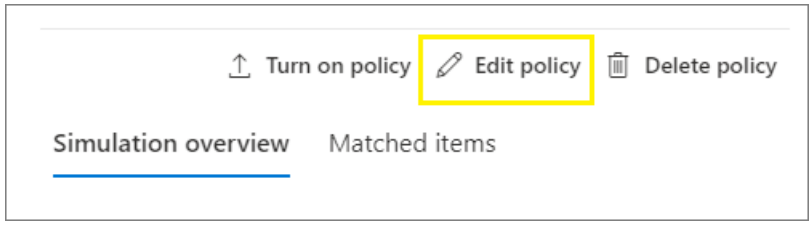

When you're ready to run the policy without simulation, select the **Turn on policy** option.

Your auto-policies run continuously until they are deleted. For example, new and modified documents will be included with the current policy settings.

You can also see the results of your auto-labeling policy by using content explorer when you have the appropriate permissions:

- **Content Explorer List Viewer** lets you see a file's label but not the file's contents.
- **Content Explorer Content Viewer** lets you see the file's contents.

> **TIP**
> You can also use content explorer to identify locations that have documents with sensitive information, but are unlabeled. Using this information, consider adding these locations to your auto-labeling policy, and include the identified sensitive information types as rules.

Use PowerShell for auto-labeling policies

You can use Security & Compliance Center PowerShell to create and configure auto-labeling policies. This means you can fully script the creation and maintenance of your auto-labeling policies, which also provides a more efficient method of specifying multiple URLs for OneDrive and SharePoint locations.

Before you run the commands in PowerShell, you must first connect to Security & Compliance Center PowerShell.

To create a new auto-labeling policy:

```
New-AutoSensitivityLabelPolicy -Name <AutoLabelingPolicyName> -SharePointLocation "<SharePointSiteLocation>"
-ApplySensitivityLabel <Label> -Mode TestWithoutNotifications
```

This command creates an auto-labeling policy for a SharePoint site that you specify. For a OneDrive location, use the *OneDriveLocation* parameter, instead.

To add additional sites to an existing auto-labeling policy:

```
$spoLocations = @("<SharePointSiteLocation1>","<SharePointSiteLocation2>")
Set-AutoSensitivityLabelPolicy -Identity <AutoLabelingPolicyName> -AddSharePointLocation $spoLocations -
ApplySensitivityLabel <Label> -Mode TestWithoutNotifications
```

This command specifies the additional SharePoint URLs in a variable that is then added to an existing auto-labeling policy. To add OneDrive locations instead, use the *AddOneDriveLocation* parameter with a different variable, such as *$OneDriveLocations*.

To create a new auto-labeling policy rule:

```
New-AutoSensitivityLabelRule -Policy <AutoLabelingPolicyName> -Name <AutoLabelingRuleName> -
ContentContainsSensitiveInformation @{"name"= "a44669fe-0d48-453d-a9b1-2cc83f2cba77"; "mincount" = "2"} -
Workload SharePoint
```

For an existing auto-labeling policy, this command creates a new policy rule to detect the sensitive information type of **U.S. social security number (SSN)**, which has an entity ID of a44669fe-0d48-453d-a9b1-2cc83f2cba77. To find the entity IDs for other sensitive information types, refer to Sensitive information type entity definitions.

For more information about the PowerShell cmdlets that support auto-labeling policies, their available parameters and some examples, see the following cmdlet help:

- Get-AutoSensitivityLabelPolicy
- New-AutoSensitivityLabelPolicy
- New-AutoSensitivityLabelRule
- Remove-AutoSensitivityLabelPolicy
- Remove-AutoSensitivityLabelRule
- Set-AutoSensitivityLabelPolicy
- Set-AutoSensitivityLabelRule

Use sensitivity labels to protect content in Microsoft Teams, Microsoft 365 groups, and SharePoint sites

2/18/2021 • 17 minutes to read • Edit Online

Microsoft 365 licensing guidance for security & compliance.

In addition to using sensitivity labels to classify and protect documents and emails, you can also use sensitivity labels to protect content in the following containers: Microsoft Teams sites, Microsoft 365 groups (formerly Office 365 groups), and SharePoint sites. For this container-level classification and protection, use the following label settings:

- Privacy (public or private) of teams sites and Microsoft 365 groups
- External user access
- External sharing from SharePoint sites (in preview)
- Access from unmanaged devices

IMPORTANT

The **Access from unmanaged devices** setting works in conjunction with the SharePoint feature to control access from unmanaged devices. You must configure this dependent SharePoint feature to use a sensitivity label that has this setting configured. Additional information is included in the instructions that follow.

When you apply this sensitivity label to a supported container, the label automatically applies the classification and configured protection settings to the site or group.

Content in these containers however, do not inherit the labels for the classification or settings for files and emails, such as visual markings and encryption. So that users can label their documents in SharePoint sites or team sites, make sure you've enabled sensitivity labels for Office files in SharePoint and OneDrive.

NOTE

Sensitivity labels for containers aren't supported with Office 365 Content Delivery Networks (CDNs).

Using sensitivity labels for Microsoft Teams, Microsoft 365 groups, and SharePoint sites

Before you enable sensitivity labels for containers and configure sensitivity labels for the new settings, users could see and apply sensitivity labels in their apps. For example, from Word:

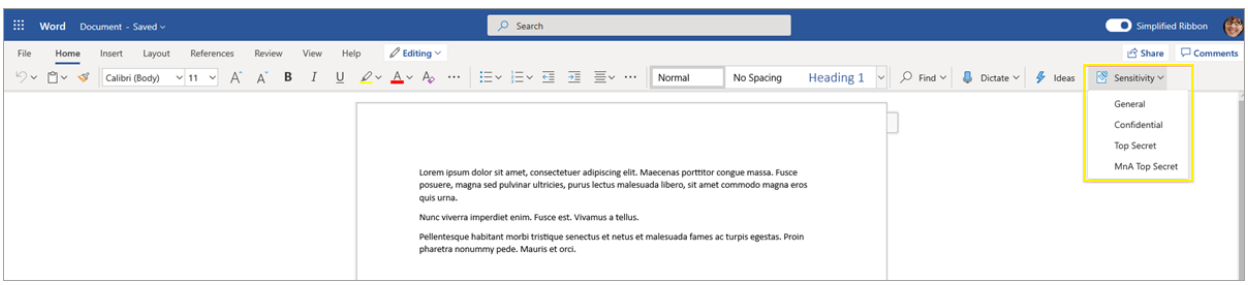

After you enable and configure sensitivity labels for containers, users can additionally see and apply sensitivity

labels to Microsoft team sites, Microsoft 365 groups, and SharePoint sites. For example, when you create a new team site from SharePoint:

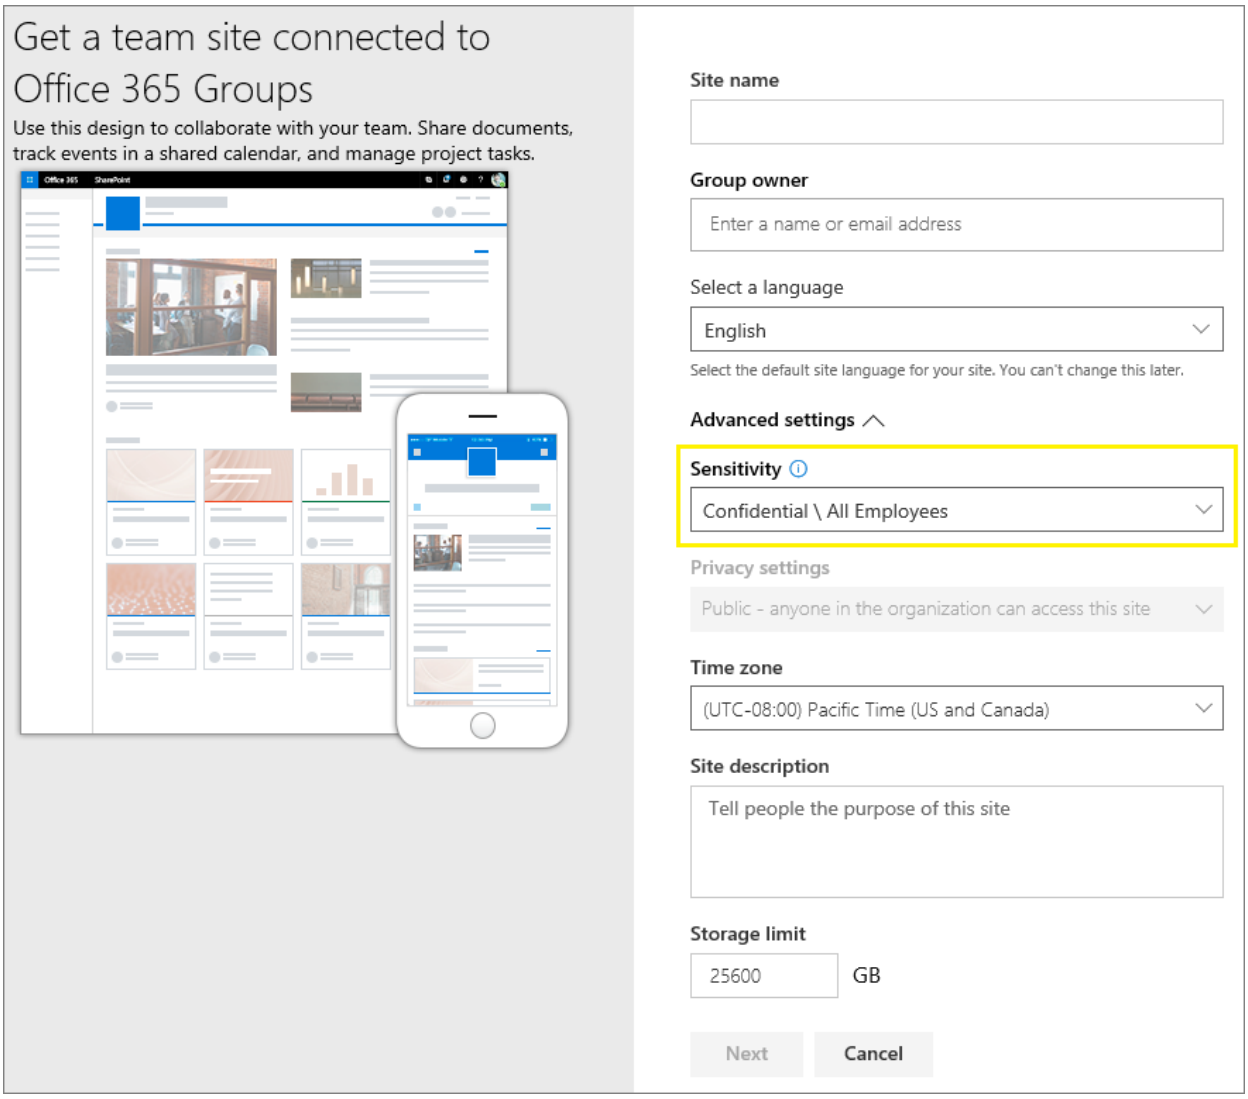

How to enable sensitivity labels for containers and synchronize labels

1. Because this feature uses Azure AD functionality, follow the instructions from the Azure AD documentation to enable sensitivity label support: Assign sensitivity labels to Microsoft 365 groups in Azure Active Directory.

2. You now need to synchronize your sensitivity labels to Azure AD. First, connect to Security & Compliance Center PowerShell.

   For example, in a PowerShell session that you run as administrator, sign in with a global administrator account.

3. Then run the following command to ensure your sensitivity labels can be used with Microsoft 365 groups:

   ```
   Execute-AzureAdLabelSync
   ```

How to configure groups and site settings

Enabling sensitivity labels for containers means that you can now configure protection settings for groups and sites in the sensitivity labeling wizard. Until you enable this support, the settings are visible in the wizard but you can't configure them.

1. Follow the general instructions to create or edit a sensitivity label and make sure you select **Groups & sites** for the label's scope:

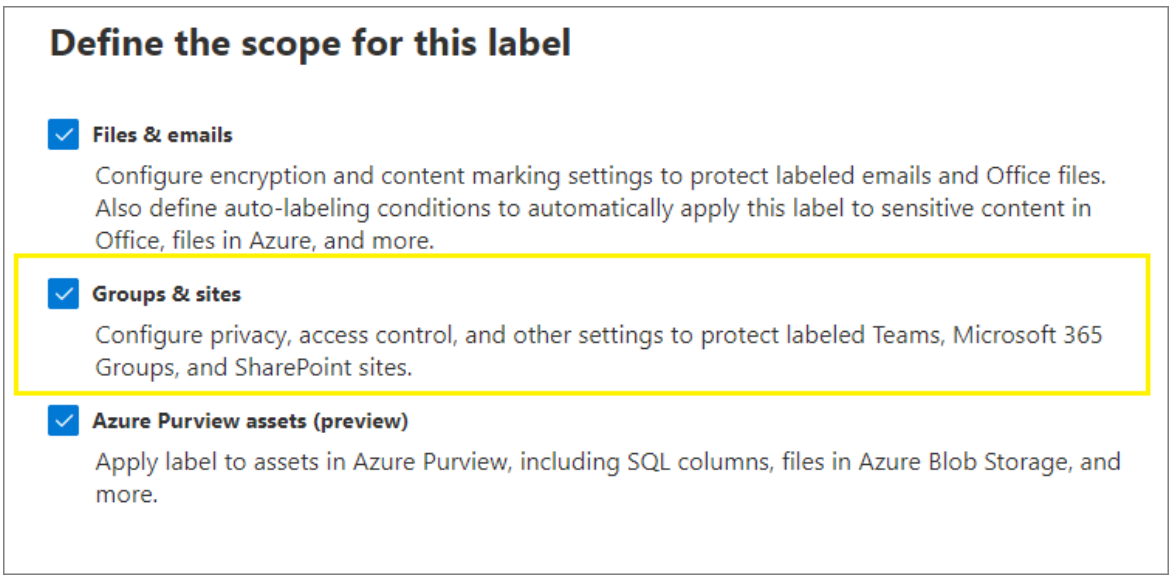

   When only this scope is selected for the label, the label won't be displayed in Office apps that support sensitivity labels and can't be applied to files and emails. Having this separation of labels can be helpful for both users and administrators, but can also add to the complexity of your label deployment.

   For example, you need to carefully review your label ordering because SharePoint detects when a labeled document is uploaded to a labeled site. In this scenario, an audit event and email are automatically generated when the document has a higher priority sensitivity label than the site's label. For more information, see the Auditing sensitivity label activities section on this page.

2. Then, on the **Define protection settings for groups and sites** page, select one or both of the available options:

   - **Privacy and external user access settings** to configure the **Privacy** and **External users access** settings.
   - **Device access and external sharing settings** to configure the **Control external sharing from labeled SharePoint sites** and **Access from unmanaged devices** setting.

3. If you selected **Privacy and external user access settings**, now configure the following settings:

   - **Privacy**: Keep the default of **Public** if you want anyone in your organization to access the team site or group where this label is applied.

     Select **Private** if you want access to be restricted to only approved members in your organization.

     Select **None** when you want to protect content in the container by using the sensitivity label, but still let users configure the privacy setting themselves.

     The settings of **Public** or **Private** set and lock the privacy setting when you apply this label to the container. Your chosen setting replaces any previous privacy setting that might be configured for the team or group, and locks the privacy value so it can be changed only by first removing the sensitivity label from the container. After you remove the sensitivity label, the privacy setting from the label remains and users can now change it again.

   - **External user access**: Control whether the group owner can add guests to the group.

4. If you selected **Device access and external sharing setting**, now configure the following settings:

   - **Control external sharing from labeled SharePoint sites**: Currently in preview, select this option to then select either external sharing for anyone, new and existing guests, existing guests,

or only people in your organization. For more information about this configuration and settings, see the SharePoint documentation, Turn external sharing on or off for a site.

- **Access from unmanaged devices**: This option uses the SharePoint feature that uses Azure AD conditional access to block or limit access to SharePoint and OneDrive content from unmanaged devices. For more information, see Control access from unmanaged devices from the SharePoint documentation. The option you specify for this label setting is the equivalent of running a PowerShell command for a site, as described in steps 3-5 from the Block or limit access to a specific SharePoint site or OneDrive section from the SharePoint instructions.

  For additional information, see More information about the dependencies for the unmanaged devices option at the end of this section.

> **IMPORTANT**
>
> Only these site and group settings take effect when you apply the label to a team, group, or site. If the label's scope includes files and emails, other label settings such as encryption and content marking aren't applied to the content within the team, group, or site.

If your sensitivity label isn't already published, now publish it by adding it to a sensitivity label policy. The users who are assigned a sensitivity label policy that includes this label will be able to select it for sites and groups.

More information about the dependencies for the unmanaged devices option

If you don't configure the dependent conditional access policy for SharePoint as documented in Use app-enforced restrictions, the option you specify here will have no effect. Additionally, it will have no effect if it's less restrictive than a configured setting at the tenant level. If you have configured an organization-wide setting for unmanaged devices, choose a label setting that's either the same or more restrictive

For example, if your tenant is configured for **Allow limited, web-only access**, the label setting that allows full access will have no effect because it's less restrictive. For this tenant-level setting, choose the label setting to block access (more restrictive) or the label setting for limited access (the same as the tenant setting).

Because you can configure the SharePoint settings separately from the label configuration, there's no check in the sensitivity label wizard that the dependencies are in place. These dependencies can be configured after the label is created and published, and even after the label is applied. However, if the label is already applied, the label setting won't take effect until after the user next authenticates.

Sensitivity label management

Use the following guidance for when you create, modify, or delete sensitivity labels that are configured for sites and groups.

Creating and publishing labels that are configured for sites and groups

When a new sensitivity label is created and published, it's visible for users in teams, groups, and sites within one hour. However, if you modify an existing label, allow up to 24 hours. Use the following guidance to publish a label for your users when that label is configured for site and group settings:

1. After you create and configure the sensitivity label, add this label to a label policy that applies to just a few test users.
2. Wait for the change to replicate:
   - New label: Wait for one hour.
   - Existing label: Wait for 24 hours.
3. After this wait period, use one of the test user accounts to create a team, Microsoft 365 group, or SharePoint site with the label that you created in step 1.

4. If there are no errors during this creation operation, you know it's safe to publish the label to all users in your tenant.

Modifying published labels that are configured for sites and groups

As a best practice, don't change the site and group settings for a sensitivity label after the label has been applied to teams, groups, or sites. If you do, remember to wait for 24 hours for the changes to replicate to all containers that have the label applied.

In addition, if your changes include the **External users access** setting:

- The new setting applies to new users but not to existing users. For example, if this setting was previously selected and as a result, guest users accessed the site, these guest users can still access the site after this setting is cleared in the label configuration.
- The privacy settings for the group properties hiddenMembership and roleEnabled aren't updated.

Deleting published labels that are configured for sites and groups

If you delete a sensitivity label that has the site and group settings enabled, and that label is included in one or more label policies, this action can result in creation failures for new teams, groups, and sites. To avoid this situation, use the following guidance:

1. Remove the sensitivity label from all label policies that include the label.
2. Wait for one hour.
3. After this wait period, try creating a team, group, or site and confirm that the label is no longer visible.
4. If the sensitivity label isn't visible, you can now safely delete the label.

How to apply sensitivity labels to containers

You're now ready to apply the sensitivity label or labels to the following containers:

- Microsoft 365 group in Azure AD
- Microsoft Teams team site
- Microsoft 365 group in Outlook on the web
- SharePoint site

You can use PowerShell if you need to apply a sensitivity label to multiple sites.

Apply sensitivity labels to Microsoft 365 groups

You're now ready to apply the sensitivity label or labels to Microsoft 365 groups. Return to the Azure AD documentation for instructions:

- Assign a label to a new group in Azure portal
- Assign a label to an existing group in Azure portal
- Remove a label from an existing group in Azure portal.

Apply a sensitivity label to a new team

Users can select sensitivity labels when they create new teams in Microsoft Teams. When they select the label from the **Sensitivity** dropdown, the privacy setting might change to reflect the label configuration. Depending on the external users access setting you selected for the label, users can or can't add people outside the organization to the team.

Learn more about sensitivity labels for Teams

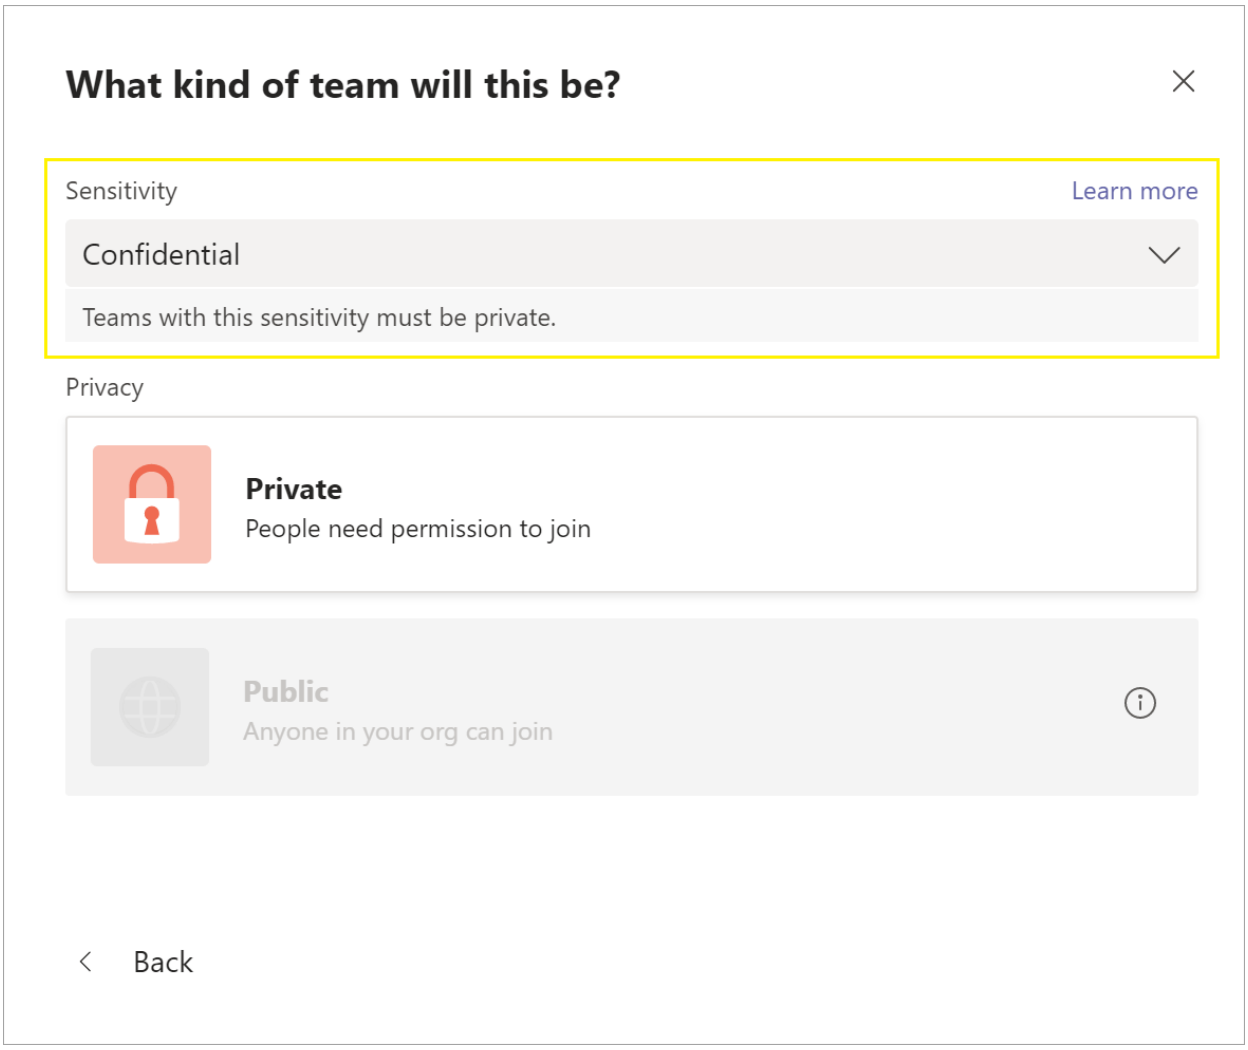

After you create the team, the sensitivity label appears in the upper-right corner of all channels.

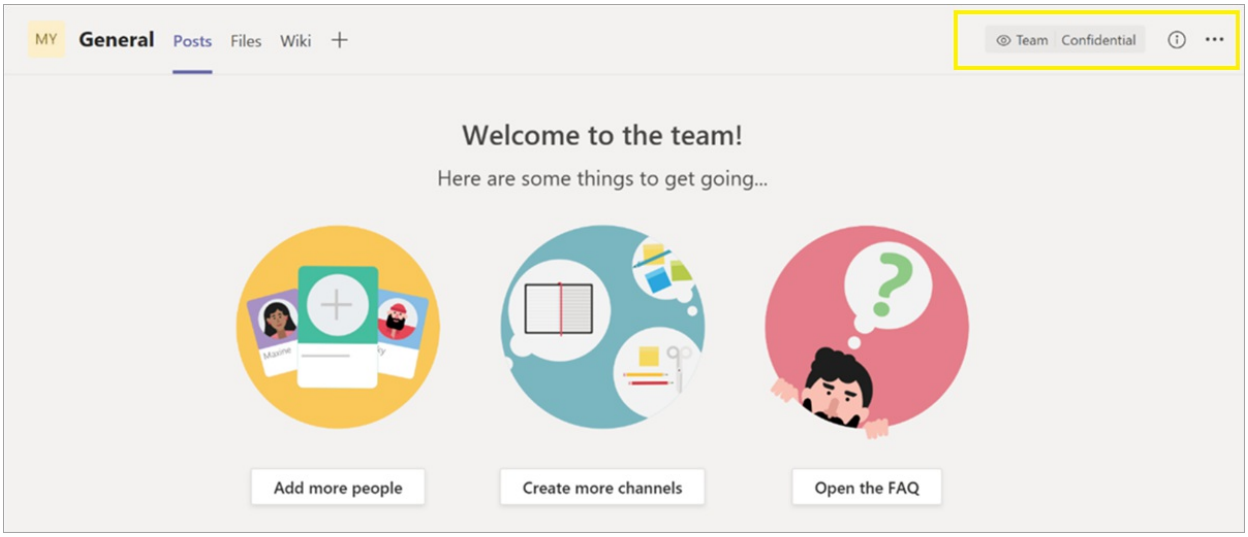

The service automatically applies the same sensitivity label to the Microsoft 365 group and the connected SharePoint team site.

Apply a sensitivity label to a new group in Outlook on the web

In Outlook on the web, when you create a new group, you can select or change the **Sensitivity** option for published labels:

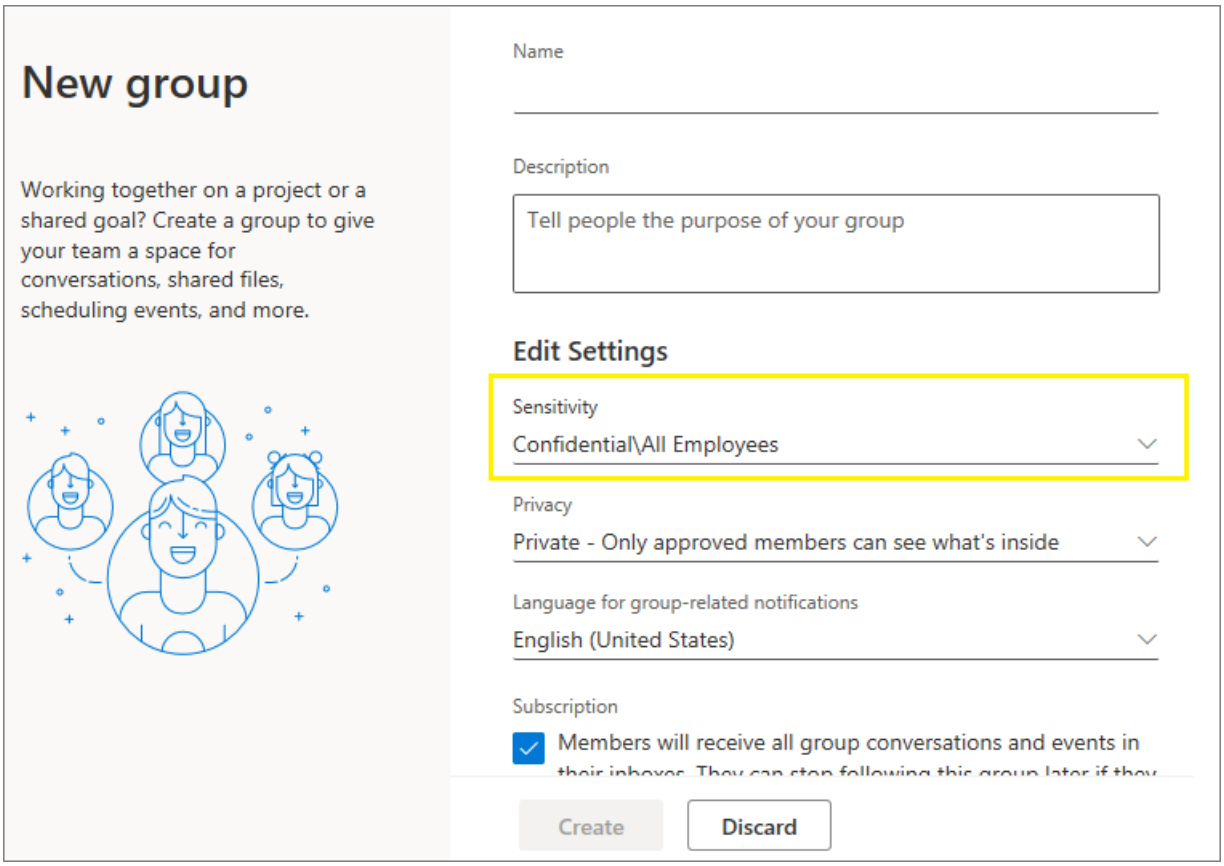

Apply a sensitivity label to a new site

Admins and end users can select sensitivity labels when they create modern team sites and communication sites, and expand **Advanced settings**:

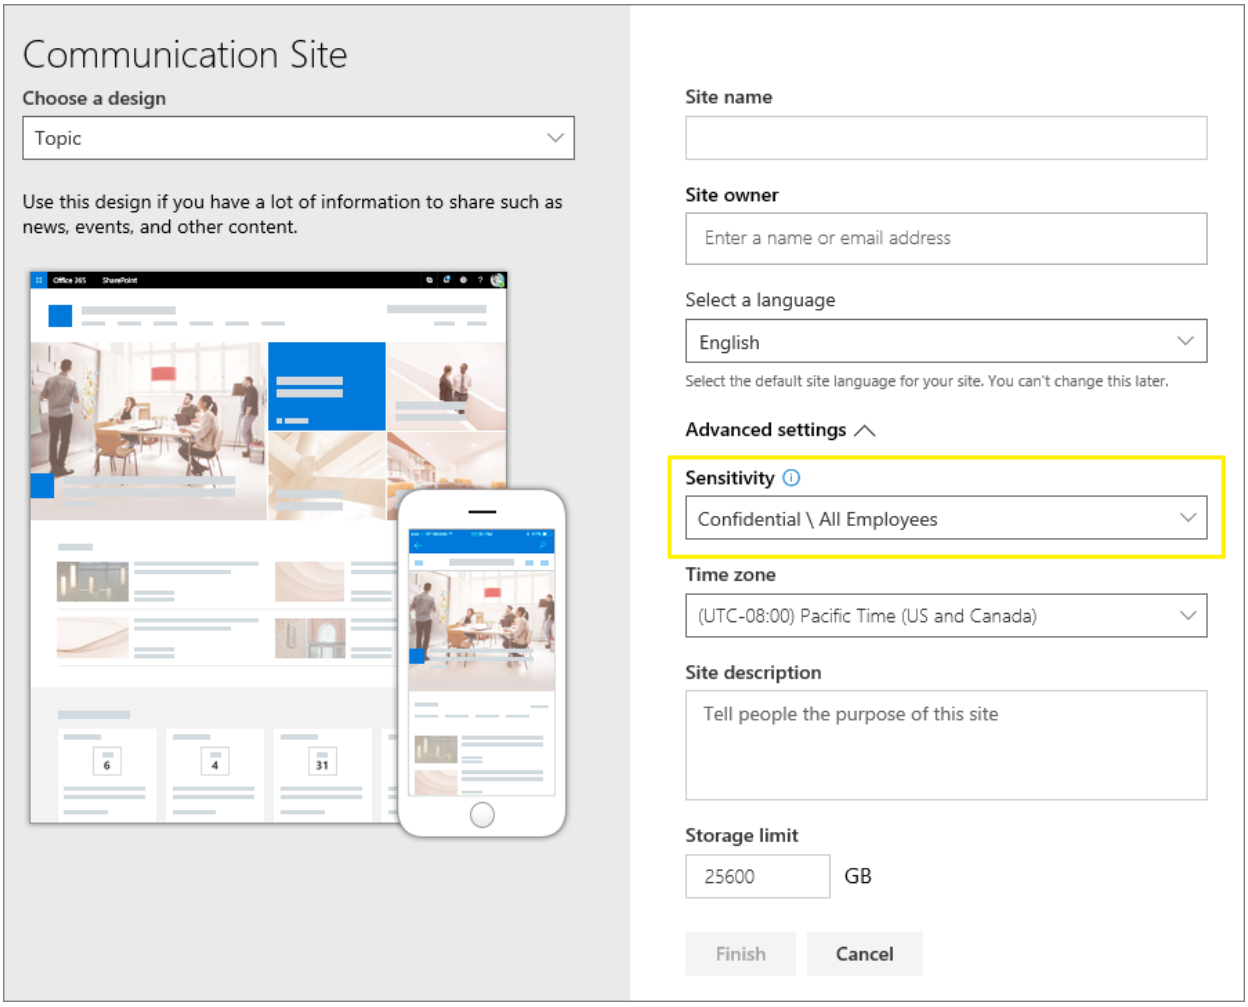

The dropdown box displays the label names for the selection, and the help icon displays all the label names with their tooltip, which can help users determine the correct label to apply.

When the label is applied, and users browse to the site, they see the name of the label and applied policies. For example, this site has been labeled as **Confidential**, and the privacy setting is set to **Private**:

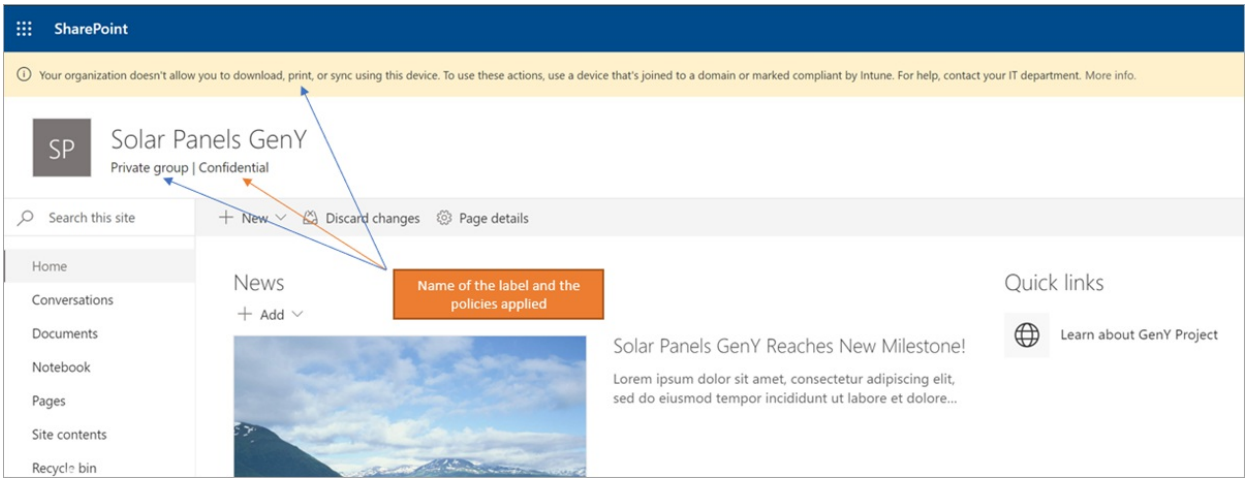

Use PowerShell to apply a sensitivity label to multiple sites

You can use the Set-SPOSite and Set-SPOTenant cmdlet with the *SensitivityLabel* parameter from the current SharePoint Online Management Shell to apply a sensitivity label to many sites. The sites can be any SharePoint site collection, or a OneDrive site.

Make sure you have version 16.0.19418.12000 or later of the SharePoint Online Management Shell.

1. Open a PowerShell session with the **Run as Administrator** option.
2. If you don't know your label GUID: Connect to Security & Compliance Center PowerShell and get the list of sensitivity labels and their GUIDs.

   ```
   Get-Label |ft Name, Guid
   ```

3. Now connect to SharePoint Online PowerShell and store your label GUID as a variable. For example:

   ```
   $Id = [GUID]("e48058ea-98e8-4940-8db0-ba1310fd955e")
   ```

4. Create a new variable that identifies multiple sites that have an identifying string in common in their URL. For example:

   ```
   $sites = Get-SPOSite -IncludePersonalSite $true -Limit all -Filter "Url -like 'documents"
   ```

5. Run the following command to apply the label to these sites. Using our examples:

   ```
   $sites | ForEach-Object {Set-SPOTenant $_.url -SensitivityLabel $Id}
   ```

To apply different labels to different sites, repeat the following command for each site:

```
Set-SPOSite -Identity <URL> -SensitivityLabel "<labelguid>"
```

View and manage sensitivity labels in the SharePoint admin center

To view, sort, and search the applied sensitivity labels, use the **Active sites** page in the new SharePoint admin center. You might need to first add the **Sensitivity** column:

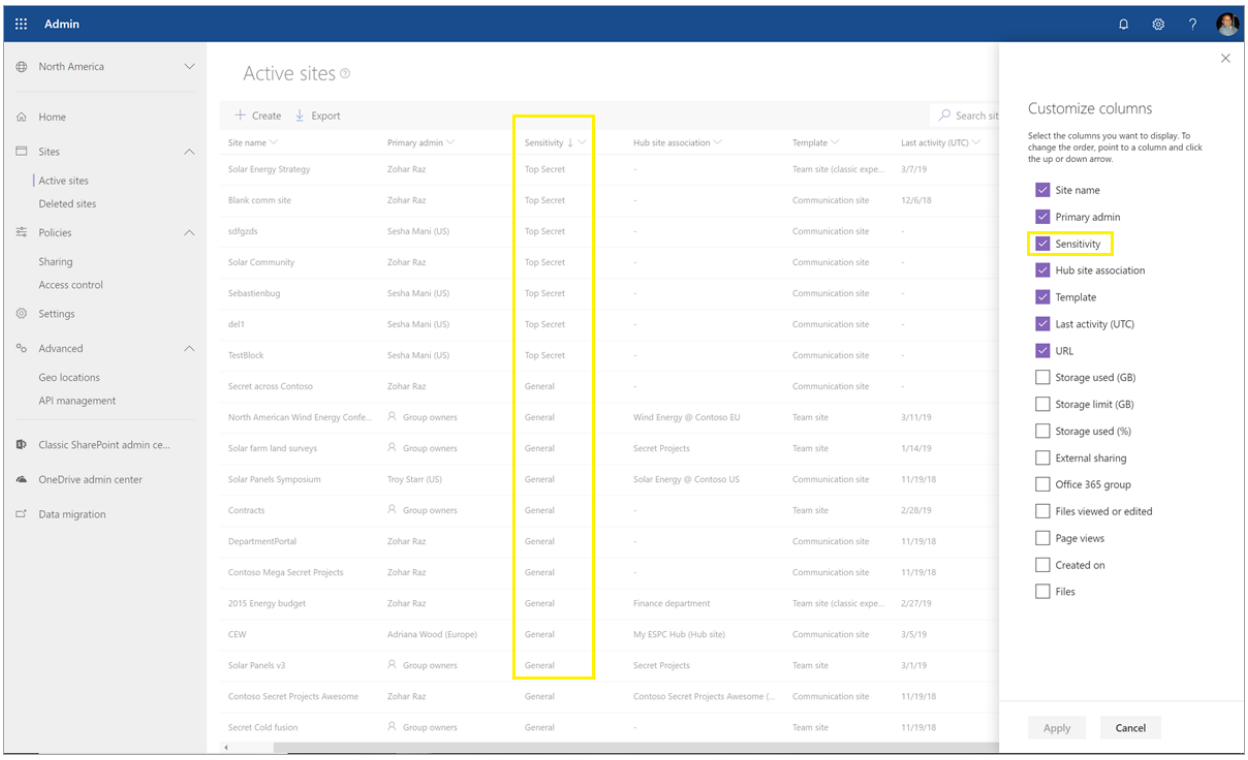

For more information about managing sites from the Active sites page, including how to add a column, see Manage sites in the new SharePoint admin center.

You can also change and apply a label from this page:

1. Select the site name to open the details pane.
2. Select the **Policies** tab, and then select **Edit** for the **Sensitivity** setting.
3. From the **Edit sensitivity setting** pane, select the sensitivity label you want to apply to the site, and then select **Save**.

Support for sensitivity labels

The following apps and services support sensitivity labels configured for sites and group settings:

- Admin centers:
  - SharePoint admin center
  - Azure Active Directory portal
  - Microsoft 365 admin center
  - Microsoft 365 compliance center, Microsoft 365 security center, Security & Compliance Center
- User apps and services:
  - SharePoint
  - Teams
  - Outlook on the web and for Windows, macOS, iOS, and Android
  - Forms
  - Stream
  - Planner

The following apps and services don't currently support sensitivity labels configured for sites and group settings:

- Admin centers:

    - Teams admin center
    - Exchange admin center
- User apps and services:
    - Dynamics 365
    - Yammer
    - Project
    - Power BI

Classic Azure AD group classification

Microsoft 365 no longer supports the old classifications for new Microsoft 365 groups and SharePoint sites after you enable sensitivity labels for containers. However, existing groups and sites that support sensitivity labels still display the old classification values until you convert them to use sensitivity labels.

As an example of how you might have used the old group classification for SharePoint, see SharePoint "modern" sites classification.

These classifications were configured by using Azure AD PowerShell or the PnP Core library and defining values for the `ClassificationList` setting. If your tenant has classification values defined, they are shown when you run the following command from the AzureADPreview PowerShell module:

```
($setting["ClassificationList"])
```

To convert your old classifications to sensitivity labels, do one of the following:

- Use existing labels: Specify the label settings you want for sites and groups by editing existing sensitivity labels that are already published.
- Create new labels: Specify the label settings you want for sites and groups by creating and publishing new sensitivity labels that have the same names as your existing classifications.

Then:

1. Use PowerShell to apply the sensitivity labels to existing Microsoft 365 groups and SharePoint sites by using name mapping. See the next section for instructions.
2. Remove the old classifications from the existing groups and sites.

Although you can't prevent users from creating new groups in apps and services that don't yet support sensitivity labels, you can run a recurring PowerShell script to look for new groups that users have created with the old classifications, and convert these to use sensitivity labels.

To help you manage the coexistence of sensitivity labels and Azure AD classifications for sites and groups, see Azure Active Directory classification and sensitivity labels for Microsoft 365 groups.

Use PowerShell to convert classifications for Microsoft 365 groups to sensitivity labels

1. First, connect to Security & Compliance Center PowerShell.

   For example, in a PowerShell session that you run as administrator, sign in with a global administrator account:

2. Get the list of sensitivity labels and their GUIDs by using the Get-Label cmdlet:

   ```
   Get-Label |ft Name, Guid
   ```

3. Make a note of the GUIDs for the sensitivity labels you want to apply to your Microsoft 365 groups.
4. Now connect to Exchange Online PowerShell in a separate Windows PowerShell window.
5. Use the following command as an example to get the list of groups that currently have the classification of "General":

```
$Groups= Get-UnifiedGroup | Where {$_.classification -eq "General"}
```

6. For each group, add the new sensitivity label GUID. For example:

```
foreach ($g in $groups)
{Set-UnifiedGroup -Identity $g.Identity -SensitivityLabelId "457fa763-7c59-461c-b402-ad1ac6b703cc"}
```

7. Repeat steps 5 and 6 for your remaining group classifications.

Auditing sensitivity label activities

> **IMPORTANT**
> If you use label separation by selecting just the **Groups & sites** scope for labels that protect containers: Because of the **Detected document sensitivity mismatch** audit event and email described in this section, consider ordering these labels before labels that have a scope for **Files & emails**.

If somebody uploads a document to a site that's protected with a sensitivity label and their document has a higher priority sensitivity label than the sensitivity label applied to the site, this action isn't blocked. For example, you've applied the **General** label to a SharePoint site, and somebody uploads to this site a document labeled **Confidential**. Because a sensitivity label with a higher priority identifies content that is more sensitivity than content that has a lower priority order, this situation could be a security concern.

Although the action isn't blocked, it is audited and automatically generates an email to the person who uploaded the document and the site administrator. As a result, both the user and administrators can identify documents that have this misalignment of label priority and take action if needed. For example, delete or move the uploaded document from the site.

It wouldn't be a security concern if the document has a lower priority sensitivity label than the sensitivity label applied to the site. For example, a document labeled **General** is uploaded to a site labeled **Confidential**. In this scenario, an auditing event and email aren't generated.

To search the audit log for this event, look for **Detected document sensitivity mismatch** from the **File and page activities** category.

The automatically generated email has the subject **Incompatible sensitivity label detected** and the email message explains the labeling mismatch with a link to the uploaded document and site. It also contains a documentation link that explains how users can change the sensitivity label. Currently, these automated emails cannot be disabled or customized.

When somebody adds or removes a sensitivity label to or from a site or group, these activities are also audited but without automatically generating an email.

All these auditing events can be found in the Sensitivity label activities category. For instructions to search the audit log, see Search the audit log in the Security & Compliance Center.

How to disable sensitivity labels for containers

You can turn off sensitivity labels for Microsoft Teams, Microsoft 365 groups, and SharePoint sites by using the same instructions from Enable sensitivity label support in PowerShell. However, to disable the feature, in step 5, specify `$setting["EnableMIPLabels"] = "False"`.

In addition to making all the settings unavailable for groups and sites when you create or edit sensitivity labels, this action reverts which property the containers use for their configuration. Enabling sensitivity labels for Microsoft Teams, Microsoft 365 groups, and SharePoint sites switches the property used from **Classification** (used for Azure AD group classification) to **Sensitivity**. When you disable sensitivity labels for containers, the containers ignore the Sensitivity property and use the Classification property again.

This means that any label settings from sites and groups previously applied to containers won't be enforced, and containers no longer display the labels.

If these containers have Azure AD classification values applied to them, the containers revert to using the classifications again. Be aware that any new sites or groups that were created after enabling the feature won't display a label or have a classification. For these containers, and any new containers, you can now apply classification values. For more information, see SharePoint "modern" sites classification and Create classifications for Office groups in your organization.

Additional resources

See the webinar recording and answered questions for Using Sensitivity labels with Microsoft Teams, O365 Groups and SharePoint Online sites.

This webinar was recorded when the feature was still in preview, so you might notice some discrepancies in the UI. However, the information for this feature is still accurate, with any new capabilities documented on this page.

Enable sensitivity labels for Office files in SharePoint and OneDrive

2/18/2021 • 15 minutes to read • Edit Online

Microsoft 365 licensing guidance for security & compliance.

Enable sensitivity labels for Office files in SharePoint and OneDrive so that users can apply your sensitivity labels in Office on the web. When this feature is enabled, users will see the **Sensitivity** button on the ribbon so they can apply labels, and see any applied label name on the status bar.

Enabling this feature also results in SharePoint and OneDrive being able to process the contents of files that have been encrypted by using a sensitivity label. The label can be applied in Office for the web, or in Office desktop apps and uploaded or saved in SharePoint and OneDrive. Until you enable this feature, these services can't process encrypted files, which means that coauthoring, eDiscovery, Data Loss Prevention, search, and other collaborative features won't work for these files.

After you enable sensitivity labels for Office files in SharePoint and OneDrive, for new and changed files that have a sensitivity label that applies encryption with a cloud-based key (and doesn't use Double Key Encryption):

- For Word, Excel, and PowerPoint files, SharePoint and OneDrive recognize the label and can now process the contents of the encrypted file.
- When users download or access these files from SharePoint or OneDrive, the sensitivity label and any encryption settings from the label are enforced and remain with the file, wherever it is stored. Ensure you provide user guidance to use only labels to protect documents. For more information, see Information Rights Management (IRM) options and sensitivity labels.
- When users upload labeled and encrypted files to SharePoint or OneDrive, they must have at least view rights to those files. For example, they can open the files outside SharePoint. If they don't have this minimum usage right, the upload is successful but the service doesn't recognize the label and can't process the file contents.
- Use Office on the web (Word, Excel, PowerPoint) to open and edit Office files that have sensitivity labels that apply encryption. The permissions that were assigned with the encryption are enforced. You can also use auto-labeling for these documents.
- External users can access documents that are labeled with encryption by using guest accounts. For more information, see Support for external users and labeled content.
- Office 365 eDiscovery supports full-text search for these files and Data Loss Prevention (DLP) policies support content in these files.

NOTE

If encryption has been applied with an on-premises key (a key management topology often referred to as "hold your own key" or HYOK), or by using Double Key Encryption, the service behavior for processing the file contents doesn't change. So for these files, coauthoring, eDiscovery, Data Loss Prevention, search, and other collaborative features won't work.

The SharePoint and OneDrive behavior also doesn't change for existing files in these locations that are labeled with encryption using a single Azure-based key. For these files to benefit from the new capabilities after you enable sensitivity labels for Office files in SharePoint and OneDrive, the files must be either downloaded and uploaded again, or edited.

After you enable sensitivity labels for Office files in SharePoint and OneDrive, three new audit events are available for monitoring sensitivity labels that are applied to documents in SharePoint and OneDrive:

- **Applied sensitivity label to file**
- **Changed sensitivity label applied to file**
- **Removed sensitivity label from file**

Watch the following video (no audio) to see the new capabilities in action:

You always have the choice to disable sensitivity labels for Office files in SharePoint and OneDrive (opt-out) at any time.

If you are currently protecting documents in SharePoint by using SharePoint Information Rights Management (IRM), be sure to check the SharePoint Information Rights Management (IRM) and sensitivity labels section on this page.

Requirements

These new capabilities work with sensitivity labels only. If you currently have Azure Information Protection labels, first migrate them to sensitivity labels so that you can enable these features for new files that you upload. For instructions, see How to migrate Azure Information Protection labels to unified sensitivity labels.

Use the OneDrive sync app version 19.002.0121.0008 or later on Windows, and version 19.002.0107.0008 or later on Mac. Both these versions were released January 28, 2019, and are currently released to all rings. For more information, see the OneDrive release notes. After you enable sensitivity labels for Office files in SharePoint and OneDrive, users who run an older version of the sync app are prompted to update it.

Limitations

- SharePoint and OneDrive don't automatically apply sensitivity labels to existing files that you've already encrypted using Azure Information Protection labels. Instead, for the features to work after you enable sensitivity labels for Office files in SharePoint and OneDrive, complete these tasks:

  1. Make sure you have migrated the Azure Information Protection labels to sensitivity labels and published them from the Microsoft 365 compliance center, or equivalent labeling admin center.
  2. Download the files and then upload them to SharePoint.

- SharePoint and OneDrive can't process encrypted files when the label that applied the encryption has any of the following configurations for encryption:

  - **Let users assign permissions when they apply the label** and the checkbox **In Word, PowerPoint, and Excel, prompt users to specify permissions** is selected. This setting is sometimes referred to as "user-defined permissions".
  - **User access to content expires** is set to a value other than **Never**.
  - **Double Key Encryption** is selected.

  For labels with any of these encryption configurations, the labels aren't displayed to users in Office on the web. Additionally, the new capabilities can't be used with labeled documents that already have these encryption settings. For example, these documents won't be returned in search results, even if they are updated.

- For encrypted documents, printing is not supported.

- For an encrypted document that grants edit permissions to a user, copying can't be blocked in the web versions of the Office apps.

- The Azure Information Protection document tracking site is not supported.
- Office desktop apps and mobile apps don't support coauthoring for files that are labeled with encryption. These apps continue to open labeled and encrypted files in exclusive editing mode.
- If an admin changes settings for a published label that's already applied to files downloaded to users' sync client, users might be unable to save changes they make to the file in their OneDrive Sync folder. This scenario applies to files that are labeled with encryption, and also when the label change is from a label that didn't apply encryption to a label that does apply encryption. Users see a red circle with a white cross icon error, and they are asked to save new changes as a separate copy. Instead, they can close and reopen the file, or use Office on the web.
- If a labeled document is uploaded to SharePoint or OneDrive and the label applied encryption by using an account from a service principal name, the document can't be opened in Office on the web. Example scenarios include Microsoft Cloud App Security and a file sent to Teams by email.
- Users can experience save problems after going offline or into a sleep mode when instead of using Office for the web, they use the desktop and mobile apps for Word, Excel, or PowerPoint. For these users, when they resume their Office app session and try to save changes, they see an upload failure message with an option to save a copy instead of saving the original file.
- Documents that have been encrypted in the following ways can't be opened in Office on the web:
  - Encryption that uses an on-premises key ("hold your own key" or HYOK)
  - Encryption that was applied by using Double Key Encryption
  - Encryption that was applied independently from a label, for example, by directly applying a Rights Management protection template.
- Labels configured for other languages are not supported and display the original language only.
- Screen captures can't be prevented for encrypted documents. For more information, see Can Rights Management prevent screen captures?
- If you delete a label that's been applied to a document in SharePoint or OneDrive, rather than remove the label from the applicable label policy, the document when downloaded won't be labeled or encrypted. In comparison, if the labeled document is stored outside SharePoint or OneDrive, the document remains encrypted if the label is deleted. Note that although you might delete labels during a testing phase, it's very rare to delete a label in a production environment.

How to enable sensitivity labels for SharePoint and OneDrive (opt-in)

You can enable the new capabilities by using the Microsoft 365 compliance center, or by using PowerShell. As with all tenant-level configuration changes for SharePoint and OneDrive, it takes about 15 minutes for the change to take effect.

Use the compliance center to enable support for sensitivity labels

This option is the easiest way to enable sensitivity labels for SharePoint and OneDrive, but you must sign in as a global administrator for your tenant.

1. Sign in to the Microsoft 365 compliance center as a global administrator, and navigate to **Solutions** > **Information protection**

   If you don't immediately see this option, first select **Show all**.

2. If you see a message to turn on the ability to process content in Office online files, select **Turn on now**:

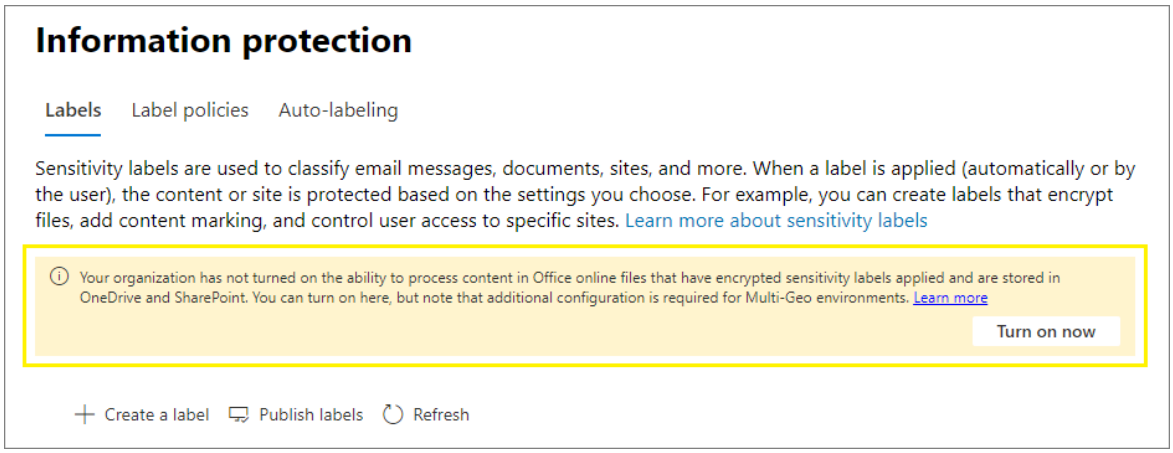

The command runs immediately and when the page is next refreshed, you no longer see the message or button.

> **NOTE**
> If you have Microsoft 365 Multi-Geo, you must use PowerShell to enable these capabilities for all your geo-locations. See the next section for details.

Use PowerShell to enable support for sensitivity labels

As an alternative to using the compliance center, you can enable support for sensitivity labels by using the Set-SPOTenant cmdlet from SharePoint Online PowerShell.

If you have Microsoft 365 Multi-Geo, you must use PowerShell to enable this support for all your geo-locations.

Prepare the SharePoint Online Management Shell

Before you run the PowerShell command to enable sensitivity labels for Office files in SharePoint and OneDrive, ensure that you're running SharePoint Online Management Shell version 16.0.19418.12000 or later. If you already have the latest version, you can skip to next procedure to run the PowerShell command.

1. If you have installed a previous version of the SharePoint Online Management Shell from PowerShell gallery, you can update the module by running the following cmdlet.

   ```
   Update-Module -Name Microsoft.Online.SharePoint.PowerShell
   ```

2. Alternatively, if you have installed a previous version of the SharePoint Online Management Shell from the Microsoft Download Center, you can also go to **Add or remove programs** and uninstall the SharePoint Online Management Shell.

3. In a web browser, go to the Download Center page and Download the latest SharePoint Online Management Shell.

4. Select your language and then click **Download**.

5. Choose between the x64 and x86 .msi file. Download the x64 file if you run the 64-bit version of Windows or the x86 file if you run the 32-bit version. If you don't know, see Which version of Windows operating system am I running?

6. After you have downloaded the file, run the file and follow the steps in the Setup Wizard.

Run the PowerShell command to enable support for sensitivity labels

To enable the new capabilities, use the Set-SPOTenant cmdlet with the *EnableAIPIntegration* parameter:

1. Using a work or school account that has global administrator or SharePoint admin privileges in Microsoft

365, connect to SharePoint. To learn how, see Getting started with SharePoint Online Management Shell.

Note: If you have Microsoft 365 Multi-Geo, use the -Url parameter with Connect-SPOService, and specify the SharePoint Online Administration Center site URL for one of your geo-locations.

2. Run the following command and press **Y** to confirm:

```
Set-SPOTenant -EnableAIPIntegration $true
```

3. For Microsoft 365 Multi-Geo: Repeat steps 1 and 2 for each of your remaining geo-locations.

Publishing and changing sensitivity labels

When you use sensitivity labels with SharePoint and OneDrive, keep in mind that you need to allow for replication time when you publish new sensitivity labels or update existing sensitivity labels. This is especially important for new labels that apply encryption.

For example: You create and publish a new sensitivity label that applies encryption and it very quickly appears in a user's desktop app. The user applies this label to a document and then uploads it to SharePoint or OneDrive. If the label replication hasn't completed for the service, the new capabilities won't be applied to that document on upload. As a result, the document won't be returned in search or for eDiscovery and the document can't be opened in Office for the web.

The following changes replicate within one hour: New and deleted sensitivity labels, and sensitivity label policy settings that include which labels are in the policy.

The following changes replicate within 24 hours: Changes to sensitivity label settings for existing labels.

Because the replication delay is only one hour for new sensitivity labels, you are unlikely to run into the scenario in the example. But as a safeguard, we recommend publishing new labels to just a few test users first, wait for an hour, and then verify the label behavior on SharePoint and OneDrive. As the final step, make the label available to more users by either adding more users to the existing label policy, or add the label to an existing label policy for your standard users. At the time your standard users see the label, it has already synchronized to SharePoint and OneDrive.

SharePoint Information Rights Management (IRM) and sensitivity labels

SharePoint Information Rights Management (IRM) is an older technology to protect files at the list and library level by applying encryption and restrictions when files are downloaded. This older protection technology is designed to prevent unauthorized users from opening the file while it's outside SharePoint.

In comparison, sensitivity labels provide the protection settings of visual markings (headers, footers, watermarks) in addition to encryption. The encryption settings support the full range of usage rights to restrict what users can do with the content, and the same sensitivity labels are supported for many scenarios. Using the same protection method with consistent settings across workloads and apps results in a consistent protection strategy.

However, you can use both protection solutions together and the behavior is as follows:

- If you upload a file with a sensitivity label that applies encryption, SharePoint can't process the content of these files so coauthoring, eDiscovery, DLP, and search are not supported for these files.
- If you label a file using Office on the web, any encryption settings from the label are enforced. For these files, coauthoring, eDiscovery, DLP, and search are supported.

- If you download a file that's labeled by using Office on the web, the label is retained and any encryption settings from the label are enforced rather than the IRM restriction settings.
- If you download an Office or PDF file that isn't encrypted with a sensitivity label, IRM settings are applied.
- If you have enabled any of the additional IRM library settings, which include preventing users from uploading documents that don't support IRM, these settings are enforced.

With this behavior, you can be assured that all Office and PDF files are protected from unauthorized access if they are downloaded, even if they aren't labeled. However, labeled files that are uploaded won't benefit from the new capabilities.

Search for documents by sensitivity label

Use the managed property **InformationProtectionLabelId** to find all documents in SharePoint or OneDrive that have a specific sensitivity label. Use the following syntax: `InformationProtectionLabelId:<GUID>`

For example, to search for all documents that have been labeled as "Confidential", and that label has a GUID of "8faca7b8-8d20-48a3-8ea2-0f96310a848e", in the search box, type:

```
InformationProtectionLabelId: 8faca7b8-8d20-48a3-8ea2-0f96310a848e
```

To get the GUIDs for your sensitivity labels, use the Get-Label cmdlet:

1. First, connect to Office 365 Security & Compliance Center PowerShell.

   For example, in a PowerShell session that you run as administrator, sign in with a global administrator account.

2. Then run the following command:

   ```
   Get-Label |ft Name, Guid
   ```

For more information about using managed properties, see Manage the search schema in SharePoint.

Remove encryption for a labeled document

There might be rare occasions when a SharePoint administrator needs to remove encryption from a document stored in SharePoint. Any user who has the Rights Management usage right of Export or Full Control assigned to them for that document can remove encryption that was applied by the Azure Rights Management service from Azure Information Protection. For example, users with either of these usage rights can replace a label that applies encryption with a label without encryption. Alternatively, a super user could download the file and save a local copy without the encryption.

As an alternative, a global admin or SharePoint admin can run the Unlock-SPOSensitivityLabelEncryptedFile cmdlet, which removes both the sensitivity label and the encryption. This cmdlet runs even if the admin doesn't have access permissions to the site or file, or if the Azure Rights Management service is unavailable.

For example:

```
Unlock-SPOSensitivityLabelEncryptedFile -FileUrl "https://contoso.com/sites/Marketing/Shared
Documents/Doc1.docx" -JustificationText "Need to decrypt this file"
```

Requirements:

- SharePoint Online Management Shell version 16.0.20616.12000 or later.

- The encryption has been applied by a sensitivity label with admin-defined encryption settings (the Assign permissions now label settings). Double Key Encryption is not supported for this cmdlet.

The justification text is added to the audit event of **Removed sensitivity label from file**, and the decryption action is also recorded in the protection usage logging for Azure Information Protection.

How to disable sensitivity labels for SharePoint and OneDrive (opt-out)

If you disable these new capabilities, files that you uploaded after you enabled sensitivity labels for SharePoint and OneDrive continue to be protected by the label because the label settings continue to be enforced. When you apply sensitivity labels to new files after you disable these new capabilities, full-text search, eDiscovery, and coauthoring will no longer work.

To disable these new capabilities, you must use PowerShell. Using the SharePoint Online Management Shell and the Set-SPOTenant cmdlet, specify the same *EnableAIPIntegration* parameter as described in the Use PowerShell to enable support for sensitivity labels section. But this time, set the parameter value to false and press **Y** to confirm:

```
Set-SPOTenant -EnableAIPIntegration $false
```

If you have Microsoft 365 Multi-Geo, you must run this command for each of your geo-locations.

Next steps

After you've enabled sensitivity labels for Office files in SharePoint and OneDrive, consider automatically labeling these files by using auto-labeling policies. For more information, see Apply a sensitivity label to content automatically.

Need to share your labeled and encrypted documents with people outside your organization? See Sharing encrypted documents with external users.

Manage sensitivity labels in Office apps

2/18/2021 • 20 minutes to read • Edit Online

> *Microsoft 365 licensing guidance for security & compliance*.

When you have published sensitivity labels from the Microsoft 365 compliance center or equivalent labeling center, they start to appear in Office apps for users to classify and protect data as it's created or edited.

Use the information in this article to help you successfully manage sensitivity labels in Office apps. For example, identify the minimum versions of apps you need to support built-in labeling, and understand interactions with the Azure Information Protection unified labeling client and compatibility with other apps and services.

Labeling client for desktop apps

To use sensitivity labels that are built into Office desktop apps for Windows and Mac, you must use a subscription edition of Office. This labeling client doesn't support standalone editions of Office, such as Office 2016 or Office 2019.

To use sensitivity labels with these standalone editions of Office on Windows computers, install the Azure Information Protection unified labeling client.

Support for sensitivity label capabilities in apps

For each capability, the following tables list the minimum Office version that you need to support sensitivity labels using built-in labeling. Or, if the label capability is in public preview or under review for a future release. Use the Microsoft 365 roadmap for details about future releases.

New versions of Office apps are made available at different times for different update channels. For more information, including how to configure your update channel so that you can test a new labeling capability that you're interested in, see Overview of update channels for Microsoft 365 Apps. New capabilities that are in private preview are not included in the table but you might be able to join these previews by nominating your organization for the Microsoft Information Protection private preview program.

> **NOTE**
> The names of the update channels for Office apps have recently changed. For example, Monthly Channel is now Current Channel, and Office Insider is now Beta Channel. For more information, see Changes to update channels for Microsoft 365 Apps.

Office for iOS and Office for Android: Sensitivity labels are built into the Office app.

Additional capabilities are available when you install the Azure Information Protection unified labeling client, which runs on Windows computers only. For these details, see Compare the labeling clients for Windows computers.

Sensitivity label capabilities in Word, Excel, and PowerPoint

The numbers listed are the minimum Office application version required for each capability.

CAPABILITY	WINDOWS	MAC	IOS	ANDROID	WEB
Manually apply, change, or remove label	1910+	16.21+	2.21+	16.0.11231+	Yes - opt-in
Apply a default label	1910+	16.21+	2.21+	16.0.11231+	Yes - opt-in
Require a justification to change a label	1910+	16.21+	2.21+	16.0.11231+	Yes - opt-in
Provide help link to a custom help page	1910+	16.21+	2.21+	16.0.11231+	Yes - opt-in
Mark the content	1910+	16.21+	2.21+	16.0.11231+	Yes - opt-in
Dynamic markings with variables	2010+	16.42+	2.42+	16.0.13328+	Under review
Assign permissions now	1910+	16.21+	2.21+	16.0.11231+	Yes - opt-in
Let users assign permissions	2004+	16.35+	Under review	Under review	Under review
Get started with data classification and send data for administrators	2011+	16.43+	Preview: Current Channel (Preview)	Preview: Current Channel (Preview)	Yes *
Require users to apply a label to their email and documents	Preview: Current Channel (Preview)	Preview: Current Channel (Preview)	Under review	Rolling out: 16.0.13628+	Under review
Apply a sensitivity label to content automatically	2009+	Rolling out: 16.44+	Under review	Under review	Yes - opt-in
Support AutoSave and coauthoring on labeled and encrypted documents	Under review	Under review	Under review	Under review	Yes - opt-in

Footnote:

* Currently, doesn't include justification text to remove a label or lower the classification level

Sensitivity label capabilities in Outlook

The numbers listed are the minimum Office application version required for each capability.

CAPABILITY	OUTLOOK FOR WINDOWS	OUTLOOK FOR MAC	OUTLOOK ON IOS	OUTLOOK ON ANDROID	OUTLOOK ON THE WEB
Manually apply, change, or remove label	1910+	16.21+	4.7.1+	4.0.39+	Yes
Apply a default label	1910+	16.21+	4.7.1+	4.0.39+	Yes
Require a justification to change a label	1910+	16.21+	4.7.1+	4.0.39+	Yes
Provide help link to a custom help page	1910+	16.21+	4.7.1+	4.0.39+	Yes
Mark the content	1910+	16.21+	4.7.1+	4.0.39+	Yes
Dynamic markings with variables	Under review	Under review	Under review	Under review	Under review
Assign permissions now	1910+	16.21+	4.7.1+	4.0.39+	Yes
Let users assign permissions	1910+	16.21+	4.7.1+	4.0.39+	Yes
Require users to apply a label to their email and documents	Preview: Current Channel (Preview))	16.43+	Under review	Under review	Yes
Get started with data classification and send data for administrators	2011+	Under review	Under review	Under review	Under review
Apply a sensitivity label to content automatically	2009+	16.44+	Under review	Under review	Yes

Office built-in labeling client and other labeling solutions

The Office built-in labeling client downloads sensitivity labels and sensitivity label policy settings from the following admin centers:

- Microsoft 365 compliance center

- Microsoft 365 security center
- Office 365 Security & Compliance Center

To use the Office built-in labeling client, you must have one or more label policies published to users from one of the listed admin centers and a supported version of Office.

If both of these conditions are met but you need to turn off the Office built-in labeling client, use the following Group Policy setting:

1. Navigate to **User Configuration/Administrative Templates/Microsoft Office 2016/Security Settings**.
2. Set **Use the Sensitivity feature in Office to apply and view sensitivity labels** to **0**.

Deploy this setting by using Group Policy, or by using the Office cloud policy service. The setting takes effect when Office apps restart.

Office built-in labeling client and the Azure Information Protection client

If users have one of the Azure Information Protection clients installed (unified labeling client or classic client), by default, the built-in labeling client is turned off in their Office apps.

To use built-in labeling rather than the Azure Information Protection client for Office apps, use the instructions from the previous section but set the Group Policy setting **Use the Sensitivity feature in Office to apply and view sensitivity labels** to **1**.

Alternatively, disable or remove the Office Add-in, **Azure Information Protection**. This method is suitable for a single computer, and ad-hoc testing. For instructions, see View, manage, and install add-ins in Office programs.

When you disable or remove this Office Add-in, the Azure Information Protection client remains installed so that you can continue to label files outside your Office apps. For example, by using File Explorer, or PowerShell.

For information about which features are supported by the Azure Information Protection clients and the Office built-in labeling client, see Choose which labeling client to use for Windows computers from the Azure Information Protection documentation.

Office file types supported

Office apps that have built-in labeling for Word, Excel, and PowerPoint files support the Open XML format (such as .docx and .xlsx) but not the Microsoft Office 97-2003 format (such as .doc and .xls), Open Document Format (such as .odt and .ods), or other formats. When a file type is not supported for built-in labeling, the **Sensitivity** button is not available in the Office app.

The Azure Information Protection unified labeling client supports both the Open XML format and the Microsoft Office 97-2003 format. For more information, see File types supported by the Azure Information Protection unified labeling client from that client's admin guide.

For other labeling solutions, check their documentation for file types supported.

Protection templates and sensitivity labels

Administrator-defined protection templates, such as those you define for Office 365 Message Encryption, aren't visible in Office apps when you're using built-in labeling. This simplified experience reflects that there's no need to select a protection template, because the same settings are included with sensitivity labels that have encryption enabled.

If you need to convert existing protection templates to labels, use the Azure portal and the following instructions: To convert templates to labels.

Information Rights Management (IRM) options and sensitivity labels

Sensitivity labels that you configure to apply encryption remove the complexity from users to specify their own encryption settings. In many Office apps, these individual encryption settings can still be manually configured by users by using Information Rights Management (IRM) options. For example, for Windows apps:

- For a document: **File** > **Info** > **Protect Document** > **Restrict Access**
- for an email: From the **Options** tab > **Encrypt**

When users initially label a document or email, they can always override your label configuration settings with their own encryption settings. For example:

- A user applies the **Confidential \ All Employees** label to a document and this label is configured to apply encryption settings for all users in the organization. This user then manually configures the IRM settings to restrict access to a user outside your organization. The end result is a document that's labeled **Confidential \ All Employees** and encrypted, but users in your organization can't open it as expected.
- A user applies the **Confidential \ Recipients Only** label to an email and this email is configured to apply the encryption setting of **Do Not Forward**. This user then manually configures the IRM settings so that the email is unrestricted. The end result is the email can be forwarded by recipients, despite having the **Confidential \ Recipients Only** label.
- A user applies the **General** label to a document, and this label isn't configured to apply encryption. This user then manually configures the IRM settings to restrict access to the document. The end result is a document that's labeled **General** but that also applies encryption so that some users can't open it as expected.

If the document or email is already labeled, a user can do any of these actions if the content isn't already encrypted, or they have the usage right Export or Full Control.

For a more consistent label experience with meaningful reporting, provide appropriate labels and guidance for users to apply only labels to protect documents. For example:

- For exception cases where users must assign their own permissions, provide labels that let users assign their own permissions.
- Instead of users manually removing encryption after selecting a label that applies encryption, provide a sublabel alternative when users need a label with the same classification, but no encryption. Such as:
  - **Confidential \ All Employees**
  - **Confidential \ Anyone (no encryption)**

> **NOTE**
> If users manually remove encryption from a labeled document that's stored in SharePoint or OneDrive and you've enabled sensitivity labels for Office files in SharePoint and OneDrive, the label encryption will be automatically restored the next time the document is accessed or downloaded.

Apply sensitivity labels to files, emails, and attachments

Users can apply just one label at a time for each document or email.

When you label an email message that has attachments, the attachments inherit the label only if the label that you apply to the email message applies encryption and the attachment is an Office document isn't already encrypted. Because the inherited label applies encryption, the attachment becomes newly encrypted.

An attachment doesn't inherit the labels from the email message when the label applied to the email message

doesn't apply encryption or the attachment is already encrypted.

Examples of label inheritance, where the label **Confidential** applies encryption and the label **General** doesn't apply encryption:

- A user creates a new email message and applies the **Confidential** label to this message. They then add a Word document that isn't labeled or encrypted. As a result of inheritance, the document is newly labeled **Confidential** and now has encryption applied from that label.
- A user creates a new email message and applies the **Confidential** label to this message. They then add a Word document that is labeled **General** and this file isn't encrypted. As a result of inheritance, the document gets relabeled as **Confidential** and now has encryption applied from that label.

Sensitivity label compatibility

With RMS-enlightened apps: If you open a labeled and encrypted document or email in an RMS-enlightened application that doesn't support sensitivity labels, the app still enforces encryption and rights management.

With the Azure Information Protection client: You can view and change sensitivity labels that you apply to documents and emails with the Office built-in labeling client by using the Azure Information Protection client, and the other way around.

With other versions of Office: Any authorized user can open labeled documents and emails in other versions of Office. However, you can only view or change the label in supported Office versions or by using the Azure Information Protection client. Supported Office app versions are listed in the previous section.

Support for SharePoint and OneDrive files protected by sensitivity labels

To use the Office built-in labeling client with Office on the web for documents in SharePoint or OneDrive, make sure you've enabled sensitivity labels for Office files in SharePoint and OneDrive.

Support for external users and labeled content

When you label a document or email, the label is stored as metadata that includes your tenant and a label GUID. When a labeled document or email is opened by an Office app that supports sensitivity labels, this metadata is read and only if the user belongs to the same tenant, the label displays in their app. For example, for built-in labeling for Word, PowerPoint, and Excel, the label name displays on the status bar.

This means that if you share documents with another organization that uses different label names, each organization can apply and see their own label applied to the document. However, the following elements from an applied label are visible to users outside your organization:

- Content markings. When a label applies a header, footer, or watermark, these are added directly to the content and remain visible until somebody modifies or deletes them.
- The name and description of the underlying protection template from a label that applied encryption. This information displays in a message bar at the top of the document, to provide information about who is authorized to open the document, and their usage rights for that document.

Sharing encrypted documents with external users

In addition to restricting access to users in your own organization, you can extend access to any other user who has an account in Azure Active Directory. However, if your organization uses Conditional Access policies, see the next section for additional considerations.

All Office apps and other RMS-enlightened application can open encrypted documents after the user has

successfully authenticated.

If external users do not have an account in Azure Active Directory, they can authenticate by using guest accounts in your tenant. These guest accounts can also be used to access shared documents in SharePoint or OneDrive when you have enabled sensitivity labels for Office files in SharePoint and OneDrive:

- One option is to create these guest accounts yourself. You can specify any email address that these users already use. For example, their Gmail address.

  The advantage of this option is that you can restrict access and rights to specific users by specifying their email address in the encryption settings. The downside is the administration overhead for the account creation and coordination with the label configuration.

- Another option is to use SharePoint and OneDrive integration with Azure AD B2B (Preview) so that guest accounts are automatically created when your users share links.

  The advantage of this option is minimum administrative overhead because the accounts are created automatically, and simpler label configuration. For this scenario, you must select the encryption option Add any authenticated user because you won't know the email addresses in advance. The downside is that this setting doesn't let you restrict access and usage rights to specific users.

External users can also use a Microsoft account to open encrypted documents when they use Windows and Microsoft 365 Apps (formerly Office 365 apps) or the standalone edition of Office 2019. More recently supported for other platforms, Microsoft accounts are also supported for opening encrypted documents on macOS (Microsoft 365 Apps, version 16.42+), Android (version 16.0.13029+), and iOS (version 2.42+). For example, a user in your organization shares an encrypted document with a user outside your organization, and the encryption settings specify a Gmail email address for the external user. This external user can create their own Microsoft account that uses their Gmail email address. Then, after signing in with this account, they can open the document and edit it, according to the usage restrictions specified for them. For a walkthrough example of this scenario, see Opening and editing the protected document.

> **NOTE**
> The email address for the Microsoft account must match the email address that's specified to restrict access for the encryption settings.

When a user with a Microsoft account opens an encrypted document in this way, it automatically creates a guest account for the tenant if a guest account with the same name doesn't already exist. When the guest account exists, it can then be used to open documents in SharePoint and OneDrive by using Office on the web, in addition to opening encrypted documents from the supported desktop and mobile Office apps.

However, the automatic guest account is not created immediately in this scenario, because of replication latency. If you specify personal email addresses as part of your label encryption settings, we recommend that you create corresponding guest accounts in Azure Active Directory. Then let these users know that they must use this account to open an encrypted document from your organization.

> **TIP**
> Because you can't be sure that external users will be using a supported Office client app, sharing links from SharePoint and OneDrive after creating guest accounts (for specific users) or when you use SharePoint and OneDrive integration with Azure AD B2B (for any authenticated user) is a more reliable method to support secure collaboration with external users.

Conditional Access policies

If your organization has implemented Azure Active Directory Conditional Access policies, check the configuration of those policies. If the policies include Azure Information Protection and the policy extends to

external users, those external users must have a guest account in your tenant even if they have an Azure AD account in their own tenant.

Without this guest account, they can't open the encrypted document and see an error message. The message text might inform them that their account needs to be added as an external user in the tenant, with the incorrect instruction for this scenario to **Sign out and sign in again with a different Azure Active Directory user account**.

If you can't create and configure guest accounts in your tenant for external users who need to open documents that are encrypted by your labels, you must either remove Azure Information Protection from the Conditional Access policies, or exclude external users from the policies.

For more information about Conditional Access and Azure Information Protection, the encryption service used by sensitivity labels, see the frequently asked question, I see Azure Information Protection is listed as an available cloud app for conditional access—how does this work?

When Office apps apply content marking and encryption

Office apps apply content marking and encryption with a sensitivity label differently, depending on the app you use.

APP	CONTENT MARKING	ENCRYPTION
Word, Excel, PowerPoint on all platforms	Immediately	Immediately
Outlook for PC and Mac	After Exchange Online sends the email	Immediately
Outlook on the web, iOS, and Android	After Exchange Online sends the email	After Exchange Online sends the email

Solutions that apply sensitivity labels to files outside Office apps do so by applying labeling metadata to the file. In this scenario, content marking from the label's configuration isn't inserted into the file but encryption is applied.

When those files are opened in an Office desktop app, the content markings are automatically applied by the Azure Information Protection unified labeling client. The content markings are not automatically applied when you use built-in labeling for desktop, mobile, or web apps.

Scenarios that include applying a sensitivity label outside Office apps include:

- The scanner, File Explorer, and PowerShell from the Azure Information Protection unified labeling client
- Auto-labeling policies for SharePoint and OneDrive
- Exported labeled and encrypted data from Power BI
- Microsoft Cloud App Security

For these scenarios, using their Office apps, a user with built-in labeling can apply the label's content markings by temporarily removing or replacing the current label and then reapplying the original label.

Dynamic markings with variables

IMPORTANT

Currently, not all apps on all platforms support dynamic content markings that you can specify for your headers, footers, and watermarks. For apps that don't support this capability, they apply the markings as the original text specified in the label configuration, rather than resolving the variables.

The Azure Information Protection unified labeling client supports dynamic markings. For labeling built in to Office, see the tables in the capabilities section on this page.

When you configure a sensitivity label for content markings, you can use the following variables in the text string for your header, footer, or watermark:

VARIABLE	DESCRIPTION	EXAMPLE WHEN LABEL APPLIED
`${Item.Label}`	Label display name of the label applied	**General**
`${Item.Name}`	File name or email subject of the content being labeled	**Sales.docx**
`${Item.Location}`	Path and file name of the document being labeled, or the email subject for an email being labeled	**\\Sales\2020\Q3\Report.docx**
`${User.Name}`	Display name of the user applying the label	**Richard Simone**
`${User.PrincipalName}`	Azure AD user principal name (UPN) of the user applying the label	**rsimone@contoso.com**
`${Event.DateTime}`	Date and time when the content is labeled, in the local time zone of the user applying the label	**8/10/2020 1:30 PM**

NOTE

The syntax for these variables is case-sensitive.

Setting different visual markings for Word, Excel, PowerPoint, and Outlook

As an additional variable, you can configure visual markings per Office application type by using an "If.App" variable statement in the text string, and identify the application type by using the values **Word**, **Excel**, **PowerPoint**, or **Outlook**. You can also abbreviate these values, which is necessary if you want to specify more than one in the same If.App statement.

NOTE

For completeness, instructions for Outlook are included, although currently supported only by the Azure Information Protection unified labeling client.

Use the following syntax:

```
${If.App.<application type>}<your visual markings text> ${If.End}
```

As with the other dynamic visual markings, the syntax is case-sensitive.

Examples:

- **Set header text for Word documents only:**

  ```
  ${If.App.Word}This Word document is sensitive ${If.End}
  ```

  In Word document headers only, the label applies the header text "This Word document is sensitive". No header text is applied to other Office applications.

- **Set footer text for Word, Excel, and Outlook, and different footer text for PowerPoint:**

  ```
  ${If.App.WXO}This content is confidential. ${If.End}${If.App.PowerPoint}This presentation is confidential. ${If.End}
  ```

  In Word, Excel, and Outlook, the label applies the footer text "This content is confidential." In PowerPoint, the label applies the footer text "This presentation is confidential."

- **Set specific watermark text for Word and PowerPoint, and then watermark text for Word, Excel, and PowerPoint:**

  ```
  ${If.App.WP}This content is ${If.End}Confidential
  ```

  In Word and PowerPoint, the label applies the watermark text "This content is Confidential". In Excel, the label applies the watermark text "Confidential". In Outlook, the label doesn't apply any watermark text because watermarks as visual markings are not supported for Outlook.

Require users to apply a label to their email and documents

> **IMPORTANT**
> Also known as mandatory labeling, not all apps on all platforms currently support the policy setting of **Require users to apply a label to their email and documents**.
>
> The Azure Information Protection unified labeling client supports mandatory labeling and for labeling built in to Office apps, see the tables in the capabilities section on this page.

When this policy setting is selected, users assigned the policy must select and apply a sensitivity label under the following scenarios:

- For the Azure Information Protection unified labeling client:
  - For documents (Word, Excel, PowerPoint): When an unlabeled document is saved or users close the document.
  - For emails (Outlook): At the time users send an unlabeled message.
- For labeling built in to Office apps:
  - For documents ((Word, Excel, PowerPoint): When an unlabeled document is opened or saved.
  - For emails (Outlook): At the time users send an unlabeled email message.

Additional information for built-in labeling:

- When users are prompted to add a sensitivity label because they open an unlabeled document, they can add a label or choose to open the document in read-only mode.
- When mandatory labeling is in effect, users can't remove sensitivity labels from documents, but can change an existing label.

For guidance about when to use this setting, see the information about policy settings.

End-user documentation

- Apply sensitivity labels to your documents and email within Office
- Known issues when you apply sensitivity labels to your Office files

Encryption

2/18/2021 • 4 minutes to read • Edit Online

Encryption is an important part of your file protection and information protection strategy. This article provides an overview of encryption for Office 365. Get help with encryption tasks like how to set up encryption for your organization and how to password-protect Office documents.

- For information about certificates and technologies like TLS, see Technical reference details about encryption in Office 365.
- For information about how to configure or set up encryption for your organization, see Set up encryption in Office 365 Enterprise.

What is encryption, and how does it work in Office 365?

The encryption process encodes your data (referred to as plaintext) into ciphertext. Unlike plaintext, ciphertext can't be used by people or computers unless and until the ciphertext is decrypted. Decryption requires an encryption key that only authorized users have. Encryption helps ensure that only authorized recipients can decrypt your content. Content includes files, email messages, calendar entries, and so on.

Encryption by itself doesn't prevent content interception. Encryption is part of a larger information protection strategy for your organization. By using encryption, you help ensure that only authorized parties can use the encrypted data.

You can have multiple layers of encryption in place at the same time. For example, you can encrypt email messages and also the communication channels through which your email flows. With Office 365, your data is encrypted at rest and in transit, using several strong encryption protocols, and technologies that include Transport Layer Security/Secure Sockets Layer (TLS/SSL), Internet Protocol Security (IPSec), and Advanced Encryption Standard (AES).

Encryption for data at rest and data in transit

Examples of data at rest include files that you've uploaded to a SharePoint library, Project Online data, documents that you've uploaded in a Skype for Business meeting, email messages and attachments that you've stored in folders in your mailbox, and files you've uploaded to OneDrive for Business.

Examples of data in transit include mail messages that are in the process of being delivered, or conversations that are taking place in an online meeting. In Office 365, data is in transit whenever a user's device is communicating with a Microsoft server, or when a Microsoft server is communicating with another server.

With Office 365, multiple layers and kinds of encryption work together to secure your data. The following table includes some examples, with links to additional information.

KINDS OF CONTENT	ENCRYPTION TECHNOLOGIES	RESOURCES TO LEARN MORE
Files on a device. These files can include email messages saved in a folder, Office documents saved on a computer, tablet, or phone, or data saved to the Microsoft cloud.	BitLocker in Microsoft datacenters. BitLocker can also be used on client machines, such as Windows computers and tablets Distributed Key Manager (DKM) in Microsoft datacenters Customer Key for Microsoft 365	Windows IT Center: BitLocker Microsoft Trust Center: Encryption Cloud security controls series: Encrypting Data at Rest How Exchange Online secures your email secrets Service encryption with Customer Key

KINDS OF CONTENT	ENCRYPTION TECHNOLOGIES	RESOURCES TO LEARN MORE
Files in transit between users. These files can include Office documents or SharePoint list items shared between users.	TLS for files in transit	Data Encryption in OneDrive for Business and SharePoint Online Skype for Business Online: Security and Archiving
Email in transit between recipients. This email includes email hosted by Exchange Online.	Office 365 Message Encryption with Azure Rights Management, S/MIME, and TLS for email in transit	Office 365 Message Encryption (OME) Email encryption in Office 365 How Exchange Online uses TLS to secure email connections in Office 365
Chats, messages, and files in transit between recipients using Microsoft Teams.	Teams uses TLS and MTLS to encrypt instant messages. Media traffic is encrypted using Secure RTP (SRTP). Teams uses FIPS (Federal Information Processing Standard) compliant algorithms for encryption key exchanges.	Encryption for Teams

What if I need more control over encryption to meet security and compliance requirements?

Microsoft 365 provides Microsoft-managed solutions for volume encryption, file encryption, and mailbox encryption in Office 365. In addition, Microsoft provides encryption solutions that you can manage and control. These encryption solutions are built on Azure.

To learn more, see the following resources:

- What is Azure Rights Management?
- Activate Rights Management in the admin center
- Set up Information Rights Management (IRM) in SharePoint admin center
- Service encryption with Customer Key in Office 365
- Double Key Encryption for Microsoft 365

How do I...

TO DO THIS TASK	SEE THESE RESOURCES
Set up encryption for my organization	Set up encryption in Office 365 Enterprise
View details about certificates, technologies, and TLS cipher suites	Technical details about encryption
Work with encrypted messages on a mobile device	View encrypted messages on your Android device View encrypted messages on your iPhone or iPad
Encrypt a document using password protection Password protection isn't supported in a browser. Use desktop versions of Word, Excel, and PowerPoint for password protection.	Add or remove protection in your document, workbook, or presentation Choose an **Add protection** section, and then see **Encrypt with Password**.

TO DO THIS TASK	SEE THESE RESOURCES
Remove encryption from a document	Add or remove protection in your document, workbook, or presentation Choose a **Remove protection** section, and then see **Remove password encryption**.

Related topics

Plan for Microsoft 365 security and information protection capabilities

Top 10 ways to secure Microsoft 365 for business plans

Microsoft Stream Video level encryption and playback flow

Double Key Encryption for Microsoft 365

2/18/2021 • 17 minutes to read • Edit Online

Applies to: Double Key Encryption for Microsoft 365, Microsoft 365 Compliance, Azure Information Protection

Instructions for: Azure Information Protection unified labeling client for Windows

Service description for: Microsoft 365 Compliance

Double Key Encryption (DKE) uses two keys together to access protected content. Microsoft stores one key in Microsoft Azure, and you hold the other key. You maintain full control of one of your keys using the Double Key Encryption service. You apply protection using The Azure Information Protection unified labeling client to your highly sensitive content.

Double Key Encryption supports both cloud and on-premises deployments. These deployments help to ensure that encrypted data remains opaque wherever you store the protected data.

For more information about the default, cloud-based tenant root keys, see Planning and implementing your Azure Information Protection tenant key.

When your organization should adopt DKE

Double Key Encryption is intended for your most sensitive data that is subject to the strictest protection requirements. DKE is not intended for all data. In general, you'll be using Double Key Encryption to protect only a small part of your overall data. You should do due diligence in identifying the right data to cover with this solution before you deploy. In some cases, you might need to narrow your scope and make use of other solutions for most your data such as Microsoft Information Protection with Microsoft-managed keys or BYOK. These solutions are sufficient for documents that aren't subject to enhanced protections and regulatory requirements. Also, these solutions enable you to use the most powerful Office 365 services; services that you can't use with DKE encrypted content. For example:

- Transport rules including anti-malware and spam that require visibility into the attachment
- Microsoft Delve
- eDiscovery
- Content search and indexing
- Office Web Apps including coauthoring functionality

Any external applications or services that are not integrated with DKE through the MIP SDK will be unable to perform actions on the encrypted data.

The Microsoft Information Protection SDK 1.7+ supports Double Key Encryption; applications that integrate with our SDK will be able to reason over this data with sufficient permissions and integrations in place.

We recommend organizations use Microsoft Information protection capabilities (classification and labeling) to protect most of their sensitive data and only use DKE for their mission-critical data. Double Key Encryption is relevant for sensitive data in highly regulated industries such as Financial services and Healthcare.

If your organizations have any of the following requirements, you can use DKE to help secure your content:

- You want to ensure that *only you* can ever decrypt protected content, under all circumstances.
- You don't want Microsoft to have access to protected data on its own.

- You have regulatory requirements to hold keys within a geographical boundary. All of the keys that you hold for data encryption and decryption are maintained in your data center.

System and licensing requirements for DKE

Double Key Encryption for Microsoft 365 comes with Microsoft 365 E5. If you don't have a Microsoft 365 E5 license, you can sign up for a trial. For more information about these licenses, see Microsoft 365 licensing guidance for security & compliance.

Azure Information Protection. DKE works with sensitivity labels and requires Azure Information Protection.

DKE sensitivity labels are made available to end users through the sensitivity ribbon in Office Desktop Apps. Install these prerequisites on each client computer where you want to protect and consume protected documents.

Microsoft Office Apps for enterprise version *.12711 or later (Desktop versions of Word, PowerPoint, and Excel) on Windows.

Azure Information Protection Unified Labeling Client versions 2.7.93.0 or later. Download and install the Unified Labeling client from the Microsoft download center.

Supported environments for storing and viewing DKE-protected content

Supported applications. Microsoft 365 Apps for enterprise clients on Windows, including Word, Excel, and PowerPoint.

Online content support. Documents and files stored online in both Microsoft SharePoint and OneDrive for Business are supported. You can share encrypted content by email, but you can't view encrypted documents and files online. Instead, you must view protected content using the desktop apps on your local computer.

Overview of deploying DKE

You'll follow these general steps to set up DKE. Once you've completed these steps, your end users will can protect your highly sensitive data with Double Key Encryption.

1. Deploy the DKE service as described in this article.
2. Create a label with Double Key Encryption. Navigate to Information protection under the Microsoft 365 compliance center and create a new label with Double Key Encryption. See Restrict access to content by using sensitivity labels to apply encryption.
3. Use Double Key Encryption labels. Protect data by selecting the Double Key Encrypted label from the Sensitivity ribbon in Microsoft Office.

There are several ways you can complete some of the steps to deploy Double Key Encryption. This article provides detailed instructions so that less experienced admins successfully deploy the service. If you're comfortable doing so, you can choose to use your own methods.

Deploy DKE

This article and the deployment video use Azure as the deployment destination for the DKE service. If you're deploying to another location, you'll need to provide your own values.

Watch the Double Key Encryption deployment video to see a step-by-step overview of the concepts in this article. The video takes about 18 minutes to complete.

You'll follow these general steps to set up Double Key Encryption for your organization.

1. Install software prerequisites for the DKE service
2. Clone the Double Key Encryption GitHub repository
3. Modify application settings
4. Generate test keys
5. Build the project
6. Deploy the DKE service and publish the key store
7. Validate your deployment
8. Register your key store
9. Create sensitivity labels using DKE
10. Enable DKE in your client
11. Migrate protected files from HYOK labels to DKE labels

When you're done, you can encrypt documents and files using DKE. For information, see Apply sensitivity labels to your files and email in Office.

Install software prerequisites for the DKE service

Install these prerequisites on the computer where you want to install the DKE service.

.NET Core 3.1 SDK. Download and install the SDK from Download .NET Core 3.1.

Visual Studio Code. Download Visual Studio Code from https://code.visualstudio.com/. Once installed, run Visual Studio Code and select **View** > **Extensions**. Install these extensions.

- C# for Visual Studio Code
- NuGet Package Manager

Git resources. Download and install one of the following.

- Git
- GitHub Desktop
- GitHub Enterprise

OpenSSL You must have OpenSSL installed to generate test keys after you deploy DKE. Make sure you're invoking it correctly from your environment variables path. For example, see "Add the installation directory to PATH" at https://www.osradar.com/install-openssl-windows/ for details.

Clone the DKE GitHub repository

Microsoft supplies the DKE source files in a GitHub repository. You clone the repository to build the project locally for your organization's use. The DKE GitHub repository is located at https://github.com/Azure-Samples/DoubleKeyEncryptionService.

The following instructions are intended for inexperienced git or Visual Studio Code users:

1. In your browser, go to: https://github.com/Azure-Samples/DoubleKeyEncryptionService.
2. Towards the right side of the screen, select **Code**. Your version of the UI might show a **Clone or download** button. Then, in the dropdown that appears, select the copy icon to copy the URL to your clipboard.

   For example:

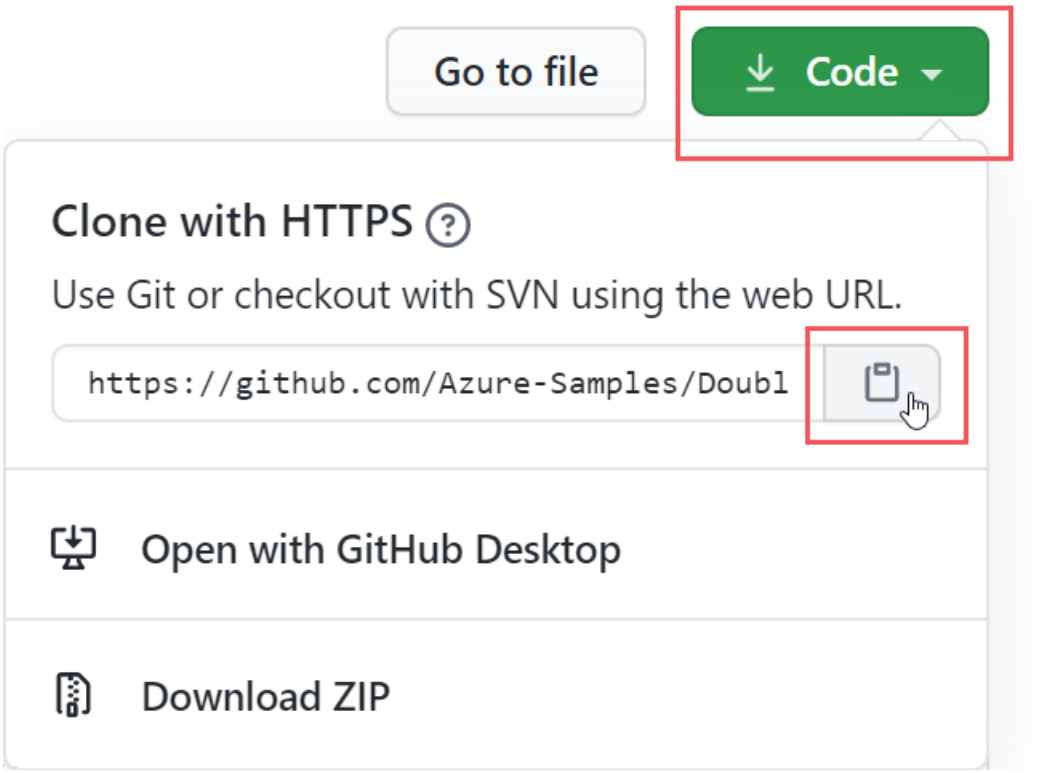

3. In Visual Studio Code, select **View** > **Command Palette** and select **Git: Clone**. To jump to the option in the list, start typing `git: clone` to filter the entries and then select it from the drop-down. For example:

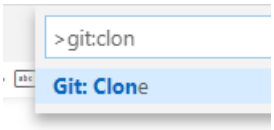

4. In the text box, paste the URL that you copied from Git and select **Clone from GitHub**.

5. In the **Select Folder** dialog that appears, browse to and select a location to store the repository. At the prompt, select **Open**.

   The repository opens in Visual Studio Code, and displays the current Git branch at the bottom left. For example, The branch should be **main**. For example:

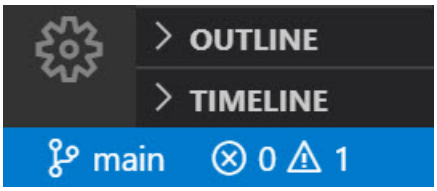

6. If you're not on the main branch, you'll need to select it. In Visual Studio Code, select the branch and choose **main** from the list of branches that displays.

   > **IMPORTANT**
   > Selecting the main branch ensures that you have the correct files to build the project. If you don't choose the correct branch your deployment will fail.

You now have your DKE source repository set up locally. Next, modify application settings for your organization.

Modify application settings

To deploy the DKE service, you must modify the following types of application settings:

- Key access settings
- Tenant and key settings

You modify application settings in the appsettings.json file. This file is located in the DoubleKeyEncryptionService repo you cloned locally under DoubleKeyEncryptionService\src\customer-key-store. For example, in Visual Studio Code, you can browse to the file as shown in the following picture.

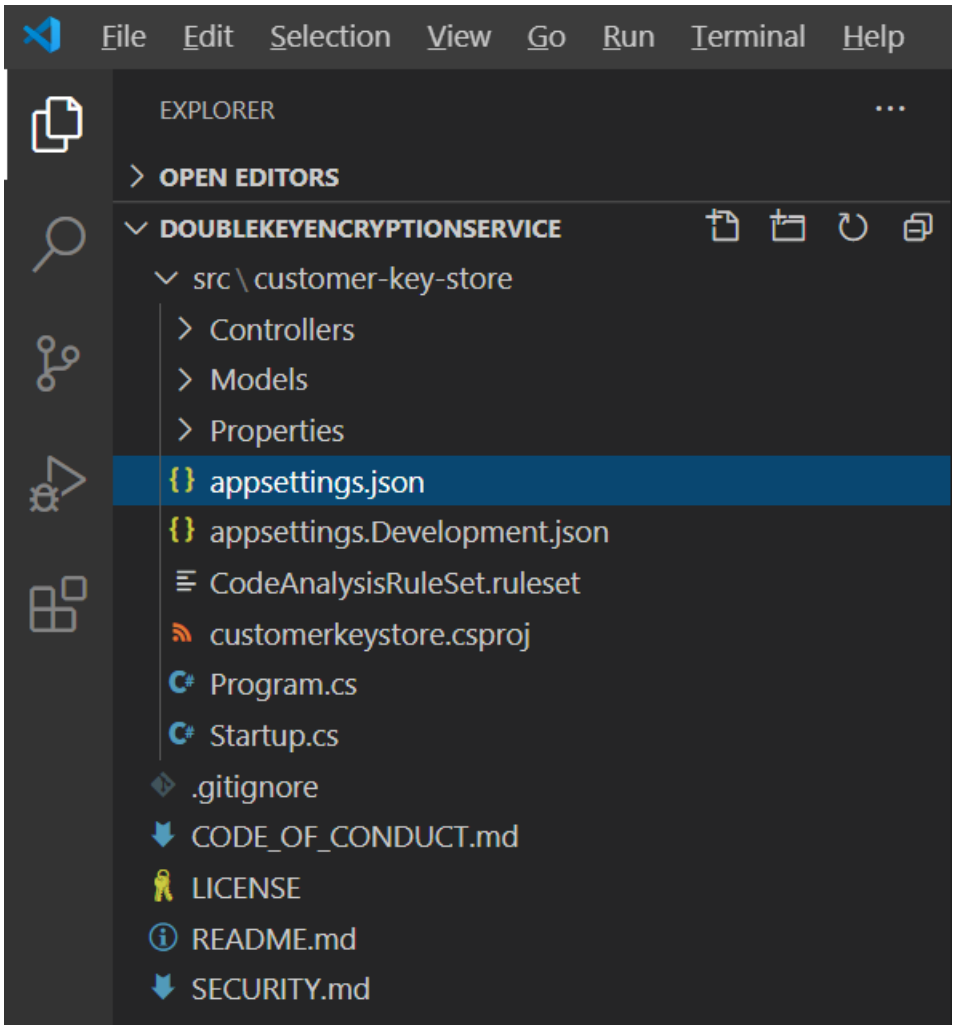

Key access settings

Choose whether to use email or role authorization. DKE supports only one of these authentication methods at a time.

- **Email authorization**. Allows your organization to authorize access to keys based on email addresses only.
- **Role authorization**. Allows your organization to authorize access to keys based on Active Directory groups, and requires that the web service can query LDAP.

To set key access settings for DKE using email authorization

1. Open the **appsettings.json** file and locate the `AuthorizedEmailAddress` setting.
2. Add the email address or addresses that you want to authorize. Separate multiple email addresses with double quotes and commas. For example:

   ```
   "AuthorizedEmailAddress": ["email1@company.com", "email2@company.com ", "email3@company.com"]
   ```

3. Locate the `LDAPPath` setting and remove the text `If you use role authorization (AuthorizedRoles) then this is the LDAP path.` between the double

quotes. Leave the double quotes in place. When you're finished, the setting should look like this.

```
"LDAPPath": ""
```

4. Locate the `AuthorizedRoles` setting and delete the entire line.

This image shows the **appsettings.json** file correctly formatted for email authorization.

To set key access settings for DKE using role authorization

1. Open the **appsettings.json** file and locate the `AuthorizedRoles` setting.
2. Add the Active Directory group names you want to authorize. Separate multiple group names with double quotes and commas. For example:

```
"AuthorizedRoles": ["group1", "group2", "group3"]
```

3. Locate the `LDAPPath` setting and add the Active Directory domain. For example:

```
"LDAPPath": "contoso.com"
```

4. Locate the `AuthorizedEmailAddress` setting and delete the entire line.

This image shows the **appsettings.json** file correctly formatted for role authorization.

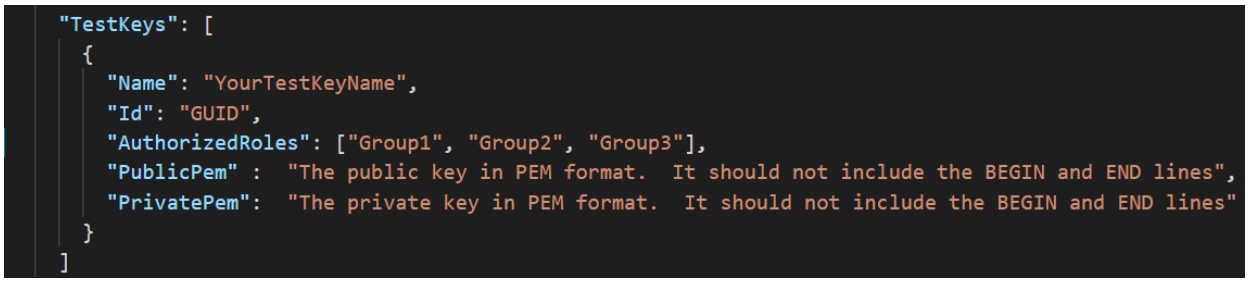

Tenant and key settings

DKE tenant and key settings are located in the **appsettings.json** file.

To configure tenant and key settings for DKE

1. Open the **appsettings.json** file.
2. Locate the `ValidIssuers` setting and replace `<tenantid>` with your tenant ID. You can locate your tenant ID by going to the Azure portal and viewing the tenant properties. For example:

```
"ValidIssuers": [
  "https://sts.windows.net/9c99431e-b513-44be-a7d9-e7b500002d4b/"
]
```

Locate the `JwtAudience`. Replace `<yourhostname>` with the hostname of the machine where the DKE service will run. For example:

> **IMPORTANT**
> The value for `JwtAudience` must match the name of your host *exactly*. You may use **localhost:5001** while debugging. However, When you're done debugging, make sure to update this value to the server's hostname.

- `TestKeys:Name`. Enter a name for your key. For example: `TestKey1`
- `TestKeys:Id`. Create a GUID and enter it as the `TestKeys:ID` value. For example, `DCE1CC21-FF9B-4424-8FF4-9914BD19A1BE`. You can use a site like Online GUID Generator to randomly generate a GUID.

This image shows the correct format for tenant and keys settings in **appsettings.json**. `LDAPPath` is configured for role authorization.

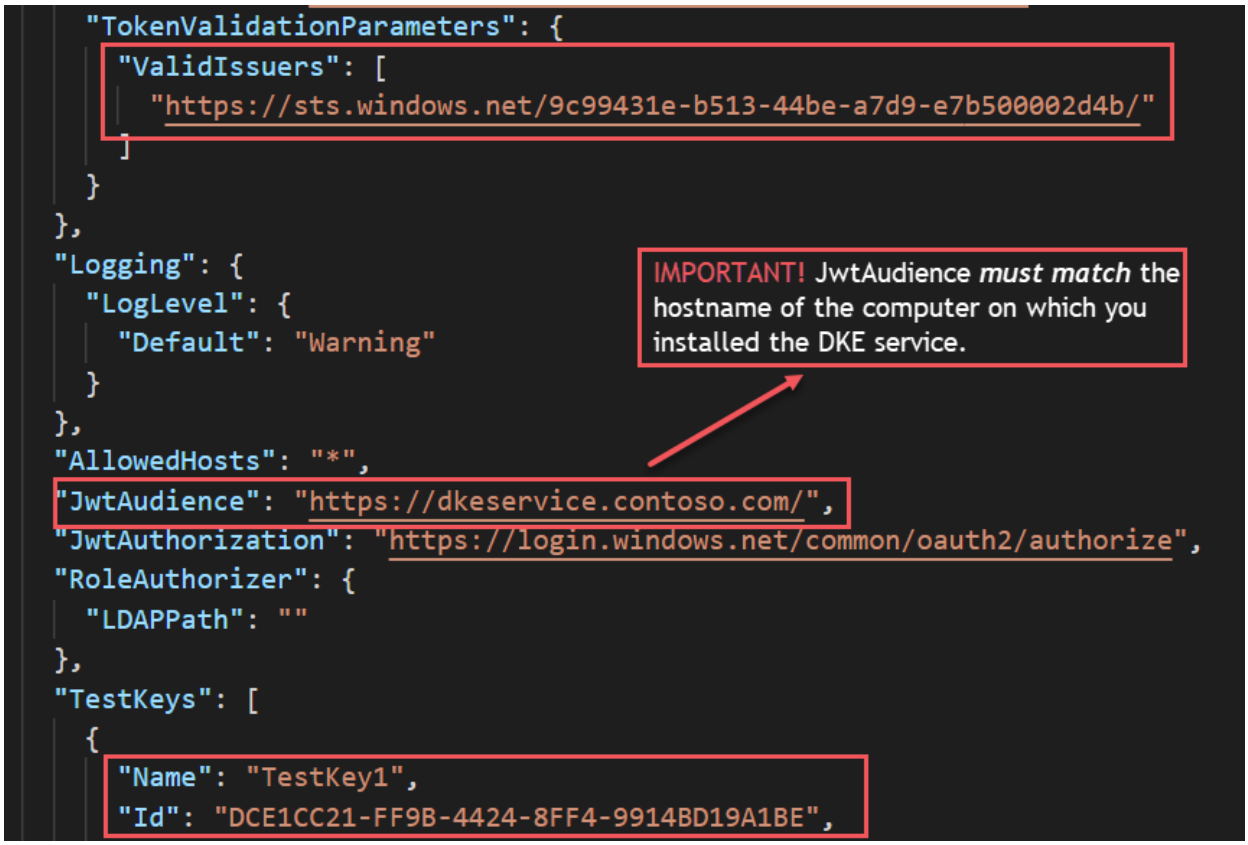

Generate test keys

Once you have your application settings defined, you're ready to generate public and private test keys.

To generate keys:

1. From the Windows Start menu, run the OpenSSL Command Prompt.
2. Change to the folder where you want to save the test keys. The files you create by completing the steps in this task are stored in the same folder.
3. Generate the new test key.

```
openssl req -x509 -newkey rsa:2048 -keyout key.pem -out cert.pem -days 365
```

4. Generate the private key.

```
openssl rsa -in key.pem -out privkeynopass.pem
```

5. Generate the public key.

```
openssl rsa -in key.pem -pubout > pubkeyonly.pem
```

6. In a text editor, open **pubkeyonly.pem**. Copy all of the content in the **pubkeyonly.pem** file, except the first and last lines, into the `PublicPem` section of the **appsettings.json** file.

7. In a text editor, open **privkeynopass.pem**. Copy all of the content in the **privkeynopass.pem** file, except the first and last lines, into the `PrivatePem` section of the **appsettings.json** file.

8. Remove all blank spaces and newlines in both the `PublicPem` and `PrivatePem` sections.

> **IMPORTANT**
> When you copy this content, do not delete any of the PEM data.

9. In Visual Studio Code, browse to the **Startup.cs** file. This file is located in the DoubleKeyEncryptionService repo you cloned locally under DoubleKeyEncryptionService\src\customer-key-store.

10. Locate the following lines:

```
    #if USE_TEST_KEYS
    #error !!!!!!!!!!!!!!!!!!!! Use of test keys is only supported for testing,
    DO NOT USE FOR PRODUCTION !!!!!!!!!!!!!!!!!!!!!!!!!!!!!
    services.AddSingleton<ippw.IKeyStore, ippw.TestKeyStore>();
    #endif
```

11. Replace these lines with the following text:

```
services.AddSingleton<ippw.IKeyStore, ippw.TestKeyStore>();
```

The end results should look similar to the following.

```
        services.AddSingleton<ippw.IKeyStore, ippw.TestKeyStore>();

        services.AddTransient<ippw.KeyManager, ippw.KeyManager>();

        services.AddMvc().SetCompatibilityVersion(CompatibilityVersion.Version_2_1);
```

Now you're ready to build your DKE project.

Build the project

Use the following instructions to build the DKE project locally:

1. In Visual Studio Code, in the DKE service repository, select **View** > **Command Palette** and then type **build** at the prompt.

2. From the list, choose **Tasks: Run build task**.

   If there are no build tasks found, select **Configure Build Task** and create one for .NET core as follows.

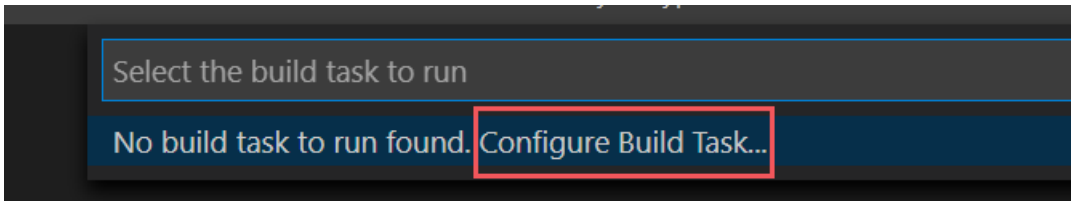

a. Choose **Create tasks.json from template**.

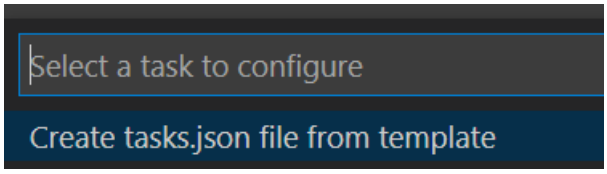

b. From the list of template types, select **.NET Core**.

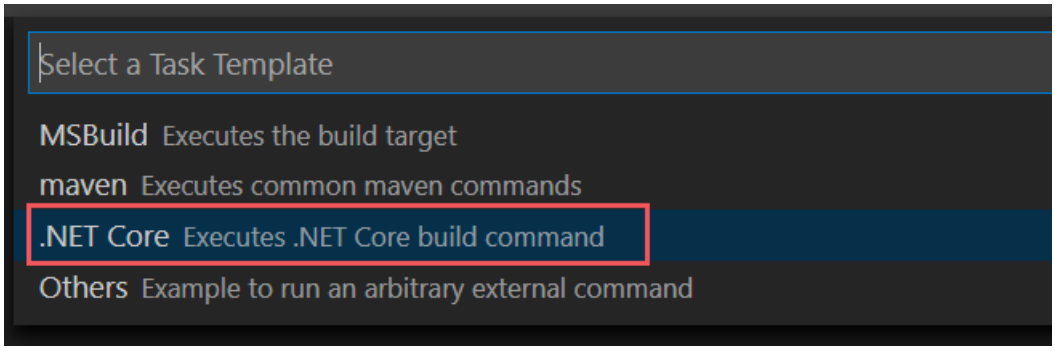

c. In the build section, locate the path to the **customerkeystore.csproj** file. If it's not there, add the following line:

```
"${workspaceFolder}/src/customer-key-store/customerkeystore.csproj",
```

d. Run the build again.

3. Verify that there are no red errors in the output window.

   If there are red errors, check the console output. Ensure that you completed all the previous steps correctly and the correct build versions are present.

4. Select **Run** > **Start Debugging** to debug the process. If you're prompted to select an environment, select **.NET core**.

   The .NET core debugger typically launches to `https://localhost:5001`. To view your test key, go to `https://localhost:5001` and append a forward slash (/) and the name of your key. For example:

   ```
   https://localhost:5001/TestKey1
   ```

   The key should display in JSON format.

Your setup is now complete. Before you publish the keystore, in appsettings.json, for the JwtAudience setting, ensure the value for hostname exactly matches your App Service host name. You may have changed it to localhost to troubleshoot the build.

Deploy the DKE service and publish the key store

For production deployments, deploy the service either in a third-party cloud or publish to an on-premises system.

You may prefer other methods to deploy your keys. Select the method that works best for your organization.

For pilot deployments, you can deploy in Azure and get started right away.

To create an Azure Web App instance to host your DKE deployment

To publish the key store, you'll create an Azure App Service instance to host your DKE deployment. Next, you'll publish your generated keys to Azure.

1. In your browser, sign in to the Microsoft Azure portal, and go to **App Services** > **Add**.
2. Select your subscription and resource group and define your instance details.
   - Enter the hostname of the computer where you want to install the DKE service. Make sure it's the same name as the one defined for the JwtAudience setting in the **appsettings.json** file. The value you provide for the name is also the WebAppInstanceName.
   - For **Publish**, select **code**, and for **Runtime stack**, select **.NET Core 3.1**.

   For example:

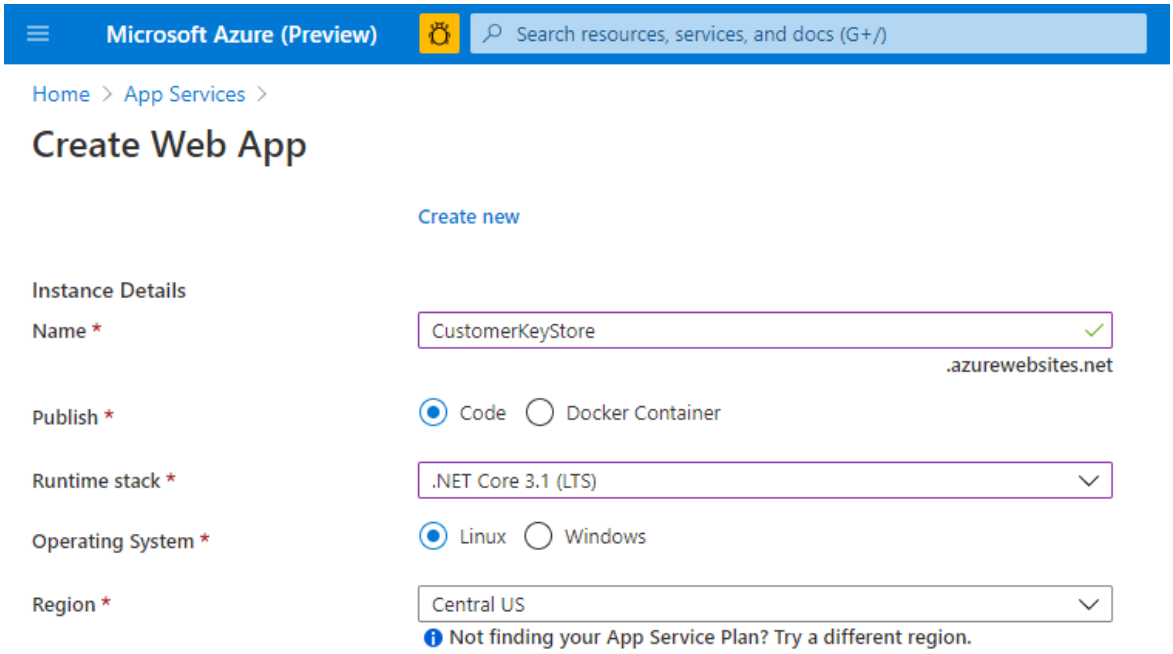

3. At the bottom of the page, select **Review + create**, and then select **Add**.
4. Do one of the following to publish your generated keys:
   - Publish via ZipDeployUI
   - Publish via FTP
   - Publish via Visual Studio 2019 or later

Publish via ZipDeployUI

1. Go to `https://<WebAppInstanceName>.scm.azurewebsites.net/ZipDeployUI`.

   For example: https://dkeservice.scm.azurewebsites.net/ZipDeployUI

2. In the codebase for the key store, go to the **customer-key-store\src\customer-key-store** folder, and verify that this folder contains the **customerkeystore.csproj** file.
3. Run: **dotnet publish**

   The output window displays the directory where the publish was deployed.

   For example: `customer-key-store\src\customer-key-store\bin\Debug\netcoreapp3.1\publish\`

4. Send all files in the publish directory to a .zip file. When creating the .zip file, make sure that all files in the directory are at the root level of the .zip file.

5. Drag and drop the .zip file you create to the ZipDeployUI site you opened above. For example: https://dkeservice.scm.azurewebsites.net/ZipDeployUI

DKE is deployed and you can browse to the test keys you've created. Continue to Validate your deployment below.

Publish via FTP

1. Connect to the App Service you created above.

   In your browser, go to: **Azure portal** > **App Service** > **Deployment Center** > **Manual Deployment** > **FTP** > **Dashboard**.

2. Copy the connection strings displayed to a local file. You'll use these strings to connect to the Web App Service and upload files via FTP.

   For example:

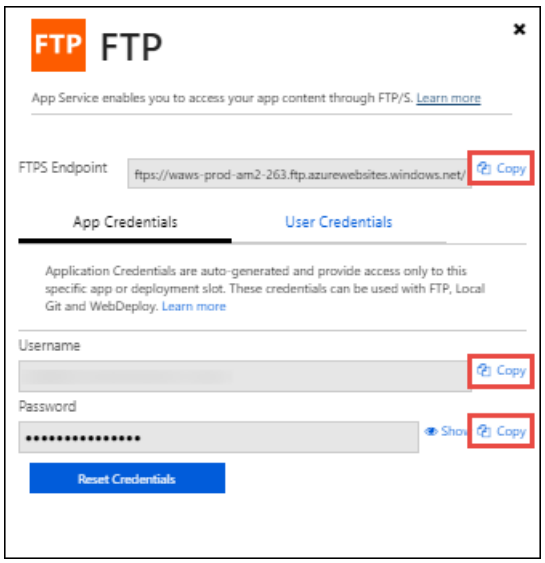

3. In the codebase for the key storage, go to the **customer-key-store\src\customer-key-store directory**.

4. Verify that this directory contains the **customerkeystore.csproj** file.

5. Run: **dotnet publish**

   The output contains the directory where the publish was deployed.

   For example: `customer-key-store\src\customer-key-store\bin\Debug\netcoreapp3.1\publish\`

6. Send all files in the publish directory to a zip file. When creating the .zip file, make sure that all files in the directory are at the root level of the .zip file.

7. From your FTP client, use the connection information you copied to connect to your App Service. Upload the .zip file you created in the previous step to the root directory of your Web App.

DKE is deployed and you can browse to the test keys you'd created. Next, Validate your deployment.

Validate your deployment

After deploying DKE using one of the methods described above, validate the deployment and the key store settings.

Run:

```
src\customer-key-store\scripts\key_store_tester.ps1 dkeserviceurl/mykey
```

For example:

```
key_store_tester.ps1 https://mydkeservice.com/mykey
```

Ensure that no errors appear in the output. When you're ready, register your key store.

The key name is case sensitive. Enter the key name as it appears in the appsettings.json file.

Register your key store

The following steps enable you to register your DKE service. Registering your DKE service is the last step in deploying DKE before you can start creating labels.

To register the DKE service:

1. In your browser, open the Microsoft Azure portal, and go to **All Services** > **Identity** > **App Registrations**.
2. Select **New registration**, and enter a meaningful name.
3. Select an account type from the options displayed.

   If you're using Microsoft Azure with a non-custom domain, such as **onmicrosoft.com**, select **Accounts in this organizational directory only (Microsoft only - Single tenant)**.

   For example:

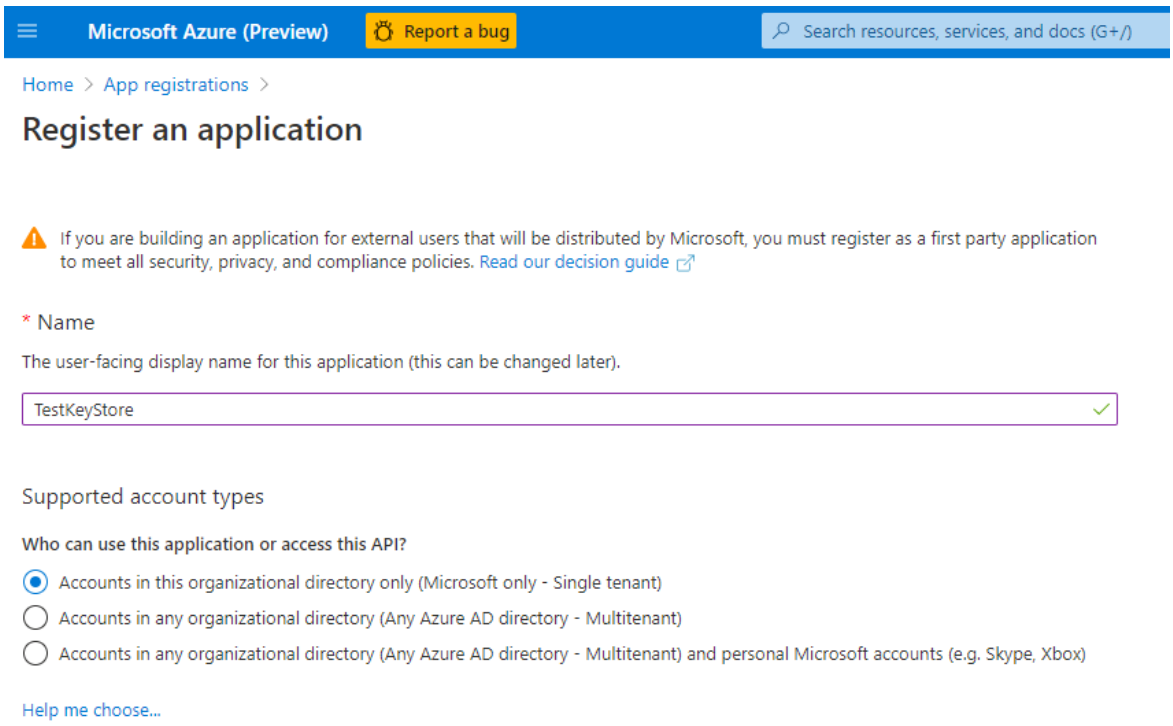

4. At the bottom of the page, select **Register** to create the new App Registration.
5. In your new App Registration, in the left pane, under **Manage**, select **Authentication**.
6. Select **Add a platform**.
7. On the **Configure platforms** popup, select **Web**.

8. Under **Redirect URIs**, enter the URI of your double key encryption service. Enter the App Service URL, including both the hostname and domain.

   For example: https://mydkeservicetest.com

   - The URL you enter must match the hostname where your DKE service is deployed.
   - If you're testing locally with Visual Studio, use **https://localhost:5001**.
   - In all cases, the scheme must be **https**.

   Ensure the hostname exactly matches your App Service hostname. You may have changed it to `localhost` to troubleshoot the build. In **appsettings.json**, this value is the hostname you set for `JwtAudience`.

9. Under **Implicit grant**, select the **ID tokens** checkbox.
10. Select **Save** to save your changes.
11. On the left pane, select **Expose an API**, then next to Application ID URI, select **Set**.
12. Still on the **Expose an API** page, in the **Scopes defined by this API** area, select **Add a scope**. In the new scope:

    a. Define the scope name as **user_impersonation**.

    b. Select the administrators and users who can consent.

    c. Define any remaining values required.

    d. Select **Add scope**.

    e. Select **Save** at the top to save your changes.

13. Still on the **Expose an API** page, in the **Authorized client applications** area, select **Add a client application**.

    In the new client application:

    a. Define the Client ID as **d3590ed6-52b3-4102-aeff-aad2292ab01c**. This value is the Microsoft Office client ID, and enables Office to obtain an access token for your key store.

    b. Under **Authorized scopes**, select the **user_impersonation** scope.

    c. Select **Add application**.

    d. Select **Save** at the top to save your changes.

Your DKE service is now registered. Continue by creating labels using DKE.

Create sensitivity labels using DKE

In the Microsoft 365 compliance center, create a new sensitivity label and apply encryption as you would otherwise. Select **Use Double Key Encryption** and enter the endpoint URL for your key.

For example:

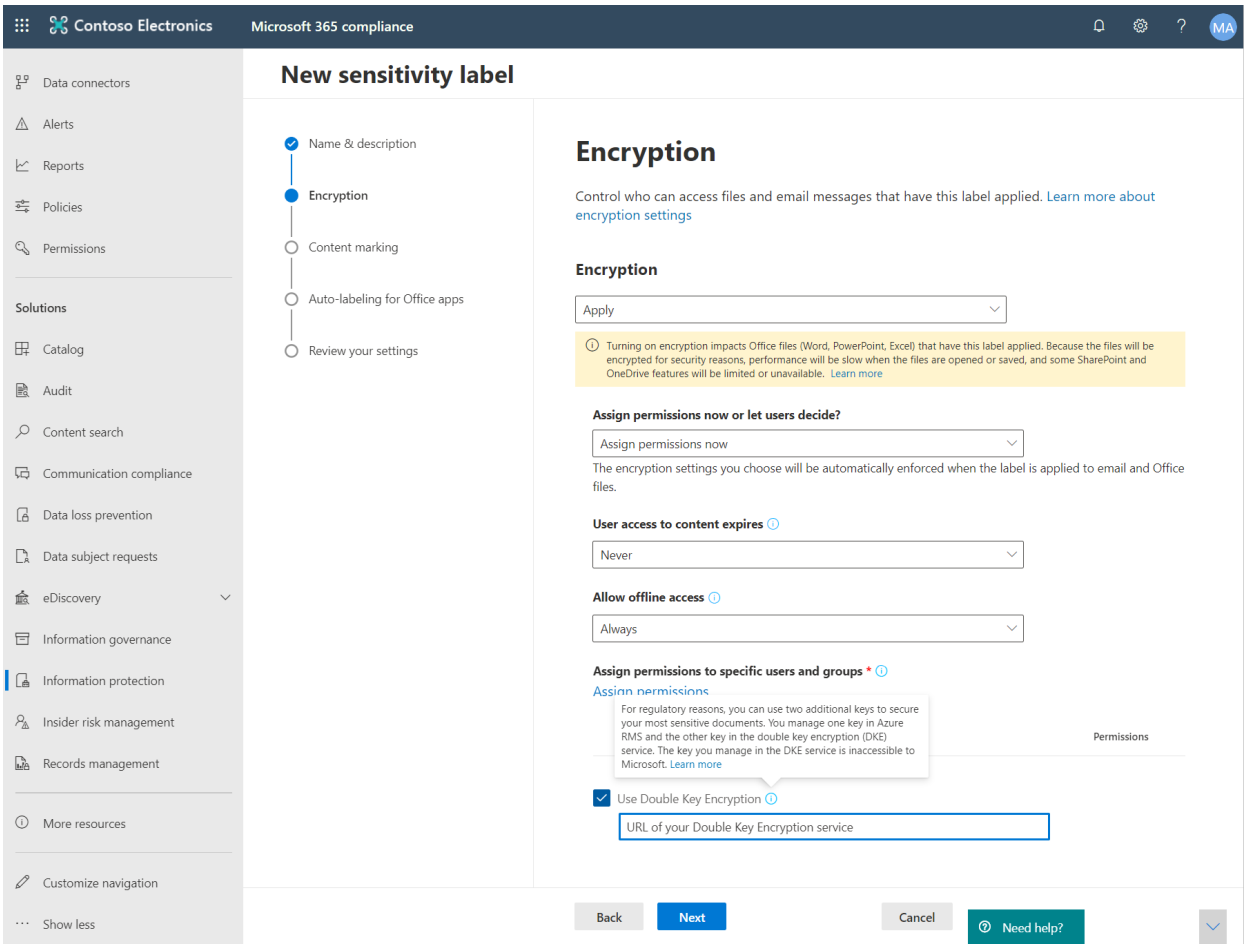

Any DKE labels you add will start appearing for users in the latest versions of Microsoft 365 Apps for enterprise.

> **NOTE**
>
> It may take up to 24 hours for the clients to refresh with the new labels.

Enable DKE in your client

If you're an Office Insider, DKE is enabled for you. Otherwise, enable DKE for your client by adding the following registry keys:

```
[HKEY_LOCAL_MACHINE\SOFTWARE\WOW6432Node\Microsoft\MSIPC\flighting]
"DoubleKeyProtection"=dword:00000001

[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\MSIPC\flighting]
"DoubleKeyProtection"=dword:00000001
```

Migrate protected files from HYOK labels to DKE labels

If you want, once you're finished setting up DKE, you can migrate content that you've protected using HYOK labels to DKE labels. To migrate, you'll use the AIP scanner. To get started using the scanner, see What is the Azure Information Protection unified labeling scanner?.

If you don't migrate content, your HYOK protected content will remain unaffected.

Double Key Encryption frequently asked questions

2/18/2021 • 3 minutes to read • Edit Online

Have a question about how Double Key Encryption works? Check for an answer here.

What is Double Key Encryption for Microsoft 365 (DKE)?

Double Key Encryption for Microsoft 365 enables customers to protect their highly sensitive data to meet specialized requirements. It helps customers maintain full control of their encryption keys. It uses two keys to protect data; one key in your control and a second key stored securely in Microsoft Azure. Viewing data protected with Double Key Encryption requires access to both keys. Since Microsoft can access only one of these keys, protected data remains inaccessible to Microsoft, ensuring that you have full control over your data privacy and security.

You can host the Double Key Encryption service used to request your key, in a location of your choice (on-premises key management server or in the cloud). You maintain the service as you would any other application. Double Key Encryption enables you to control access to the Double Key Encryption service. You can store your highly sensitive data on-premises or move it to the cloud. You can be confident about preventing third-party access because you maintain full control of your key. Double Key Encryption allows you to store your data and key in the same location.

DKE helps you meet regulatory requirements across several regulations and standards such as the General Data Protection Regulation (GDPR), the Health Insurance Portability and Accountability Act (HIPAA), the Gramm-Leach-Bliley Act (GLBA), Russia's data localization law – Federal Law No. 242-FZ, Australia's Federal Privacy Act 1988, and New Zealand's Privacy Act 1993.

Can I use Double Key Encryption with Microsoft Office built-in sensitivity labeling?

You'll need to use the Azure Information Protection unified labeling client to protect documents with Double Key Encryption. Currently, you can't use Microsoft Office built-in sensitivity labeling.

What Microsoft 365 Apps can I use with DKE?

You can use DKE labels to protect documents using the desktop versions of Word, Excel, and PowerPoint on Windows. Ensure that you're using *.12711 or later (Desktop versions of Word, PowerPoint, and Excel) on Windows.

How is Double Key Encryption different from the existing hold your own key (HYOK) solution?

Double Key Encryption encrypts your data with two keys. Your encryption key is in your control and the second key is stored in Microsoft Azure, allowing you to move your encrypted data to the cloud. HYOK protects your content with only one key and the key is always on premises.

Can Double Key Encrypted documents be shared externally?

You can share Double Key Encrypted documents with users on a separate tenant as long as those users:

- Have the required permission to access your key in your Double Key Encryption service.

- Have the required permission to access your key in Microsoft Azure.

What happens to documents that are protected with HYOK?

Deploying Double Key Encryption won't affect your existing HYOK setup. However, we recommend that you start using Double Key Encryption in parallel with HYOK.

Can I run Double Key Encryption in my non-Microsoft air-gapped environment?

DKE doesn't support these environments because the service requires access to Microsoft Azure.

Where can I store Double Key Encrypted documents?

You can store Double Key Encrypted documents on-premises or in the cloud. In the cloud, you can move encrypted content to SharePoint Online and OneDrive for Business. Since Microsoft doesn't have access to your private key, the encrypted data remains opaque to Microsoft. This also means that you can't view the encrypted documents online in Office Web Apps.

What regions and languages is Double Key Encryption available in? Is Double Key Encryption available worldwide?

DKE labels are localized to the same languages as other sensitivity labels in Microsoft Information Protection. Double Key Encryption is available worldwide.

Can I convert a non-DKE label to a DKE label?

No. You can't add DKE to a label after you create it. Instead, you must choose **Use Double Key Encryption** and provide the URL of your Double Key Encryption service when you create the label.

How do I roll my DKE keys?

For instructions on rolling (also called rotating or rekeying) the key you store in Azure, see Operations for your Azure Information Protection tenant key.

See Tenant and key settings for information on creating a new key for the DKE service.

When you create a key, you set up a name and a GUID. Then, if you rotate a key, you keep the old record with the name and GUID but add a new record with the same name but different GUID. The new key gets set as active so that the public key API starts returning it for new encryption. Both keys are available for decryption so that new content and old content can be decrypted.

Service Encryption

2/18/2021 • 2 minutes to read • Edit Online

In addition to using volume-level encryption, Exchange Online, Microsoft Teams, SharePoint Online, and OneDrive for Business also use Service Encryption to encrypt customer data. Service Encryption allows for two key management options:

Microsoft-managed keys

Microsoft manages all cryptographic keys including the root keys for service encryption. This option is currently enabled by default for Exchange Online, SharePoint Online, OneDrive for Business. Microsoft-managed keys provide default service encryption unless you decide to onboard using Customer Key. If, at a later date, you decide to stop using Customer Key without following the data purge path, then your data stays encrypted using the Microsoft-managed keys. Your data is always encrypted at this default level at a minimum.

Customer Key

You supply root keys used with service encryption and you manage these keys using Azure Key Vault. Microsoft manages all other keys. This option is called Customer Key, and it is currently available for Exchange Online, SharePoint Online, and OneDrive for Business. (Previously referred to as Advanced Encryption with BYOK. See Enhancing transparency and control for Office 365 customers for the original announcement.)

Service encryption provides multiple benefits:

- Provides an added layer of protection on top of BitLocker.
- Provides separation of Windows operating system administrators from access to application data stored or processed by the operating system.
- Includes a Customer Key option that enables multi-tenant services to provide per-tenant key management.
- Enhances the ability of Microsoft 365 to meet the demands of customers that have specific compliance requirements regarding encryption.

Using Customer Key, you can generate your own cryptographic keys using either an on-premises Hardware Service Module (HSM) or Azure Key Vault (AKV). Regardless of how you generate the key, you use AKV to control and manage the cryptographic keys used by Office 365. Once your keys are stored in AKV, they can be used as the root of one of the keychains that encrypts your mailbox data or files.

Another benefit of Customer Key is the control you have over the ability of Microsoft to process your data. If you want to remove data from Office 365, such as if you want to terminate service with Microsoft or remove a portion of your data stored in the cloud, you can do so and use Customer Key as a technical control. Removing data ensures that no one, including Microsoft, can access or process the data. Customer Key is in addition and complementary to Customer Lockbox that you use to control access to your data by Microsoft personnel.

To learn how to set up Customer Key for Microsoft 365 for Exchange Online, Microsoft Teams, SharePoint Online, including Team Sites, and OneDrive for Business, see these articles:

- Service encryption with Customer Key
- Set up Customer Key
- Manage Customer Key

- Roll or rotate a customer key or an availability key
- Understand the availability key

Overview of Customer Key for Microsoft 365 at the tenant level (public preview)

2/18/2021 • 19 minutes to read • Edit Online

Using keys you provide, you can create a data encryption policy (DEP) and assign it to the tenant. The DEP encrypts data across the tenant for these workloads:

- Teams chat messages (1:1 chats, group chats, meeting chats and channel conversations)
- Teams media messages (images, code snippets, videos messages, audio messages, wiki images)
- Teams call and meeting recordings stored in Teams storage
- Teams chat notifications
- Teams chat suggestions by Cortana
- Teams status messages
- User and signal information for Exchange Online

For Microsoft Teams, Customer Key at the tenant level encrypts new data from the time the DEP is assigned to the tenant. Public preview does not support encrypting past data. For Exchange Online, Customer Key encrypts all existing and new data.

You can create multiple DEPs per tenant but can only assign one DEP at any point in time. When you assign the DEP, encryption begins automatically but can take some time to complete depending on the size of your tenant.

Tenant level policies add broader control to Customer Key for Microsoft 365

If you already have Customer Key set up for Exchange Online and Sharepoint Online, here's how the new tenant-level public preview fits in.

The tenant-level encryption policy you create encrypts all data for the Microsoft Teams and Exchange Online workloads in Microsoft 365. This policy doesn't interfere with finely tuned DEPs you've already created in Customer Key.

Examples:

Microsoft Teams files and some Teams call and meeting recordings that are saved in OneDrive for Business and SharePoint are encrypted by a SharePoint Online DEP. A single SharePoint Online DEP encrypts content within a single geo.

For Exchange Online, you can create a DEP that encrypts one or more user mailboxes with Customer Key. When you create a tenant-level policy, that policy will not encrypt the encrypted mailboxes. However, the tenant-level key will encrypt the mailboxes that are not affected by a DEP already.

Set up Customer Key at the tenant level (public preview)

These steps are similar but not identical to the steps for setting up Customer Key at the application level. You should only use this public preview with test data in test tenants. Do not use this release with production data or in your production environment. If you already have a production deployment of Customer Key, use these steps to set up Customer Key at the tenant level in a test environment.

You'll complete most of these tasks by remotely connecting to Azure PowerShell. For best results, use version

4.4.0 or later of Azure PowerShell.

Before you get started, make sure of the following:

- You'll need to use a work or school account that has the compliance admin role to set up Customer Key at the tenant level.
- Ensure that you have the appropriate licensing for your organization. Use a paid, invoiced Azure Subscription using either an Enterprise Agreement or a Cloud Service Provider. Azure Subscriptions purchased using Pay As You Go plans or using a credit card aren't supported for Customer Key. Starting April 1, 2020, Customer Key in Office 365 is offered in Office 365 E5, M365 E5, M365 E5 Compliance, and M365 E5 Information Protection & Governance SKUs. Office 365 Advanced Compliance SKU is no longer available for procuring new licenses. Existing Office 365 Advanced Compliance licenses will continue to be supported. While the service can be enabled with a minimum of one license under the tenant having the appropriate license, you should still make sure all users that benefit from the service have appropriate licenses.

Create two new Azure subscriptions

Customer Key requires two keys for each data encryption policy (DEP). To achieve this, you must create two Azure subscriptions. As a best practice, Microsoft recommends that you have separate members of your organization configure one key in each subscription. Only use these Azure subscriptions to administer encryption keys for Microsoft 365. This protects your organization in case one of your operators accidentally, intentionally, or maliciously deletes or otherwise mismanages the keys for which they are responsible.

There is no practical limit to the number of Azure subscriptions that you can create for your organization. Following this best practice helps minimize the impact of human error while helping to manage the resources used by Customer Key.

Register Azure subscriptions to use a mandatory retention period

The temporary or permanent loss of root encryption keys can be disruptive or even catastrophic to service operation and can result in data loss. For this reason, the resources used with Customer Key require strong protection. All the Azure resources that are used with Customer Key offer protection mechanisms beyond the default configuration. Azure subscriptions can be tagged or registered in a way that will prevent immediate and irrevocable cancellation. This is referred to as registering for a mandatory retention period. The steps required to register Azure subscriptions for a mandatory retention period require collaboration with the Microsoft. This process can take up to five business days. Previously, this was sometimes referred to as "Do Not Cancel".

Before contacting the Microsoft 365 team, you must perform the following steps for each Azure subscription that you use with Customer Key. Ensure that you have the Azure PowerShell Az module installed before you start.

1. Sign in with Azure PowerShell. For instructions, see Sign in with Azure PowerShell.
2. Run the Register-AzProviderFeature cmdlet to register your subscriptions to use a mandatory retention period. Perform this action for each subscription.

   ```
   Set-AzContext -SubscriptionId <SubscriptionId>
   Register-AzProviderFeature -FeatureName mandatoryRetentionPeriodEnabled -ProviderNamespace
   Microsoft.Resources
   ```

3. Contact Microsoft to have the process finalized at m365ck@microsoft.com. Include the following in your email:

   Subject: Customer Key for < *Your tenant's fully-qualified domain name*>

   Body: Subscription IDs for which you want to have the mandatory retention period finalized. The output of Get-AzProviderFeature for each subscription.

The Service Level Agreement (SLA) for completion of this process is five business days once Microsoft has been notified (and verified) that you have registered your subscriptions to use a mandatory retention period.

4. Once you receive notification from Microsoft that registration is complete, verify the status of your registration by running the Get-AzProviderFeature command as follows. If verified, the Get-AzProviderFeature command returns a value of **Registered** for the **Registration State** property. Perform this action for each subscription.

```
Set-AzContext -SubscriptionId <SubscriptionId>
Get-AzProviderFeature -ProviderNamespace Microsoft.Resources -FeatureName
mandatoryRetentionPeriodEnabled
```

5. To complete the process, run the Register-AzResourceProvider command. Perform this action for each subscription.

```
Set-AzContext -SubscriptionId <SubscriptionId>
Register-AzResourceProvider -ProviderNamespace Microsoft.KeyVault
```

Create a premium Azure Key Vault in each subscription

The steps to create a key vault are documented in Getting Started with Azure Key Vault, which guides you through installing and launching Azure PowerShell, connecting to your Azure subscription, creating a resource group, and creating a key vault in that resource group.

When you create a key vault, you must choose a SKU: either Standard or Premium. The Standard SKU allows Azure Key Vault keys to be protected with software - there is no Hardware Security Module (HSM) key protection - and the Premium SKU allows the use of HSMs for protection of Key Vault keys. Customer Key accepts key vaults that use either SKU, though Microsoft strongly recommends that you use only the Premium SKU. The cost of operations with keys of either type is the same, so the only difference in cost is the cost per month for each HSM-protected key. See Key Vault pricing for details.

> **IMPORTANT**
> Use the Premium SKU key vaults and HSM-protected keys for production data, and only use Standard SKU key vaults and keys for testing and validation purposes.

Use a common prefix for key vaults and include an abbreviation of the use and scope of the key vault and keys. For example, for the Contoso service where the vaults will be located in North America, a possible pair of names is Contoso-O365-NA-VaultA1 and Contoso-O365-NA-VaultA2. Vault names are globally unique strings within Azure, so you may need to try variations of your desired names in case the desired names are already claimed by other Azure customers. Once configured, vault names cannot be changed, so the best practice is to have a written plan for setup and use a second person to verify the plan is executed correctly.

If possible, create your vaults in non-paired regions. Paired Azure regions provide high availability across service failure domains. Therefore, regional pairs can be thought of as each other's backup region. This means that an Azure resource that is placed in one region is automatically gaining fault tolerance through the paired region. For this reason, choosing regions for two vaults used in a data encryption policy where the regions are paired means that only a total of two regions of availability are in use. Most geographies only have two regions, so it's not yet possible to select non-paired regions. If possible, choose two non-paired regions for the two vaults used with a data encryption policy. This benefits from a total of four regions of availability. For more information, see Business continuity and disaster recovery (BCDR): Azure Paired Regions for a current list of regional pairs.

Assign permissions to each key vault

For each key vault, you will need to define three separate sets of permissions for Customer Key, depending on your implementation. For example, you will need to define one set of permissions for each of the following:

- **Key vault administrators** that will perform day-to-day management of your key vault for your organization. These tasks include backup, create, get, import, list, and restore.

  > **IMPORTANT**
  > The set of permissions assigned to key vault administrators does not include the permission to delete keys. This is intentional and an important practice. Deleting encryption keys is not typically done, since doing so permanently destroys data. As a best practice, do not grant this permission to key vault administrators by default. Instead, reserve this for key vault contributors and only assign it to an administrator on a short term basis once a clear understanding of the consequences is understood.

  To assign these permissions to a user in your organization, log in to your Azure subscription with Azure PowerShell. For instructions, see Sign in with Azure PowerShell.

  Run the Set-AzKeyVaultAccessPolicy cmdlet to assign the necessary permissions.

  ```
  Set-AzKeyVaultAccessPolicy -VaultName <vault name> -UserPrincipalName <UPN of user> -
  PermissionsToKeys create,import,list,get,backup,restore
  ```

  For example:

  ```
  Set-AzKeyVaultAccessPolicy -VaultName Contoso-O365EX-NA-VaultA1 -UserPrincipalName alice@contoso.com
  -PermissionsToKeys create,import,list,get,backup,restore
  ```

- **Key vault contributors** that can change permissions on the Azure Key Vault itself. You'll need to change these permissions as employees leave or join your team, or in the rare situation that the key vault administrators legitimately need permission to delete or restore a key. This set of key vault contributors needs to be granted the Contributor role on your key vault. You can assign this role by using Azure Resource Manager. For detailed steps, see Use Role-Based Access Control to manage access to your Azure subscription resources. The administrator who creates a subscription has this access by default, and the ability to assign other administrators to the Contributor role.

- **Microsoft 365 data at rest encryption service** that does the work of Customer Key at the tenant level. To give permission to Microsoft 365, run the **Set-AzKeyVaultAccessPolicy** cmdlet using the following syntax:

  ```
  Set-AzKeyVaultAccessPolicy -VaultName <vault name> -PermissionsToKeys wrapKey,unwrapKey,get -
  ServicePrincipalName <Microsoft 365 appID>
  ```

  Where:

  - *vault name* is the name of the key vault you created.

  Example: For the Microsoft 365 Data at Rest Encryption service, replace *Microsoft 365 appID* with `c066d759-24ae-40e7-a56f-027002b5d3e4`

  ```
  Set-AzKeyVaultAccessPolicy -VaultName Contoso-O365EX-NA-VaultA1 -PermissionsToKeys
  wrapKey,unwrapKey,get -ServicePrincipalName c066d759-24ae-40e7-a56f-027002b5d3e4
  ```

Enable and then confirm soft delete on your key vaults

When you can quickly recover your keys, you are less likely to experience an extended service outage due to

accidentally or maliciously deleted keys. Enable this configuration, referred to as Soft Delete, before you can use your keys with Customer Key. Enabling Soft Delete allows you to recover keys or vaults within 90 days of deletion without having to restore them from backup.

To enable Soft Delete on your key vaults, complete these steps:

1. Sign in to your Azure subscription with Windows PowerShell. For instructions, see Sign in with Azure PowerShell.
2. Run the Get-AzKeyVault cmdlet. In this example, *vault name* is the name of the key vault for which you are enabling soft delete:

```
$v = Get-AzKeyVault -VaultName <vault name>
$r = Get-AzResource -ResourceId $v.ResourceId
$r.Properties | Add-Member -MemberType NoteProperty -Name enableSoftDelete -Value 'True'
Set-AzResource -ResourceId $r.ResourceId -Properties $r.Properties
```

3. Confirm soft delete is configured for the key vault by running the **Get-AzKeyVault** cmdlet. If soft delete is configured properly for the key vault, then the *Soft Delete Enabled* property returns a value of **True**:

```
Get-AzKeyVault -VaultName <vault name> | fl
```

Add a key to each key vault either by creating or importing a key

There are two ways to add keys to an Azure Key Vault; you can create a key directly in Key Vault, or you can import a key. Creating a key directly in Key Vault is the less complicated method, while importing a key provides total control over how the key is generated. Use the RSA keys. Azure Key Vault doesn't support wrapping and unwrapping with elliptical curve keys.

To create a key directly in your key vault, run the Add-AzKeyVaultKey cmdlet as follows:

```
Add-AzKeyVaultKey -VaultName <vault name> -Name <key name> -Destination <HSM|Software> -KeyOps
wrapKey,unwrapKey
```

Where:

- *vault name* is the name of the key vault in which you want to create the key.
- *key name* is the name you want to give the new key.

> **TIP**
> Name keys using a similar naming convention as described above for key vaults. This way, in tools that show only the key name, the string is self-describing.

If you intend to protect the key with an HSM, ensure that you specify **HSM** as the value of the *Destination* parameter, otherwise, specify **Software**.

For example,

```
Add-AzKeyVaultKey -VaultName Contoso-O365EX-NA-VaultA1 -Name Contoso-O365EX-NA-VaultA1-Key001 -Destination
HSM -KeyOps wrapKey,unwrapKey
```

Check the recovery level of your keys

Microsoft 365 requires that the Azure Key Vault subscription is set to Do Not Cancel and that the keys used by

Customer Key have soft delete enabled. You can confirm this by looking at the recovery level on your keys.

To check the recovery level of a key, in Azure PowerShell, run the Get-AzKeyVaultKey cmdlet as follows:

```
(Get-AzKeyVaultKey -VaultName <vault name> -Name <key name>).Attributes
```

If the *Recovery Level* property returns anything other than a value of **Recoverable+ProtectedSubscription**, you will need to review this article and ensure that you have followed all of the steps to put the subscription on the Do Not Cancel list and that you enabled "soft delete" on each of your key vaults. Next, send a screenshot of the output of `(Get-AzKeyVaultKey -VaultName <vault name> -Name <key name>).Attributes` in email to m365ck@microsoft.com.

Back up Azure Key Vault

Immediately following creation or any change to a key, perform a backup and store copies of the backup, both online and offline. Don't connect offline copies to any network. Instead, store them in a physical safe or commercial storage facility. At least one copy of the backup should be stored in a location that will be accessible if a disaster happens. The backup blobs are the sole means of restoring key material should a Key Vault key be permanently destroyed or otherwise rendered inoperable. Keys that are external to Azure Key Vault and were imported to Azure Key Vault do not qualify as a backup because the metadata necessary for Customer Key to use the key does not exist with the external key. Only a backup taken from Azure Key Vault can be used for restore operations with Customer Key. Therefore, it is essential that you make a backup of Azure Key Vault once a key is uploaded or created.

To create a backup of an Azure Key Vault key, run the Backup-AzKeyVaultKey cmdlet as follows:

```
Backup-AzKeyVaultKey -VaultName <vault name> -Name <key name>
-OutputFile <filename.backup>
```

Ensure that your output file uses the suffix `.backup`.

The output file resulting from this cmdlet is encrypted and cannot be used outside of Azure Key Vault. The backup can be restored only to the Azure subscription from which the backup was taken.

For example:

```
Backup-AzKeyVaultKey -VaultName Contoso-O365EX-NA-VaultA1 -Name Contoso-O365EX-NA-VaultA1-Key001 -OutputFile
Contoso-O365EX-NA-VaultA1-Key001-Backup-20170802.backup
```

Validate Azure Key Vault configuration settings

Performing validation before using keys in a DEP is optional, but highly recommended. In particular, if you use steps to set up your keys and vaults other than the ones described in this topic, you should validate the health of your Azure Key Vault resources before you configure Customer Key.

To verify that your keys have get, wrapKey, and unwrapKey operations enabled:

Run the Get-AzKeyVault cmdlet as follows:

```
Get-AzKeyVault -VaultName <vault name>
```

In the output, look for the Access Policy and for the Microsoft 365 app ID (GUID) as appropriate. All three operations, get, wrapKey, and unwrapKey, must be shown under Permissions to Keys.

If the access policy configuration is incorrect, run the Set-AzKeyVaultAccessPolicy cmdlet as follows:

```
Set-AzKeyVaultAccessPolicy -VaultName <vault name> -PermissionsToKeys wrapKey,unwrapKey,get -
ServicePrincipalName <Microsoft 365 appID>
```

Example: For the Microsoft 365 Data at Rest Encryption service, replace *Microsoft 365 appID* with `c066d759-24ae-40e7-a56f-027002b5d3e4`

```
Set-AzKeyVaultAccessPolicy -VaultName Contoso-O365EX-NA-VaultA1 -PermissionsToKeys wrapKey,unwrapKey,get -
ServicePrincipalName c066d759-24ae-40e7-a56f-027002b5d3e4
```

To verify that an expiration date is not set for your keys, run the Get-AzKeyVaultKey cmdlet as follows:

```
Get-AzKeyVaultKey -VaultName <vault name>
```

An expired key cannot be used by Customer Key and operations attempted with an expired key will fail and possibly result in a service outage. We strongly recommend that keys used with Customer Key do not have an expiration date. An expiration date, once set, cannot be removed, but can be changed to a different date. If a key must be used that has an expiration date set, change the expiration value to 12/31/9999. Keys with an expiration date set to a date other than 12/31/9999 will not pass Microsoft 365 validation.

To change an expiration date that has been set to any value other than 12/31/9999, run the Update-AzKeyVaultKey cmdlet as follows:

```
Update-AzKeyVaultKey -VaultName <vault name> -Name <key name> -Expires (Get-Date -Date "12/31/9999")
```

Obtain the URI for each Azure Key Vault key

Once you've completed all the steps in Azure to set up your key vaults and added your keys, run the following command to get the URI for the key in each key vault. You will need to use these URIs when you create and assign each DEP later, so save this information in a safe place. Remember to run this command once for each key vault.

In Azure PowerShell:

```
(Get-AzKeyVaultKey -VaultName <vault name>).Id
```

Set up the Customer Key encryption policy for your tenant

You need to be assigned permissions before you can run these cmdlets. Although this article lists all parameters for the cmdlets, you may not have access to some parameters if they're not included in the permissions assigned to you. To find the permissions required to run any cmdlet or parameter in your organization, see Find the permissions required to run any Exchange cmdlet.

Create policy

```
   New-M365DataAtRestEncryptionPolicy [-Name] <String> -AzureKeyIDs <MultiValuedProperty> [-Description
<String>] [-Enabled <Boolean>]
```

Description: Enable compliance admin to create a new data encryption policy (DEP) using two AKV root keys. Once created, a policy can then be assigned using Set-M365DataAtRestEncryptionPolicy cmdlet. Upon first assignment of keys or after you rotate keys, it can take up to 24 hours for the new keys to take effect. If the new DEP takes more than 24 hours to take effect, contact Microsoft.

Example:

```
New-M365DataAtRestEncryptionPolicy -Name "Default_Policy" -AzureKeyIDs
"https://contosoWestUSvault01.vault.azure.net/keys/Key_01","https://contosoEastUSvault01.vault.azure.net/key
s/Key_02" -Description "Tenant default policy"
```

Parameters:

NAME	DESCRIPTION	OPTIONAL (Y/N)
Name	Friendly name of the data encryption policy	N
AzureKeyIDs	Specifies two URI values of the Azure Key Vault keys, separated by a comma, to associate with the data encryption policy	N
Description	Description of the data encryption policy	N

Assign policy

```
Set-M365DataAtRestEncryptionPolicyAssignment -Policy “<Default_PolicyName or Default_PolicyID>”
```

Description: This cmdlet is used for configuring default Data Encryption Policy. This policy will be used to then encrypt data across all support workloads.

Example:

```
Set-M365DataAtRestEncryptionPolicyAssignment -Policy “Tenant default policy”
```

Parameters: | Name | Description | Optional (Y/N) | |----------|----------|---------| -Policy|Specifies the data encryption policy that needs to be assigned; specify either the Policy Name or the Policy ID.|N|

Modify or Refresh policy

```
Set-M365DataAtRestEncryptionPolicy [-Identity] < M365DataAtRestEncryptionPolicy
DataEncryptionPolicyIdParameter> -Refresh [-Enabled <Boolean>] [-Name <String>] [-Description <String>]
```

Description: The cmdlet can be used either to modify or refresh an existing policy. It can also be used to enable or disable a policy. Upon first assignment of keys or after you rotate keys, it can take up to 24 hours for the new keys to take effect. If the new DEP takes more than 24 hours to take effect, contact Microsoft.

Examples:

Disable a data encryption policy.

```
Set-M365DataAtRestEncryptionPolicy -Identity "NAM Policy" -Enabled $false
```

Refresh a data encryption policy.

```
Set-M365DataAtRestEncryptionPolicy -Identity “EUR Policy” -Refresh
```

Parameters: | Name | Description | Optional (Y/N) | |----------|----------|---------| |-Identity|Specifies the data encryption policy that you want to modify.|N| |-Refresh|Use the Refresh switch to update the data encryption policy after you rotate any of the associated keys in the Azure Key Vault. You don't need to specify a value with this switch.|Y| |-Enabled|The Enabled parameter enables or disable the data encryption policy. Before you disable a policy, you must unassign it from your tenant. Valid values are:
$true: The policy is enabled
$false: The policy is disabled.|Y| |-Name|The Name parameter specifies the unique name for the data encryption policy.|Y |-Description|The Description parameter specifies an optional description for the data encryption policy.|Y|

Get policy details

```
Get-M365DataAtRestEncryptionPolicy [-Identity] < M365DataAtRestEncryptionPolicy
DataEncryptionPolicyIdParameter>
```

Description: This cmdlet lists all of M365DataAtRest encryption policies that are created for the tenant or details about a specific policy.

Examples:

This example returns a summary list of M365DatAtRest Encryption policies in the organization.

```
Get-M365DataAtRestEncryptionPolicy
```

This example returns detailed information for the data encryption policy named "NAM Policy".

```
Get-M365DataAtRestEncryptionPolicy -Identity "NAM Policy"
```

Parameters:

NAME	DESCRIPTION	OPTIONAL (Y/N)
-Identity	Specifies the data encryption policy that you want to list the details for.	Y

Get policy assignment info

```
Get-M365DataAtRestEncryptionPolicyAssignment
```

Description: This cmdlet lists the policy that’s currently assigned to the tenant.

Offboarding from Customer Key

If you need to revert back to Microsoft-managed keys, you can. When you offboard, your data is re-encrypted using default encryption supported by each individual workload. For example, Exchange Online supports default encryption using Microsoft-managed keys.

If you decided to offboard your tenant from Customer Key at the tenant level, reach out to Microsoft with a request through email to “disable” the service for the tenant at m365ck@microsoft.com.

IMPORTANT

Offboarding is not the same as a data purge. A data purge permanently crypto-deletes your organization's data from Microsoft 365, offboarding does not. You can't perform a data purge for a tenant-level policy. For information about data purge path, see Revoke your keys and start the data purge path process.

About the availability key

For information about the availability key, see Learn about the availability key.

Key rotation

For information about rotating or rolling keys used with Customer Key, see Roll or rotate a Customer Key or an availability key. When you update the DEP to use the new version of the keys, you'll run the Set-M365DataAtRestEncryptionPolicy cmdlet as described earlier in this article.

Related articles:

- Service encryption with Customer Key
- Roll or rotate a Customer Key or an availability key
- Learn about the availability key
- Service Encryption

Service encryption with Customer Key

2/18/2021 • 4 minutes to read • Edit Online

Microsoft 365 provides baseline, volume-level encryption enabled through BitLocker and Distributed Key Manager (DKM). Microsoft 365 offers an added layer of encryption at the application layer for your content. This content includes data from Exchange Online, Skype for Business,
SharePoint Online, OneDrive for Business, and Teams files. This added layer of encryption is called service encryption.

How service encryption, BitLocker, and Customer Key work together

Service encryption ensures that content at rest is encrypted at the service layer. **Your data is always encrypted at rest in the Microsoft 365 service with BitLocker and DKM**. For more information, see the "Security, Privacy, and Compliance Information", and How Exchange Online secures your email secrets.
Customer Key provides additional protection against viewing of data by unauthorized systems or personnel, and complements BitLocker disk encryption in Microsoft datacenters. Service encryption is not meant to prevent Microsoft personnel from accessing customer data. The primary purpose is to assist customers in meeting regulatory or compliance obligations for controlling root keys. Customers explicitly authorize O365 services to use their encryption keys to provide value added cloud services, such as eDiscovery, anti-malware, anti-spam, search indexing, etc.

Customer Key is built on service encryption and lets you provide and control encryption keys. Microsoft 365 then uses these keys to encrypt your data at rest as described in the Online Services Terms (OST). Customer Key helps you meet compliance obligations because you control the encryption keys that Microsoft 365 uses to encrypt and decrypt data.

Customer Key enhances the ability of your organization to meet the demands of compliance requirements that specify key arrangements with the cloud service provider. With Customer Key, you provide and control the root encryption keys for your Microsoft 365 data at-rest at the application level. As a result, you exercise control over your organization's keys. If you decide to exit the service, you revoke access to your organization's root keys. For all Microsoft 365 services, revoking access to the keys is the first step on the path towards data deletion. By revoking access to the keys, the data is unreadable to the service.

Customer Key encrypts data at rest in Office 365

Using keys you provide, Customer Key encrypts:

- SharePoint Online, OneDrive for Business, and Teams files.
- Files uploaded to OneDrive for Business.
- Exchange Online mailbox content including e-mail body content, calendar entries, and the content within email attachments.
- Text conversations from Skype for Business.

We don't currently offer customer control of the encryption keys for Skype Meeting Broadcast and Skype Meeting content uploads. Instead, this content is encrypted along with all other content in Office 365.

Customer Key with hybrid deployments

Customer Key only encrypts data at rest in the cloud. Customer Key does not work to protect your on-premises mailboxes and files. You can encrypt your on-premises data using another method, such as BitLocker.

About the data encryption policy (DEP)

A data encryption policy defines the encryption hierarchy to encrypt data using each of the keys you provide as well as the availability key protected by Microsoft. You create DEPs using PowerShell cmdlets, which are different for each service, and assign those DEPs to encrypt application data. For example:

Exchange Online and Skype for Business You can create up to 50 DEPs per tenant. You associate DEPs to your Customer Keys in Azure Key Vault and then assign DEPs to individual mailboxes. When you assign a DEP to a mailbox:

- the mailbox is marked for a mailbox move. Based on priorities in Microsoft 365 as described here Move requests in the Microsoft 365 service.
- The encryption takes place while the mailbox is moved. Allow 72 hours for the mailbox to become encrypted with the new DEP. If the mailboxes aren't encrypted after waiting 72 hours from the time you assigned the DEP, contact Microsoft.

Later, you can either refresh the DEP or assign a different DEP to the mailbox as described in Manage Customer Key for Office 365. Each mailbox must have appropriate licenses in order to assign a DEP. For more information about licensing, see Before you set up Customer Key.

> **NOTE**
> The DEP can be applied to a shared mailbox, public folder mailbox, and Microsoft 365 group mailbox for tenants that meet the licensing requirement for user mailboxes, even though some of these mailbox types cannot be an assigned license (public folder mailbox and Microsoft 365 group mailbox) or need a license for increasing storage (shared mailbox).

SharePoint Online, OneDrive for Business, and Teams files If you're using the multi-geo feature, you can create up to one DEP per geo for your organization. You can use different Customer Keys for each geo. If you're not using the multi-geo feature, you can only create one DEP per tenant. When you assign the DEP, encryption begins automatically but can take some time to complete. Refer to the details in Set up Customer Key.

Leaving the service

Customer Key assists you in meeting compliance obligations by allowing you to revoke your keys when you leave the Microsoft 365 service. When you revoke your keys as part of leaving the service, the availability key is deleted resulting in cryptographic deletion of your data. Cryptographic deletion mitigates the risk of data remanence which is important for meeting both security and compliance obligations. For information about the data purge process and key revocation, see Revoke your keys and start the data purge path process.

Encryption ciphers used by Customer Key

Customer Key uses a variety of encryption ciphers to encrypt keys as shown in the following figures.

Encryption ciphers used to encrypt keys for Exchange Online and Skype for Business

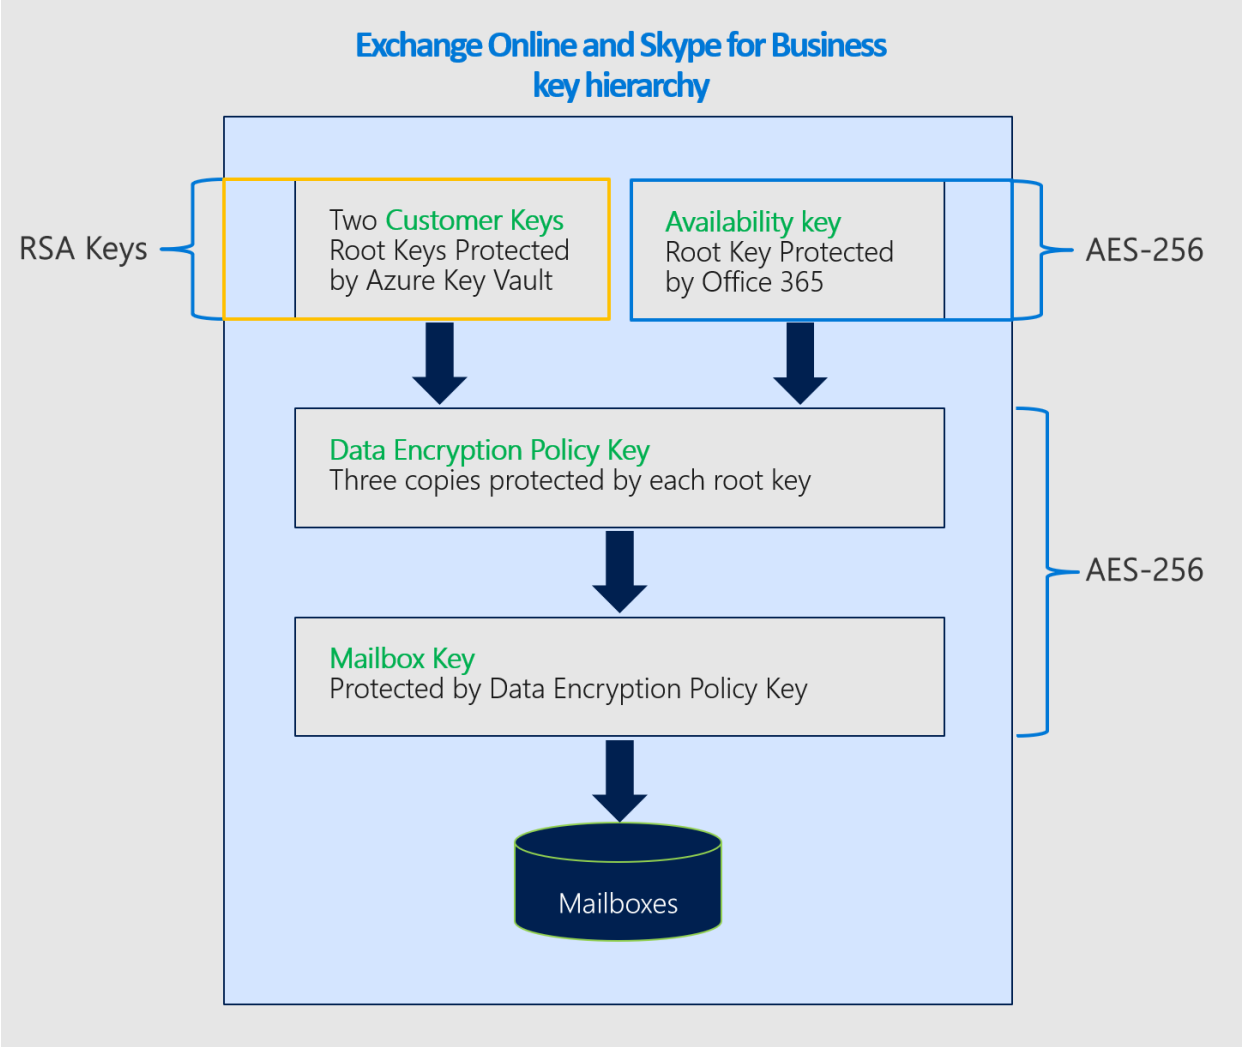

Encryption ciphers used to encrypt keys for SharePoint Online, OneDrive for Business, and Teams files

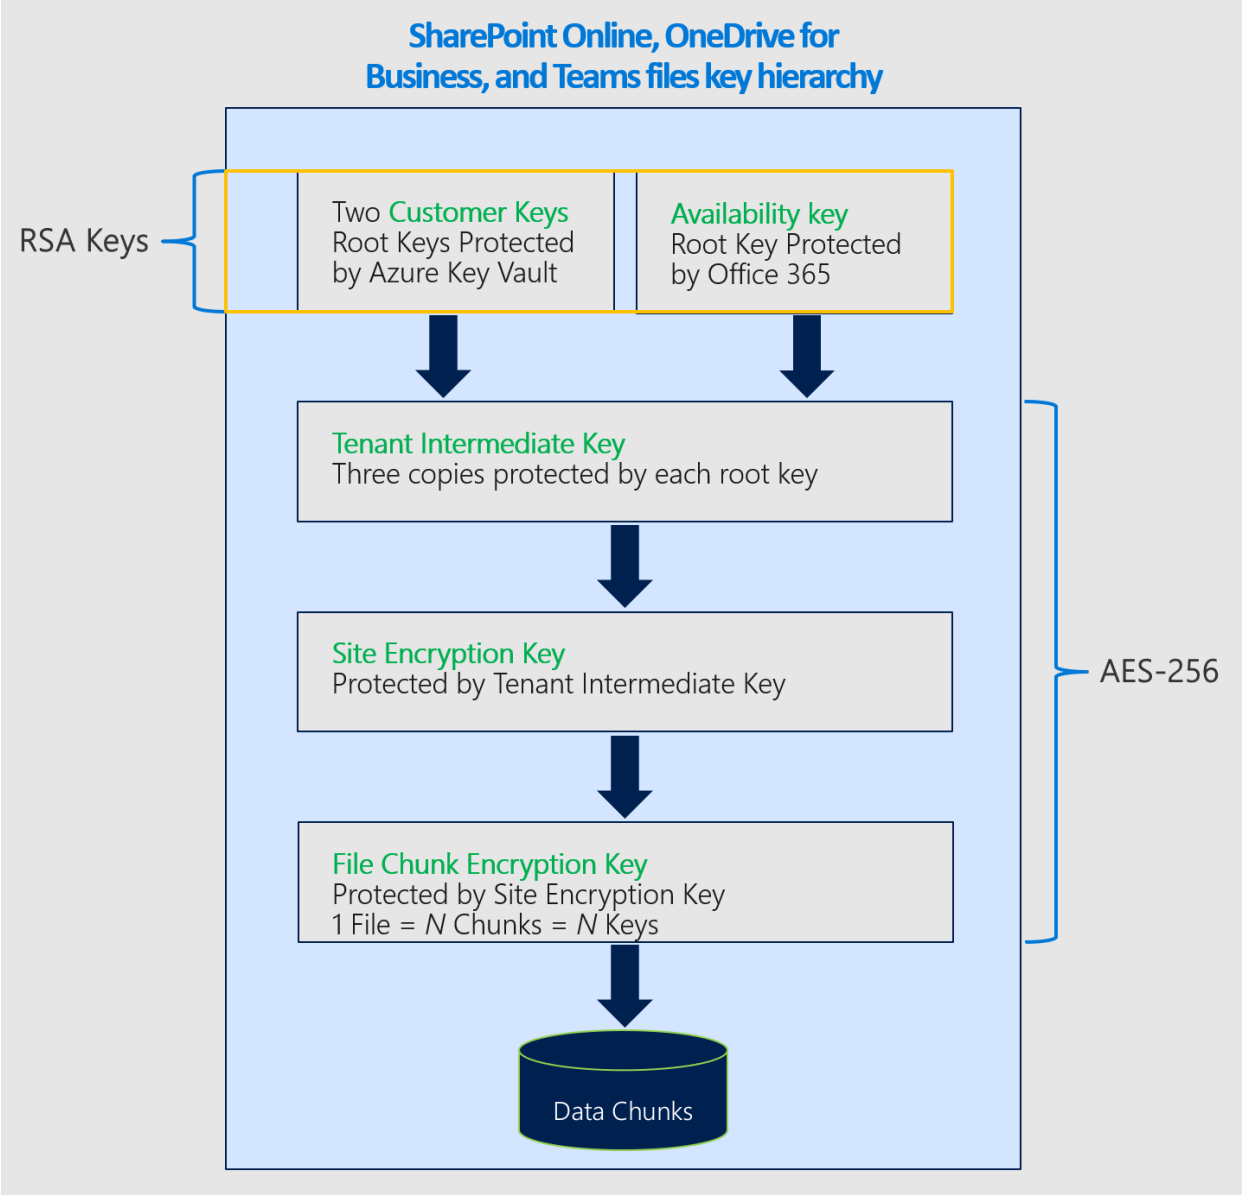

Related articles

- Set up Customer Key
- Manage Customer Key
- Roll or rotate a Customer Key or an availability key
- Learn about the availability key
- Customer Lockbox
- Service Encryption

Set up Customer Key at the application level

2/18/2021 • 26 minutes to read • Edit Online

With Customer Key, you control your organization's encryption keys and then configure Microsoft 365 to use them to encrypt your data at rest in Microsoft's data centers. In other words, Customer Key allows customers to add a layer of encryption that belongs to them, with their keys. Data at rest includes data from Exchange Online and Skype for Business that is stored in mailboxes and files that are stored in SharePoint Online and OneDrive for Business.

You must set up Azure before you can use Customer Key for Office 365. This article describes the steps you need to follow to create and configure the required Azure resources and then provides the steps for setting up Customer Key in Office 365. After you have completed Azure setup, you determine which policy, and therefore, which keys, to assign to mailboxes and files in your organization. Mailboxes and files for which you don't assign a policy will use encryption policies that are controlled and managed by Microsoft. For more information about Customer Key, or for a general overview, see Service encryption with Customer Key in Office 365.

IMPORTANT

We strongly recommend that you follow the best practices in this article. These are called out as **TIP** and **IMPORTANT**. Customer Key gives you control over root encryption keys whose scope can be as large as your entire organization. This means that mistakes made with these keys can have a broad impact and may result in service interruptions or irrevocable loss of your data.

Before you set up Customer Key

Before you get started, ensure that you have the appropriate licensing for your organization. Use a paid, invoiced Azure Subscription using either an Enterprise Agreement or a Cloud Service Provider. Azure Subscriptions purchased using Pay As You Go plans or using a credit card aren't supported for Customer Key. Starting April 1, 2020, Customer Key in Office 365 is offered in Office 365 E5, M365 E5, M365 E5 Compliance, and M365 E5 Information Protection & Governance SKUs. Office 365 Advanced Compliance SKU is no longer available for procuring new licenses. Existing Office 365 Advanced Compliance licenses will continue to be supported.

To understand the concepts and procedures in this article, review the Azure Key Vault documentation. Also, become familiar with the terms used in Azure, for example, Azure AD tenant.

FastTrack is only used to collect the required tenant and service configuration information used to register for Customer Key. The Customer Key Offers are published via FastTrack so that it is convenient for you and our partners to submit the required information using the same method. FastTrack also makes it easy to archive the data that you provide in the Offer.

If you need more support beyond the documentation, contact Microsoft Consulting Services (MCS), Premier Field Engineering (PFE), or a Microsoft partner for assistance. To provide feedback on Customer Key, including the documentation, send your ideas, suggestions, and perspectives to customerkeyfeedback@microsoft.com.

Overview of steps to set up Customer Key

To set up Customer Key, complete these tasks in the listed order. The rest of this article provides detailed instructions for each task, or links out to more information for each step in the process.

In Azure and Microsoft FastTrack:

You will complete most of these tasks by remotely connecting to Azure PowerShell. For best results, use version 4.4.0 or later of Azure PowerShell.

- Create two new Azure subscriptions
- Register Azure subscriptions to use a mandatory retention period

  Registration can take from one to five business days.
- Submit a request to activate Customer Key for Office 365

Once you've created the two new Azure subscriptions, you'll need to submit the appropriate Customer Key offer request by completing a web form that is hosted in the Microsoft FastTrack portal. **The FastTrack team doesn't provide assistance with Customer Key. Office simply uses the FastTrack portal to allow you to submit the form and to help us track the relevant offers for Customer Key**.

- Create a premium Azure Key Vault in each subscription
- Assign permissions to each key vault
- Enable and then confirm soft delete on your key vaults
- Add a key to each key vault either by creating or importing a key
- Check the recovery level of your keys
- Back up Azure Key Vault
- Validate Azure Key Vault configuration settings
- Obtain the URI for each Azure Key Vault key

In Office 365:

Exchange Online and Skype for Business:

- Create a data encryption policy (DEP) for use with Exchange Online and Skype for Business
- Assign a DEP to a mailbox
- Validate mailbox encryption

SharePoint Online and OneDrive for Business:

- Create a data encryption policy (DEP) for each SharePoint Online and OneDrive for Business geo
- Validate file encryption for SharePoint Online, OneDrive for Business, and Teams files

Complete tasks in Azure Key Vault and Microsoft FastTrack for Customer Key

Complete these tasks in Azure Key Vault. You'll need to complete these steps regardless of whether you intend to set up Customer Key for Exchange Online and Skype for Business or for SharePoint Online, OneDrive for Business, and Teams files, or for all supported services in Office 365.

Create two new Azure subscriptions

Customer Key requires two Azure subscriptions. As a best practice, Microsoft recommends that you create new Azure subscriptions for use with Customer Key. Azure Key Vault keys can only be authorized for applications in the same Azure Active Directory (Microsoft Azure Active Directory) tenant, you must create the new subscriptions using the same Azure AD tenant used with your organization where the DEPs will be assigned. For example, using your work or school account that has global administrator privileges in your organization. For

detailed steps, see Sign up for Azure as an organization.

> **IMPORTANT**
> Customer Key requires two keys for each data encryption policy (DEP). In order to achieve this, you must create two Azure subscriptions. As a best practice, Microsoft recommends that you have separate members of your organization configure one key in each subscription. You should only use these Azure subscriptions to administer encryption keys for Office 365. This protects your organization in case one of your operators accidentally, intentionally, or maliciously deletes or otherwise mismanages the keys for which they are responsible.

There is no practical limit to the number of Azure subscriptions that you can create for your organization. Following these best practices will minimize the impact of human error while helping to manage the resources used by Customer Key.

Submit a request to activate Customer Key for Office 365

Once you've completed the Azure steps, you'll need to submit an offer request in the Microsoft FastTrack portal. Once you've submitted a request through the FastTrack web portal, Microsoft verifies the Azure Key Vault configuration data and contact information you provided. The selections that you make in the offer form about the authorized officers of your organization is critical and necessary for completion of Customer Key registration. The officers of your organization ensure the authenticity of any request to revoke and destroy all keys used with a Customer Key data encryption policy. You'll need to do this step once to activate Customer Key for Exchange Online and Skype for Business coverage and a second time to activate Customer Key for SharePoint Online and OneDrive for Business.

To submit an offer to activate Customer Key, complete these steps:

1. Using a work or school account that has global administrator permissions in your organization, sign in to the Microsoft FastTrack portal.
2. Once you're logged in, browse to the **Dashboard**.
3. Choose **Deploy** from the navigation bar **OR** select **View all deployment resources** on the **Deploy** information card, and review the list of current offers.
4. Choose the information card for the offer that applies to you:
   - **Exchange Online and Skype for Business:** Choose the **Request encryption key help for Exchange online** offer.
   - **SharePoint Online, OneDrive, and Teams files:** Choose the **Request encryption key help for Sharepoint and OneDrive** offer.
5. Once you've reviewed the offer details, choose **Continue to step 2**.
6. Fill out all applicable details and requested information on the offer form. Pay particular attention to your selections for which officers of your organization you want to authorize to approve the permanent and irreversible destruction of encryption keys and data. Once you've completed the form, choose **Submit**.

Register Azure subscriptions to use a mandatory retention period

The temporary or permanent loss of root encryption keys can be disruptive or even catastrophic to service operation and can result in data loss. For this reason, the resources used with Customer Key require strong protection. All the Azure resources that are used with Customer Key offer protection mechanisms beyond the default configuration. You can tag or register Azure subscriptions for a *mandatory retention period*. A mandatory retention period prevents immediate and irrevocable cancellation of your Azure subscription. The steps required to register Azure subscriptions for a mandatory retention period require collaboration with the Microsoft 365 team. This process can take from one to five business days. Previously, mandatory retention period was sometimes referred to as "Do Not Cancel".

Before contacting the Microsoft 365 team, you must do the following steps for each Azure subscription that you use with Customer Key. Ensure that you have the Azure PowerShell Az module installed before you start.

1. Sign in with Azure PowerShell. For instructions, see Sign in with Azure PowerShell.
2. Run the Register-AzProviderFeature cmdlet to register your subscriptions to use a mandatory retention period. Complete this action for each subscription.

   ```
   Set-AzContext -SubscriptionId <SubscriptionId>
   Register-AzProviderFeature -FeatureName mandatoryRetentionPeriodEnabled -ProviderNamespace
   Microsoft.Resources
   ```

3. Contact Microsoft to complete the process. For the SharePoint and OneDrive for Business team, contact spock@microsoft.com. For Exchange Online and Skype for Business, contact exock@microsoft.com. Include the following information in your email:

   Subject: Customer Key for *< Your tenant's fully qualified domain name>*

   Body: Include the subscription IDs for which you want to complete the mandatory retention period and the output of Get-AzProviderFeature for each subscription.

   The Service Level Agreement (SLA) for completion of this process is five business days once Microsoft has been notified (and verified) that you have registered your subscriptions to use a mandatory retention period.

4. Once you receive notification from Microsoft that registration is complete, verify the status of your registration by running the Get-AzProviderFeature command as follows. If verified, the Get-AzProviderFeature command returns a value of **Registered** for the **Registration State** property. Complete this step for each subscription.

   ```
   Set-AzContext -SubscriptionId <SubscriptionId>
   Get-AzProviderFeature -ProviderNamespace Microsoft.Resources -FeatureName
   mandatoryRetentionPeriodEnabled
   ```

5. To complete the process, run the Register-AzResourceProvider command. Complete this step for each subscription.

   ```
   Set-AzContext -SubscriptionId <SubscriptionId>
   Register-AzResourceProvider -ProviderNamespace Microsoft.KeyVault
   ```

Create a premium Azure Key Vault in each subscription

The steps to create a key vault are documented in Getting Started with Azure Key Vault, which guides you through installing and launching Azure PowerShell, connecting to your Azure subscription, creating a resource group, and creating a key vault in that resource group.

When you create a key vault, you must choose a SKU: either Standard or Premium. The Standard SKU allows Azure Key Vault keys to be protected with software - there's no Hardware Security Module (HSM) key protection - and the Premium SKU allows the use of HSMs for protection of Key Vault keys. Customer Key accepts key vaults that use either SKU, though Microsoft strongly recommends that you use only the Premium SKU. The cost of operations with keys of either type is the same, so the only difference in cost is the cost per month for each HSM-protected key. See Key Vault pricing for details.

> **IMPORTANT**
>
> Use the Premium SKU key vaults and HSM-protected keys for production data, and only use Standard SKU key vaults and keys for testing and validation purposes.

For each Microsoft 365 service with which you will use Customer Key, create a key vault in each of the two Azure subscriptions that you created. For example, for Exchange Online and Skype for Business only or SharePoint Online and OneDrive for Business only, you'll create only one pair of vaults. To enable Customer Key for both Exchange Online and SharePoint Online, you will create two pairs of key vaults.

Use a naming convention for key vaults that reflects the intended use of the DEP with which you will associate the vaults. See the Best Practices section below for naming convention recommendations.

Create a separate, paired set of vaults for each data encryption policy. For Exchange Online, the scope of a data encryption policy is chosen by you when you assign the policy to mailbox. A mailbox can have only one policy assigned, and you can create up to 50 policies. The scope of a SharePoint Online policy includes all of the data within an organization in a geographic location, or *geo*.

The creation of key vaults also requires the creation of Azure resource groups, since key vaults need storage capacity (though small) and Key Vault logging, if enabled, also generates stored data. As a best practice Microsoft recommends using separate administrators to manage each resource group, with the administration that's aligned with the set of administrators that will manage all related Customer Key resources.

> **IMPORTANT**
>
> To maximize availability, your key vaults should be in regions close to your Microsoft 365 service. For example, if your Exchange Online organization is in North America, place your key vaults in North America. If your Exchange Online organization is in Europe, place your key vaults in Europe.
>
> Use a common prefix for key vaults, and include an abbreviation of the use and scope of the key vault and keys (e.g., for the Contoso SharePoint service where the vaults will be located in North America, a possible pair of names is Contoso-O365SP-NA-VaultA1 and Contoso-O365SP-NA-VaultA2. Vault names are globally unique strings within Azure, so you may need to try variations of your desired names in case the desired names are already claimed by other Azure customers. As of July 2017 vault names cannot be changed, so a best practice is to have a written plan for setup and use a second person to verify the plan is executed correctly.
>
> If possible, create your vaults in non-paired regions. Paired Azure regions provide high availability across service failure domains. Therefore, regional pairs can be thought of as each other's backup region. This means that an Azure resource that is placed in one region is automatically gaining fault tolerance through the paired region. For this reason, choosing regions for two vaults used in a data encryption policy where the regions are paired means that only a total of two regions of availability are in use. Most geographies only have two regions, so it's not yet possible to select non-paired regions. If possible, choose two non-paired regions for the two vaults used with a data encryption policy. This benefits from a total of four regions of availability. For more information, see Business continuity and disaster recovery (BCDR): Azure Paired Regions for a current list of regional pairs.

Assign permissions to each key vault

You'll need to define three separate sets of permissions for each key vault, depending on your implementation. For example, you will need to define one set of permissions for each of the following:

- **Key vault administrators** that do day-to-day management of your key vault for your organization. These tasks include backup, create, get, import, list, and restore.

> **IMPORTANT**
> The set of permissions assigned to key vault administrators does not include the permission to delete keys. This is intentional and an important practice. Deleting encryption keys is not typically done, since doing so permanently destroys data. As a best practice, do not grant this permission to key vault administrators by default. Instead, reserve this for key vault contributors and only assign it to an administrator on a short term basis once a clear understanding of the consequences is understood.

To assign these permissions to a user in your organization, sign in to your Azure subscription with Azure PowerShell. For instructions, see Sign in with Azure PowerShell.

- Run the Set-AzKeyVaultAccessPolicy cmdlet to assign the necessary permissions.

```
Set-AzKeyVaultAccessPolicy -VaultName <vault name> -UserPrincipalName <UPN of user> -
PermissionsToKeys create,import,list,get,backup,restore
```

For example:

```
Set-AzKeyVaultAccessPolicy -VaultName Contoso-O365EX-NA-VaultA1 -UserPrincipalName alice@contoso.com
-PermissionsToKeys create,import,list,get,backup,restore
```

- **Key vault contributors** that can change permissions on the Azure Key Vault itself. You'll need to change these permissions as employees leave or join your team. In the rare situation that the key vault administrators legitimately need permission to delete or restore a key you'll also need to change the permissions. This set of key vault contributors needs to be granted the **Contributor** role on your key vault. You can assign this role by using Azure Resource Manager. For detailed steps, see Use Role-Based Access Control to manage access to your Azure subscription resources. The administrator who creates a subscription has this access implicitly, and the ability to assign other administrators to the Contributor role.

- If you intend to use Customer Key with Exchange Online and Skype for Business, you need to give permission to Microsoft 365 to use the key vault on behalf of Exchange Online and Skype for Business. Likewise, if you intend to use Customer Key with SharePoint Online and OneDrive for Business, you need to add permission for the Microsoft 365 to use the key vault on behalf of SharePoint Online and OneDrive for Business. To give permission to Microsoft 365, run the **Set-AzKeyVaultAccessPolicy** cmdlet using the following syntax:

```
Set-AzKeyVaultAccessPolicy -VaultName <vault name> -PermissionsToKeys wrapKey,unwrapKey,get -
ServicePrincipalName <Office 365 appID>
```

Where:

  - *vault name* is the name of the key vault you created.
  - For Exchange Online and Skype for Business, replace *Office 365 appID* with `00000002-0000-0ff1-ce00-000000000000`
  - For SharePoint Online, OneDrive for Business, and Teams files, replace *Office 365 appID* with `00000003-0000-0ff1-ce00-000000000000`

Example: Setting permissions for Exchange Online and Skype for Business:

```
Set-AzKeyVaultAccessPolicy -VaultName Contoso-O365EX-NA-VaultA1 -PermissionsToKeys
wrapKey,unwrapKey,get -ServicePrincipalName 00000002-0000-0ff1-ce00-000000000000
```

Example: Setting permissions for SharePoint Online, OneDrive for Business, and Teams files:

```
Set-AzKeyVaultAccessPolicy -VaultName Contoso-O365SP-NA-VaultA1 -PermissionsToKeys
wrapKey,unwrapKey,get -ServicePrincipalName 00000003-0000-0ff1-ce00-000000000000
```

Enable and then confirm soft delete on your key vaults

When you can quickly recover your keys, you are less likely to experience an extended service outage due to accidentally or maliciously deleted keys. You need to enable this configuration, referred to as Soft Delete, before you can use your keys with Customer Key. Enabling Soft Delete allows you to recover keys or vaults within 90 days of deletion without having to restore them from backup.

To enable Soft Delete on your key vaults, complete these steps:

1. Sign in to your Azure subscription with Windows PowerShell. For instructions, see Sign in with Azure PowerShell.
2. Run the Get-AzKeyVault cmdlet. In this example, *vault name* is the name of the key vault for which you are enabling soft delete:

```
$v = Get-AzKeyVault -VaultName <vault name>
$r = Get-AzResource -ResourceId $v.ResourceId
$r.Properties | Add-Member -MemberType NoteProperty -Name enableSoftDelete -Value 'True'
Set-AzResource -ResourceId $r.ResourceId -Properties $r.Properties
```

3. Confirm soft delete is configured for the key vault by running the **Get-AzKeyVault** cmdlet. If soft delete is configured properly for the key vault, then the *Soft Delete Enabled* property returns a value of **True**:

```
Get-AzKeyVault -VaultName <vault name> | fl
```

Add a key to each key vault either by creating or importing a key

There are two ways to add keys to an Azure Key Vault; you can create a key directly in Key Vault, or you can import a key. Creating a key directly in Key Vault is the less complicated method, while importing a key provides total control over how the key is generated. Use the RSA keys. Azure Key Vault doesn't support wrapping and unwrapping with elliptical curve keys.

To create a key directly in your key vault, run the Add-AzKeyVaultKey cmdlet as follows:

```
Add-AzKeyVaultKey -VaultName <vault name> -Name <key name> -Destination <HSM|Software> -KeyOps
wrapKey,unwrapKey
```

Where:

- *vault name* is the name of the key vault in which you want to create the key.
- *key name* is the name you want to give the new key.

> **TIP**
> Name keys using a similar naming convention as described above for key vaults. This way, in tools that show only the key name, the string is self-describing.

If you intend to protect the key with an HSM, ensure that you specify **HSM** as the value of the *Destination* parameter, otherwise, specify **Software**.

For example,

```
Add-AzKeyVaultKey -VaultName Contoso-O365EX-NA-VaultA1 -Name Contoso-O365EX-NA-VaultA1-Key001 -Destination
HSM -KeyOps wrapKey,unwrapKey
```

To import a key directly into your key vault, you need to have a nCipher nShield Hardware Security Module.

Some organizations prefer this approach to establish the provenance of their keys, and then this method also provides the following attestations:

- The toolset used for import includes attestation from nCipher that the Key Exchange Key (KEK) that is used to encrypt the key you generate is not exportable and is generated inside a genuine HSM that was manufactured by nCipher.
- The toolset includes attestation from nCipher that the Azure Key Vault security world was also generated on a genuine HSM manufactured by nCipher. This attestation proves to you that Microsoft is also using genuine nCipher hardware.

Check with your security group to determine if the above attestations are required. For detailed steps to create a key on-premises and import it into your key vault, see How to generate and transfer HSM-protected keys for Azure Key Vault. Use the Azure instructions to create a key in each key vault.

Check the recovery level of your keys

Microsoft 365 requires that the Azure Key Vault subscription is set to Do Not Cancel and that the keys used by Customer Key have soft delete enabled. You can confirm you subscriptions settings by looking at the recovery level on your keys.

To check the recovery level of a key, in Azure PowerShell, run the Get-AzKeyVaultKey cmdlet as follows:

```
(Get-AzKeyVaultKey -VaultName <vault name> -Name <key name>).Attributes
```

If the *Recovery Level* property returns anything other than a value of **Recoverable+ProtectedSubscription**, ensure that you have put the subscription on the Do Not Cancel list and that you have soft delete enabled on each of your key vaults.

Back up Azure Key Vault

Immediately following creation or any change to a key, perform a backup and store copies of the backup, both online and offline. Offline copies should not be connected to any network, such as in a physical safe or commercial storage facility. At least one copy of the backup should be stored in a location that will be accessible in the event of a disaster. The backup blobs are the sole means of restoring key material should a Key Vault key be permanently destroyed or otherwise rendered inoperable. Keys that are external to Azure Key Vault and were imported to Azure Key Vault do not qualify as a backup because the metadata necessary for Customer Key to use the key does not exist with the external key. Only a backup taken from Azure Key Vault can be used for restore operations with Customer Key. Therefore, you must create a backup of Azure Key Vault after you upload or create a key.

To create a backup of an Azure Key Vault key, run the Backup-AzKeyVaultKey cmdlet as follows:

```
Backup-AzKeyVaultKey -VaultName <vault name> -Name <key name>
-OutputFile <filename.backup>
```

Ensure that your output file uses the suffix `.backup`.

The output file resulting from this cmdlet is encrypted and cannot be used outside of Azure Key Vault. The backup can be restored only to the Azure subscription from which the backup was taken.

> **TIP**
> For the output file, choose a combination of your vault name and key name. This will make the file name self-describing. It will also ensure that backup file names do not collide.

For example:

```
Backup-AzKeyVaultKey -VaultName Contoso-O365EX-NA-VaultA1 -Name Contoso-O365EX-NA-VaultA1-Key001 -OutputFile
Contoso-O365EX-NA-VaultA1-Key001-Backup-20170802.backup
```

Validate Azure Key Vault configuration settings

Validating before using keys in a DEP is optional, but highly recommended. If you use steps to set up your keys and vaults other than the ones described in this article, validate the health of your Azure Key Vault resources before you configure Customer Key.

To verify that your keys have `get`, `wrapKey`, and `unwrapKey` operations enabled:

Run the Get-AzKeyVault cmdlet as follows:

```
Get-AzKeyVault -VaultName <vault name>
```

In the output, look for the Access Policy and for the Exchange Online identity (GUID) or the SharePoint Online identity (GUID) as appropriate. All three of the above permissions must be shown under Permissions to Keys.

If the access policy configuration is incorrect, run the Set-AzKeyVaultAccessPolicy cmdlet as follows:

```
Set-AzKeyVaultAccessPolicy -VaultName <vault name> -PermissionsToKeys wrapKey,unwrapKey,get -
ServicePrincipalName <Office 365 appID>
```

For example, for Exchange Online and Skype for Business:

```
Set-AzKeyVaultAccessPolicy -VaultName Contoso-O365EX-NA-VaultA1
-PermissionsToKeys wrapKey,unwrapKey,get -ServicePrincipalName 00000002-0000-0ff1-ce00-000000000000
```

For example, for SharePoint Online and OneDrive for Business:

```
Set-AzKeyVaultAccessPolicy -VaultName Contoso-O365SP-NA-VaultA1
-PermissionsToKeys wrapKey,unwrapKey,get -ServicePrincipalName 00000003-0000-0ff1-ce00-000000000000
```

To verify that an expiration date isn't set for your keys, run the Get-AzKeyVaultKey cmdlet as follows:

```
Get-AzKeyVaultKey -VaultName <vault name>
```

Customer Key can't use an expired key. Operations attempted with an expired key will fail, and possibly result in a service outage. We strongly recommend that keys used with Customer Key do not have an expiration date. An expiration date, once set, cannot be removed, but can be changed to a different date. If a key must be used that has an expiration date set, change the expiration value to 12/31/9999. Keys with an expiration date set to a date other than 12/31/9999 will not pass Microsoft 365 validation.

To change an expiration date that has been set to any value other than 12/31/9999, run the Update-AzKeyVaultKey cmdlet as follows:

```
Update-AzKeyVaultKey -VaultName <vault name> -Name <key name> -Expires (Get-Date -Date "12/31/9999")
```

Caution

Don't set expiration dates on encryption keys you use with Customer Key.

Obtain the URI for each Azure Key Vault key

Once you've set up your key vaults and added your keys, run the following command to get the URI for the key in each key vault. You'll need to use these URIs when you create and assign each DEP later, so save this information in a safe place. Run this command once for each key vault.

In Azure PowerShell:

```
(Get-AzKeyVaultKey -VaultName <vault name>).Id
```

Office 365: Setting up Customer Key for Exchange Online and Skype for Business

Before you begin, ensure that you've completed the tasks required to set up Azure Key Vault. See Complete tasks in Azure Key Vault and Microsoft FastTrack for Customer Key for information.

To set up Customer Key for Exchange Online and Skype for Business, you'll complete these steps by remotely connecting to Exchange Online with Windows PowerShell.

Create a data encryption policy (DEP) for use with Exchange Online and Skype for Business

A DEP is associated with a set of keys stored in Azure Key Vault. You assign a DEP to a mailbox in Microsoft 365. Microsoft 365 will then use the keys identified in the policy to encrypt the mailbox. To create the DEP, you need the Key Vault URIs you obtained earlier. See Obtain the URI for each Azure Key Vault key for instructions.

Remember! When you create a DEP, you specify two keys in two different Azure Key Vaults. Create these keys in two separate Azure regions to ensure geo-redundancy.

To create the DEP, follow these steps:

1. On your local computer, using a work or school account that has global administrator permissions in your organization, connect to Exchange Online PowerShell in a Windows PowerShell window.
2. To create a DEP, use the New-DataEncryptionPolicy cmdlet by typing the following command.

   ```
   New-DataEncryptionPolicy -Name <PolicyName> -Description "Policy Description" -AzureKeyIDs
   <KeyVaultURI1>, <KeyVaultURI2>
   ```

   Where:

   - *PolicyName* is the name you want to use for the policy. Names can't contain spaces. For example, USA_mailboxes.
   - *Policy Description* is a user-friendly description of the policy that will help you remember what the policy is for. You can include spaces in the description. For example, "Root key for mailboxes in USA and its territories".
   - *KeyVaultURI1* is the URI for the first key in the policy. For example, https://contoso_EastUSvault01.vault.azure.net/keys/USA_key_01.

- *KeyVaultURI2* is the URI for the second key in the policy. For example, https://contoso_EastUS2vault01.vault.azure.net/keys/USA_Key_02. Separate the two URIs by a comma and a space.

Example:

```
New-DataEncryptionPolicy -Name USA_mailboxes -Description "Root key for mailboxes in USA and its territories" -AzureKeyIDs https://contoso_EastUSvault01.vault.azure.net/keys/USA_key_01, https://contoso_EastUS2vault01.vault.azure.net/keys/USA_Key_02
```

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

Assign a DEP to a mailbox

Assign the DEP to a mailbox by using the Set-Mailbox cmdlet. Once you assign the policy, Microsoft 365 can encrypt the mailbox with the key identified in the DEP.

```
Set-Mailbox -Identity <MailboxIdParameter> -DataEncryptionPolicy <PolicyName>
```

Where *MailboxIdParameter* specifies a user mailbox. For more information about the Set-Mailbox cmdlet, see Set-Mailbox.

In hybrid environments, you can assign a DEP to the on-premises mailbox data that is synchronized into your Exchange Online tenant. To assign a DEP to this synchronized mailbox data, you'll use the Set-MailUser cmdlet. For more information about mailbox data in the hybrid environment, see on-premises mailboxes using Outlook for iOS and Android with hybrid Modern Authentication.

```
Set-MailUser -Identity <MailUserIdParameter> -DataEncryptionPolicy <PolicyName>
```

Where *MailUserIdParameter* specifies a mail user (also known as a mail-enabled user). For more information about the Set-MailUser cmdlet, see Set-MailUser.

Validate mailbox encryption

Encrypting a mailbox can take some time. For first-time policy assignment, the mailbox must also completely move from one database to another before the service can encrypt the mailbox. We recommend that you wait 72 hours before you attempt to validate encryption after you change a DEP or the first time you assign a DEP to a mailbox.

Use the Get-MailboxStatistics cmdlet to determine if a mailbox is encrypted.

```
Get-MailboxStatistics -Identity <GeneralMailboxOrMailUserIdParameter> | fl IsEncrypted
```

The IsEncrypted property returns a value of **true** if the mailbox is encrypted and a value of **false** if the mailbox isn't encrypted. The time to complete mailbox moves depends on the number of mailboxes to which you assign a DEP for the first time, and the size of the mailboxes. If the mailboxes haven't been encrypted after a week from the time you assigned the DEP, contact Microsoft.

Office 365: Setting up Customer Key for SharePoint Online, OneDrive for Business, and Teams files

Before you begin, ensure that you've completed the tasks required to set up Azure Key Vault. See Complete tasks in Azure Key Vault and Microsoft FastTrack for Customer Key for information.

To set up Customer Key for SharePoint Online, OneDrive for Business, and Teams files you complete these steps

by remotely connecting to SharePoint Online with Windows PowerShell.

Create a data encryption policy (DEP) for each SharePoint Online and OneDrive for Business geo

You associate a DEP with a set of keys stored in Azure Key Vault. You apply a DEP to all of your data in one geographic location, also called a geo. If you use the multi-geo feature of Office 365, you can create one DEP per geo with the capability to use different keys per geo. If you aren't using multi-geo, you can create one DEP in your organization for use with SharePoint Online, OneDrive for Business, and Teams files. Microsoft 365 uses the keys identified in the DEP to encrypt your data in that geo. To create the DEP, you need the Key Vault URIs you obtained earlier. See Obtain the URI for each Azure Key Vault key for instructions.

Remember! When you create a DEP, you specify two keys in two different Azure Key Vaults. Create these keys in two separate Azure regions to ensure geo-redundancy.

To create a DEP, you need to remotely connect to SharePoint Online by using Windows PowerShell.

1. On your local computer, using a work or school account that has global administrator permissions in your organization, Connect to SharePoint Online PowerShell.
2. In the Microsoft SharePoint Online Management Shell, run the Register-SPODataEncryptionPolicy cmdlet as follows:

   ```
   Register-SPODataEncryptionPolicy -Identity <adminSiteCollectionURL> -PrimaryKeyVaultName
   <PrimaryKeyVaultName> -PrimaryKeyName <PrimaryKeyName> -PrimaryKeyVersion <PrimaryKeyVersion> -
   SecondaryKeyVaultName <SecondaryKeyVaultName> -SecondaryKeyName <SecondaryKeyName> -
   SecondaryKeyVersion <SecondaryKeyVersion>
   ```

   Example:

   ```
   Register-SPODataEncryptionPolicy -Identity https://contoso.sharepoint.com -PrimaryKeyVaultName
   'stageRG3vault' -PrimaryKeyName 'SPKey3' -PrimaryKeyVersion 'f635a23bd4a44b9996ff6aadd88d42ba' -
   SecondaryKeyVaultName 'stageRG5vault' -SecondaryKeyName 'SPKey5' -SecondaryKeyVersion
   '2b3e8f1d754f438dacdec1f0945f251a’
   ```

   When you register the DEP, encryption begins on the data in the geo. Encryption can take some time. For more information on using this parameter, see Register-SPODataEncryptionPolicy.

Validate file encryption

To validate encryption of SharePoint Online, OneDrive for Business, and Teams files, connect to SharePoint Online PowerShell, and then use the Get-SPODataEncryptionPolicy cmdlet to check the status of your tenant. The *State* property returns a value of **registered** if Customer Key encryption is enabled and all files in all sites have been encrypted. If encryption is still in progress, this cmdlet returns a value of **registering**.

Related articles

- Service encryption with Customer Key
- Manage Customer Key
- Roll or rotate a Customer Key or an availability key
- Learn about the availability key
- Service Encryption

Manage Customer Key

11/2/2020 • 9 minutes to read • Edit Online

After you've set up Customer Key for Office 365, you can manage your keys as described in this article. Learn more about Customer Key in the related topics.

Restore Azure Key Vault keys

Before performing a restore, use the recovery capabilities provided by soft delete. All keys that are used with Customer Key are required to have soft delete enabled. Soft delete acts like a recycle bin and allows recovery for up to 90 days without the need to restore. Restore should only be required in extreme or unusual circumstances, for example if the key or key vault is lost. If you must restore a key for use with Customer Key, in Azure PowerShell, run the Restore-AzureKeyVaultKey cmdlet as follows:

```
Restore-AzKeyVaultKey -VaultName <vault name> -InputFile <filename>
```

For example:

```
Restore-AzKeyVaultKey -VaultName Contoso-O365EX-NA-VaultA1 -InputFile Contoso-O365EX-NA-VaultA1-Key001-
Backup-20170802.backup
```

If the key vault already contains a key with the same name, the restore operation fails. Restore-AzKeyVaultKey restores all key versions and all metadata for the key including the key name.

Manage key vault permissions

Several cmdlets are available that enable you to view and, if necessary, remove key vault permissions. You might need to remove permissions, for example, when an employee leaves the team. For each of these tasks, you will use Azure PowerShell. For information about Azure Powershell, see Overview of Azure PowerShell.

To view key vault permissions, run the Get-AzKeyVault cmdlet.

```
Get-AzKeyVault -VaultName <vault name>
```

For example:

```
Get-AzKeyVault -VaultName Contoso-O365EX-NA-VaultA1
```

To remove an administrator's permissions, run the Remove-AzKeyVaultAccessPolicy cmdlet:

```
Remove-AzKeyVaultAccessPolicy -VaultName <vault name> -UserPrincipalName <UPN of user>
```

For example:

```
Remove-AzKeyVaultAccessPolicy -VaultName Contoso-O365EX-NA-VaultA1 -UserPrincipalName alice@contoso.com
```

Manage data encryption policies (DEPs) with Customer Key

Customer Key handles DEPs differently between the different services. For example, you can create a different number of DEPs for the different services.

Exchange Online and Skype for Business: You can create up to 50 DEPs. For instructions, see Create a data encryption policy (DEP) for use with Exchange Online and Skype for Business.

SharePoint Online, OneDrive for Business, and Teams files: A DEP applies to data in one geographic location, also called a *geo*. If you use the multi-geo feature of Office 365, you can create one DEP per geo. If you are not using multi-geo, you can create one DEP. Normally, you create the DEP when you set up Customer Key. For instructions, see Create a data encryption policy (DEP) for each SharePoint Online and OneDrive for Business geo.

View the DEPs you've created for Exchange Online and Skype for Business

To view a list of all the DEPs you've created for Exchange Online and Skype for Business using the Get-DataEncryptionPolicy PowerShell cmdlet, complete these steps.

1. Using a work or school account that has global administrator permissions in your organization, connect to Exchange Online PowerShell.
2. To return all DEPs in your organization, run the Get-DataEncryptionPolicy cmdlet without any parameters.

   ```
   Get-DataEncryptionPolicy
   ```

   For more information about the Get-DataEncryptionPolicy cmdlet, see Get-DataEncryptionPolicy.

Assign a DEP before you migrate a mailbox to the cloud

When you assign the DEP, Microsoft 365 encrypts the contents of the mailbox using the assigned DEP during the migration. This process is more efficient than migrating the mailbox, assigning the DEP, and then waiting for encryption to take place, which can take hours or possibly days.

To assign a DEP to a mailbox before you migrate it to Office 365, run the Set-MailUser cmdlet in Exchange Online PowerShell:

1. Using a work or school account that has global administrator permissions in your organization, connect to Exchange Online PowerShell.
2. Run the Set-MailUser cmdlet.

   ```
   Set-MailUser -Identity <GeneralMailboxOrMailUserIdParameter> -DataEncryptionPolicy
   <DataEncryptionPolicyIdParameter>
   ```

   Where *GeneralMailboxOrMailUserIdParameter* specifies a mailbox, and *DataEncryptionPolicyIdParameter* is the ID of the DEP. For more information about the Set-MailUser cmdlet, see Set-MailUser.

Determine the DEP assigned to a mailbox

To determine the DEP assigned to a mailbox, use the Get-MailboxStatistics cmdlet. The cmdlet returns a unique identifier (GUID).

1. Using a work or school account that has global administrator permissions in your organization, connect to Exchange Online PowerShell.

   ```
   Get-MailboxStatistics -Identity <GeneralMailboxOrMailUserIdParameter> | fl DataEncryptionPolicyID
   ```

Where *GeneralMailboxOrMailUserIdParameter* specifies a mailbox and DataEncryptionPolicyID returns the GUID of the DEP. For more information about the Get-MailboxStatistics cmdlet, see Get-MailboxStatistics.

2. Run the Get-DataEncryptionPolicy cmdlet to find out the friendly name of the DEP to which the mailbox is assigned.

```
Get-DataEncryptionPolicy <GUID>
```

Where *GUID* is the GUID returned by the Get-MailboxStatistics cmdlet in the previous step.

Verify that Customer Key has finished encryption

Whether you've just rolled a Customer Key, assigned a new DEP, or migrated a mailbox, use the steps in this section to ensure that encryption completes.

Verify encryption completes for Exchange Online and Skype for Business

Encrypting a mailbox can take some time. We recommend that you wait 72 hours before you attempt to validate encryption after you change a DEP or the first time you assign a DEP to a mailbox.

Use the Get-MailboxStatistics cmdlet to determine if a mailbox is encrypted.

```
Get-MailboxStatistics -Identity <GeneralMailboxOrMailUserIdParameter> | fl IsEncrypted
```

The IsEncrypted property returns a value of **true** if the mailbox is encrypted and a value of **false** if the mailbox is not encrypted.

The time to complete mailbox moves depends on the size of the mailbox. If Customer Key hasn't completely encrypted the mailbox after 72 hours from the time you assign a new DEP, contact Microsoft support for help. The New-MoveRequest cmdlet is no longer available for local mailbox moves. Refer to this announcement for additional information.

Verify encryption completes for SharePoint Online, OneDrive for Business, and Teams files

Check on the status of encryption by running the Get-SPODataEncryptionPolicy cmdlet as follows:

```
Get-SPODataEncryptionPolicy -Identity <SPOAdminSiteUrl>
```

The output from this cmdlet includes:

- The URI of the primary key.
- The URI of the secondary key.
- The encryption status for the geo. Possible states include:
  - **Unregistered:** Customer Key encryption has not yet been applied.
  - **Registering:** Customer Key encryption has been applied and your files are in the process of being encrypted. If the key for the geo is registering, you'll also be shown information on what percentage of sites in the geo are complete so that you can monitor encryption progress.
  - **Registered:** Customer Key encryption has been applied, and all files in all sites have been encrypted.
  - **Rolling:** A key roll is in progress. If the key for the geo is rolling, you'll also be shown information on what percentage of sites have completed the key roll operation so that you can monitor

progress.

Unassign a DEP from a mailbox

You unassign a DEP from a mailbox using the Set-mailbox PowerShell cmdlet and setting the `DataEncryptionPolicy` to `$NULL`. Running this cmdlet unassigns the currently assigned DEP and reencrypts the mailbox using the DEP associated with default Microsoft managed keys. You can't unassign the DEP used by Microsoft managed keys. If you don't want to use Microsoft managed keys, you can assign another DEP to the mailbox.

To unassign the DEP from a mailbox using the Set-Mailbox PowerShell cmdlet, complete these steps.

1. Using a work or school account that has global administrator permissions in your organization, connect to Exchange Online PowerShell.
2. Run the Set-Mailbox cmdlet.

```
Set-Mailbox -Identity <mailbox> -DataEncryptionPolicy $NULL
```

Revoke your keys and start the data purge path process

You control the revocation of all root keys including the availability key. Customer Key provides control of the exit planning aspect of the regulatory requirements for you. If you decide to revoke your keys to purge your data and exit the service, the service deletes the availability key once the data purge process completes.

Microsoft 365 audits and validates the data purge path. For more information, see the SSAE 18 SOC 2 Report available on the Service Trust Portal. In addition, Microsoft recommends the following documents:

- Risk Assessment and Compliance Guide for Financial Institutions in the Microsoft Cloud
- O365 Exit Planning Considerations

The data purge path differs slightly between the different services.

Revoke your Customer Keys and the availability key for Exchange Online and Skype for Business

When you initiate the data purge path for Exchange Online and Skype for Business, you set a permanent data purge request on a DEP. Doing so permanently deletes encrypted data within the mailboxes to which that DEP is assigned.

Since you can only run the PowerShell cmdlet against one DEP at a time, consider reassigning a single DEP to all of your mailboxes before you initiate the data purge path.

> **WARNING**
> Do not use the data purge path to delete a subset of your mailboxes. This process is only intended for customers who are exiting the service.

To initiate the data purge path, complete these steps:

1. Remove wrap and unwrap permissions for "O365 Exchange Online" from Azure Key Vaults.
2. Using a work or school account that has global administrator privileges in your organization, connect to Exchange Online PowerShell.
3. For each DEP that contains mailboxes that you want to delete, run the Set-DataEncryptionPolicy cmdlet as follows.

```
Set-DataEncryptionPolicy <Policy ID> -PermanentDataPurgeRequested -PermanentDataPurgeReason <Reason>
-PermanentDataPurgeContact <ContactName>
```

If the command fails, ensure that you've removed the Exchange Online permissions from both keys in Azure Key Vault as specified earlier in this task.Once you've set the PermanentDataPurgeRequested switch using the Set-DataEncryptionPolicy cmdlet, you'll no longer be able to assign this DEP to mailboxes.

4. Contact Microsoft support and request the Data Purge eDocument.

   At your request, Microsoft sends you a legal document to acknowledge and authorize data deletion. The person in your organization who signed up as an approver in the FastTrack offer during onboarding needs to sign this document. Normally, this is an executive or other designated person in your company who is legally authorized to sign the paperwork on behalf of your organization.

5. Once your representative has signed the legal document, return it to Microsoft (usually through an eDoc signature).

   Once Microsoft receives the legal document, Microsoft runs cmdlets to trigger the data purge which first deletes the policy, marks the mailboxes for permanent deletion, then deletes the availability key. Once the data purge process completes, the data has been purged, is inaccessible to Exchange Online, and is not recoverable.

Revoke your Customer Keys and the availability key for SharePoint Online, OneDrive for Business, and Teams files

To initiate the data purge path for SharePoint Online, OneDrive for Business, and Teams files, complete these steps:

1. Revoke Azure Key Vault access. All key vault admins must agree to revoke access.

   You do not delete the Azure Key Vault for SharePoint Online. Key vaults may be shared among several SharePoint Online tenants and DEPs.

2. Contact Microsoft to delete the availability key.

   When you contact Microsoft to delete the availability key, we'll send you a legal document. The person in your organization who signed up as an approver in the FastTrack offer during onboarding needs to sign this document. Normally, this is an executive or other designated person in your company who's legally authorized to sign the paperwork on behalf of your organization.

3. Once your representative signs the legal document, return it to Microsoft (usually through an eDoc signature).

   Once Microsoft receives the legal document, we run cmdlets to trigger the data purge which performs crypto deletion of the tenant key, site key, and all individual per-document keys, irrevocably breaking the key hierarchy. Once the data purge cmdlets complete, your data has been purged.

Related articles

- Service encryption with Customer Key
- Learn about the availability key
- Set up Customer Key
- Roll or rotate a Customer Key or an availability key
- Customer Lockbox

- Service Encryption

Roll or rotate a Customer Key or an availability key

4/21/2020 • 4 minutes to read • Edit Online

Caution

Only roll an encryption key that you use with Customer Key when your security or compliance requirements dictate that you must roll the key. In addition, do not delete any keys that are or were associated with policies. When you roll your keys, there will be content encrypted with the previous keys. For example, while active mailboxes will be re-encrypted frequently, inactive, disconnected, and disabled mailboxes may still be encrypted with the previous keys. SharePoint Online performs backup of content for restore and recovery purposes, so there may still be archived content using older keys.

About rolling the availability key

Microsoft does not expose direct control of the availability key to customers. For example, you can only roll (rotate) the keys that you own in Azure Key Vault. Microsoft 365 rolls the availability keys on an internally-defined schedule. There is no customer-facing, service-level agreement (SLA) for these key rolls. Microsoft 365 rotates the availability key using Microsoft 365 service code in an automated, non-manual process. Microsoft administrators may initiate the roll process. The key is rolled using automated mechanisms without direct access to the key store. Access to the availability key secret store is not provisioned to Microsoft administrators. Availability key rolling leverages the same mechanism used to initially generate the key. For more information about the availability key, see Understand the availability key.

IMPORTANT

Exchange Online and Skype for Business availability keys can be effectively rolled by customers creating a new DEP, since a unique availability key is generated for each DEP you create. Availability keys for SharePoint Online, OneDrive for Business, and Teams files exist at the forest level and are shared across DEPs and customers, which means rolling only occurs at a Microsoft internally defined schedule. To mitigate the risk of not rolling the availability key each time a new DEP is created, SharePoint, OneDrive, and Teams roll the tenant intermediate key (TIK), the key wrapped by the customer root keys and availability key, each time a new DEP is created.

Request a new version of each existing root key you want to roll

When you roll a key, you request a new version of an existing key. To request a new version of an existing key, you use the same cmdlet, Add-AzKeyVaultKey, with the same syntax that you used to create the key in the first place. After you've finished rolling any key associated with a Data Encryption Policy (DEP), you run another cmdlet to ensure that Customer Key begins using the new key. Do this step in each Azure Key Vault (AKV).

For example:

1. Sign in to your Azure subscription with Azure PowerShell. For instructions, see Sign in with Azure PowerShell.
2. Run the Add-AzKeyVaultKey cmdlet as shown in the following example:

```
Add-AzKeyVaultKey -VaultName Contoso-O365EX-NA-VaultA1 -Name Contoso-O365EX-NA-VaultA1-Key001 -
Destination HSM -KeyOps @('wrapKey','unwrapKey') -NotBefore (Get-Date -Date "12/27/2016 12:01 AM")
```

In this example, since a key named **Contoso-O365EX-NA-VaultA1-Key001** exists in the **Contoso-O365EX-NA-VaultA1** vault, the cmdlet creates a new version of the key. This operation preserves the

previous key versions in the version history for the key. You need the previous key version to decrypt the data that it still encrypts. Once you complete rolling any key associated with a DEP, run an extra cmdlet to ensure that Customer Key begins using the new key. The following sections describe the cmdlets in more detail.

Update the Customer Key for Exchange Online and Skype for Business

When you roll either of the Azure Key Vault keys associated with a DEP used with Exchange Online and Skype for Business, you must update the DEP to point to the new key. This does not rotate the availability key.

To instruct Customer Key to use the new key to encrypt mailboxes, run the Set-DataEncryptionPolicy cmdlet as follows:

1. Run the Set-DataEncryptionPolicy cmdlet in Azure PowerShell:

   ```
   Set-DataEncryptionPolicy -Identity <DataEncryptionPolicyID> -Refresh
   ```

   Within 72 hours, the active mailboxes associated with this DEP become encrypted with the new key.

2. To check the value for the DataEncryptionPolicyID property for the mailbox, use the steps in Determine the DEP assigned to a mailbox. The value for this property changes once the service applies the updated key.

Update the Customer Key for SharePoint Online, OneDrive for Business, and Teams files

SharePoint Online only allows you to roll one key at a time. If you want to roll both keys in a key vault, wait for the first operation to complete. Microsoft recommends that you stagger your operations to avoid this issue. When you roll either of the Azure Key Vault keys associated with a DEP used with SharePoint Online and OneDrive for Business, you must update the DEP to point to the new key. This does not rotate the availability key.

1. Run the Update-SPODataEncryptionPolicy cmdlet as follows:

   ```
   Update-SPODataEncryptionPolicy -Identity <SPOAdminSiteUrl> -KeyVaultName <ReplacementKeyVaultName> -
   KeyName <ReplacementKeyName> -KeyVersion <ReplacementKeyVersion> -KeyType <Primary | Secondary>
   ```

   While this cmdlet starts the key roll operation for SharePoint Online and OneDrive for Business, the action doesn't complete immediately.

2. To see the progress of the key roll operation, run the Get-SPODataEncryptionPolicy cmdlet as follows:

   ```
   Get-SPODataEncryptionPolicy -Identity <SPOAdminSiteUrl>
   ```

Related articles

- Service encryption with Customer Key for Office 365
- Set up Customer Key for Office 365
- Manage Customer Key for Office 365
- Learn about the availability key

Learn about the availability key for Customer Key

4/21/2020 • 15 minutes to read • Edit Online

The availability key is a root key automatically generated and provisioned when you create a data encryption policy. Microsoft 365 stores and protects the availability key. The availability key is functionally like the two root keys that you supply for service encryption with Customer Key. The availability key wraps the keys one tier lower in the key hierarchy. Unlike the keys that you provide and manage in Azure Key Vault, you can't directly access the availability key. Microsoft 365 automated services manage the availability key programatically. These services initiate automated operations that never involve direct access to the availability key.

The primary purpose of the availability key is to provide recovery capability from the unanticipated loss of root keys that you manage. Loss could be a result of mismanagement or malicious action. If you lose control of your root keys, contact Microsoft Support and Microsoft will assist you through the process of recovery using the availability key. You'll use the availability key to migrate to a new Data Encryption Policy with new root keys you provision.

Storage and control of the availability key are deliberately different from Azure Key Vault keys for three reasons:

- The availability key provides a recovery, "break-glass" capability if control over both Azure Key Vault keys is lost.
- The separation of logical controls and secure storage locations provides defense-in-depth and protects against the loss of all keys, and your data, from a single attack or point of failure.
- The availability key provides a high-availability capability if Microsoft 365 services are unable to reach keys hosted in Azure Key Vault due to transient errors. This rule only applies to Exchange Online and Skype for Business service encryption. SharePoint Online, OneDrive for Business, and Teams files never use the availability key unless you explicitly instruct Microsoft to initiate the recovery process.

Sharing the responsibility to protect your data, using a variety of protections and processes for key management, ultimately reduces the risk that all keys (and therefore your data) will be permanently lost or destroyed. Microsoft provides you with sole authority over the disablement or destruction of the availability key when you leave the service. By design, no one at Microsoft has access to the availability key: it is only accessible by Microsoft 365 service code.

See the Microsoft Trust Center for more information about how we secure keys.

Availability key uses

The availability key provides recovery capability for scenarios in which an external malefactor or malicious insider steals control of your key vault, or when inadvertent mismanagement results in loss of root keys. This recovery capability applies to all Microsoft 365 services compatible with Customer Key. Individual services use the availability key differently. Microsoft 365 only uses the availability key in the ways described below.

Exchange Online and Skype for Business uses

In addition to the recovery capability, Exchange Online and Skype for Business use the availability key to ensure data availability during transient, or intermittent operational issues, related to the service accessing root keys. When the service cannot reach either of your Customer Keys in Azure Key Vault due to transient errors, the service automatically uses the availability key. The service NEVER goes directly to the availability key.

Automated systems in Exchange Online and Skype for Business may use the availability key during transient errors to support automated back-end services such as anti-virus, e-discovery, data loss prevention, mailbox moves, and data indexing.

SharePoint Online, OneDrive for Business, and Teams files uses

For SharePoint Online, OneDrive for Business, and Teams files, the availability key is NEVER used outside of the recovery capability and customers must explicitly instruct Microsoft to initiate use of the availability key during a recovery scenario. Automated service operations solely rely on your Customer Keys in Azure Key vault. For in-depth information about how the key hierarchy works for these services, see How SharePoint Online, OneDrive for Business, and Teams files use the availability key.

Availability key security

Microsoft shares the responsibility of data protection with you by instantiating the availability key and taking extensive measures to protect it. Microsoft does not expose direct control of the availability key to customers. For example, you can only roll (rotate) the keys that you own in Azure Key Vault. For more information, see Roll or rotate a customer key or an availability key.

Availability key secret stores

Microsoft protects availability keys in access-controlled, internal secret stores like the customer-facing Azure Key Vault. We implement access controls to prevent Microsoft administrators from directly accessing the secrets contained within. Secret Store operations, including key rotation and deletion, occur through automated commands that never involve direct access to the availability key. Secret store management operations are limited to specific engineers and require privilege escalation through an internal tool, Lockbox. Privilege escalation requires manager approval and justification prior to being granted. Lockbox ensures access is time bound with automatic access revocation upon time expiration or engineer log out.

Exchange Online and Skype for Business availability keys are stored in an Exchange Online Active Directory secret store. Availability keys are securely stored inside tenant specific containers within the Active Directory Domain Controller. This secure storage location is separate and isolated from the SharePoint Online, OneDrive for Business, and Teams files secret store.

SharePoint Online, OneDrive for Business, and Teams files availability keys are stored in an internal secret store managed by the service team. This secured, secrets storage service has front-end servers with application endpoints and a SQL Database as the back end. Availability keys are stored in the SQL Database and are wrapped (encrypted) by secret store encryption keys that use a combination of AES-256 and HMAC to encrypt the availability key at rest. The secret store encryption keys are stored in a logically isolated component of the same SQL Database and are further encrypted with RSA-2048 keys contained in certificates managed by the Microsoft certificate authority (CA). These certificates are stored in the secret store front-end servers that perform operations against the database.

Defense-in-depth

Microsoft employs a defense-in-depth strategy to prevent malicious actors from impacting the confidentiality, integrity, or availability of customer data stored in the Microsoft Cloud. Specific preventive and detective controls are implemented to protect the secret store and the availability key as part of the overarching security strategy.

Microsoft 365 is built to prevent misuse of the availability key. The application layer is the only method through which keys, including the availability key, can be used to encrypt and decrypt data. Only Microsoft 365 service code has the ability to interpret and traverse the key hierarchy for encryption and decryption activities. Logical isolation exists between the storage locations of Customer Keys, availability keys, other hierarchical keys, and customer data. This isolation mitigates the risk of data exposure in the event one or more locations are compromised. Each layer in the hierarchy has built in 24x7 intrusion detection capabilities to protect data and secrets stored.

Access controls are implemented to prevent unauthorized access to internal systems, including availability key secret stores. Microsoft engineers don't have direct access to the availability key secret stores. For additional detail on access controls, review Administrative Access Controls in Microsoft 365.

Technical controls prevent Microsoft personnel from logging into highly-privileged service accounts, which might otherwise be used by attackers to impersonate Microsoft services. For example, these controls prevent interactive logon.

Security logging and monitoring controls are another defense-in-depth safeguard implemented that mitigate risk to Microsoft services and your data. Microsoft service teams have deployed active monitoring solutions that generate alerts and audit logs. All service teams upload their logs to a central repository where the logs are aggregated and processed. Internal tools automatically examine records to confirm that services are functioning in an optimal, resilient, and secure state. Unusual activity is flagged for further review.

Any log event that indicates a potential violation of the Microsoft Security Policy is immediately brought to the attention of Microsoft security teams. Microsoft 365 security has configured alerts to detect attempted access to availability key secret stores. Alerts are also generated if Microsoft personnel attempt interactive logon to service accounts, which is prohibited and protected by access controls. Microsoft 365 security also detects and alerts upon deviations of the Microsoft 365 service from normal baseline operations. Malefactors attempting to misuse Microsoft 365 services would trigger alerts resulting in the offender's eviction from the Microsoft cloud environment.

Use the availability key to recover from key loss

If you lose control of your Customer Keys, the availability key provides you the ability to recover and re-encrypt your data.

Recovery procedure for Exchange Online and Skype for Business

If you lose control of your Customer Keys, the availability key gives you the capability to recover your data and bring your impacted Microsoft 365 resources back online. The availability key continues to protect your data while you recover.At a high level, to fully recover from key loss, you'll need to create a new DEP and move impacted resources to the new policy.

To encrypt your data with new Customer Keys, create new keys in Azure Key Vault, create a new DEP using the new Customer Keys, then assign the new DEP to the mailboxes currently encrypted with the previous DEP for which the keys were lost or compromised.

This re-encryption process can take up to 72 hours. This is the standard duration when you change a DEP.

Recovery procedure for SharePoint Online, OneDrive for Business, and Teams files

For SharePoint Online, OneDrive for Business, and Teams files, the availability key is NEVER used outside of the recovery capability. You must explicitly instruct Microsoft to initiate use of the availability key during a recovery scenario. To initiate the recovery process, contact Microsoft to activate the availability key. Once activated, the availability key is automatically used to decrypt your data allowing you to encrypt the data with a newly-created DEP associated to new Customer Keys.

This operation is proportional to the number of sites in your organization. Once you call Microsoft to use the availability key, you should be fully online within about four hours.

How Exchange Online and Skype for Business use the availability key

When you create a DEP with Customer Key, Microsoft 365 generates a Data Encryption Policy Key (DEP Key) associated with that DEP. The service encrypts the DEP Key three times: once with each of the customer keys and once with the availability key. Only the encrypted versions of the DEP Key are stored, and a DEP Key can only be decrypted with the customer keys or the availability key. The DEP Key is then used to encrypt Mailbox Keys, which encrypt individual mailboxes.

Microsoft 365 follows this process to decrypt and provide data when customers are using the service:

1. Decrypt the DEP Key using the Customer Key.

2. Use the decrypted DEP Key to decrypt a Mailbox Key.
3. Use the decrypted Mailbox Key to decrypt the mailbox itself, allowing you to access the data within the mailbox.

How SharePoint Online, OneDrive for Business, and Teams files use the availability key

The SharePoint Online and OneDrive for Business architecture and implementation for Customer Key and availability key are different from Exchange Online and Skype for Business.

When an organization moves to customer-managed keys, Microsoft 365 creates an organization-specific intermediate key (TIK). Microsoft 365 encrypts the TIK twice, once with each of the customer keys, and stores the two encrypted versions of the TIK. Only the encrypted versions of the TIK are stored, and a TIK can only be decrypted with the customer keys. The TIK is then used to encrypt site keys, which are then used to encrypt blob keys (also called file chunk keys). Depending on file size, the service may split a file into multiple file chunks each with a unique key. The blobs (file chunks) themselves are encrypted with the blob keys and stored in the Microsoft Azure Blob storage service.

Microsoft 365 follows this process to decrypt and provide customer files when customers are using the service:

1. Decrypt the TIK using the Customer Key.
2. Use the decrypted TIK to decrypt a site key.
3. Use the decrypted site key to decrypt a blob key.
4. Use the decrypted blob key to decrypt the blob.

Microsoft 365 decrypts a TIK by issuing two decryption requests to Azure Key Vault with a slight offset. The first one to finish furnishes the result, canceling the other request.

In case you lose access to your customer keys, Microsoft 365 also encrypts the TIK with an availability key and stores this along with the TIKs encrypted with each customer key. The TIK encrypted with the availability key is used only when the customer calls Microsoft to enlist the recovery path when they have lost access to their keys, maliciously or accidentally.

For availability and scale reasons, decrypted TIKs are cached in a time-limited memory cache. Two hours before a TIK cache is set to expire, Microsoft 365 attempts to decrypt each TIK. Decrypting the TIKs extends the lifetime of the cache. If TIK decryption fails for a significant amount of time, Microsoft 365 generates an alert to notify engineering prior to the cache expiration. Only if the customer calls Microsoft will Microsoft 365 initiate the recovery operation, which involves decrypting the TIK with the availability key stored in Microsoft's secret store and onboarding the tenant again using the decrypted TIK and a new set of customer-supplied Azure Key Vault keys.

As of today, Customer Key is involved in the encryption and decryption chain of SharePoint Online file data stored in the Azure blob store, but not SharePoint Online list items or metadata stored in the SQL Database. Microsoft 365 does not use the availability key for Exchange Online, Skype for Business, SharePoint Online, OneDrive for Business, and Teams files other than the case described above, which is customer-initiated. Human access to customer data is protected by Customer Lockbox.

Availability key triggers

Microsoft 365 triggers the availability key only in specific circumstances. These circumstances differ by service.

Triggers for Exchange Online and Skype for Business

1. Microsoft 365 reads the DEP to which the mailbox is assigned in order to determine the location of the

two Customer Keys in Azure Key Vault.

2. Microsoft 365 randomly chooses one of the two Customer Keys from the DEP and sends a request to Azure Key Vault to unwrap the DEP key using the Customer Key.
3. If the request to unwrap the DEP key using the Customer Key fails, Microsoft 365 sends a second request to Azure Key Vault, this time instructing it to use the alternate (second) Customer Key.
4. If the second request to unwrap the DEP key using the Customer Key fails, Microsoft 365 examines the results of both requests.
    - If the examination determines that the requests failed returning a system ERROR:
        - Microsoft 365 triggers the availability key to decrypt the DEP key.
        - Microsoft 365 then uses the DEP key to decrypt the mailbox key and complete the user request.
        - In this case, Azure Key Vault is either unable to respond or unreachable due to a transient ERROR.
    - If the examination determines that the requests failed returning ACCESS DENIED:
        - This means deliberate, inadvertent, or malicious action has been taken to render the customer keys unavailable (for example, during the data purge process as part of leaving the service).
        - In this case, the availability key will be used only for system actions and not for user actions, the user request fails, and the user receives an error message.

> **IMPORTANT**
> Microsoft 365 service code always has a valid login token for reasoning over customer data to provide value-adding cloud services. Therefore, until the availability key has been deleted, it can be used as a fallback for actions initiated by, or internal to, Exchange Online and Skype for Business such as search index creation or moving mailboxes. This applies to both transient ERRORS and ACCESS DENIED requests to Azure Key Vault.

Triggers for SharePoint Online, OneDrive for Business, and Teams files

For SharePoint Online, OneDrive for Business, and Teams files, the availability key is NEVER used outside of the recovery capability and customers must explicitly instruct Microsoft to initiate use of the availability key during a recovery scenario.

Audit logs and the availability key

Automated systems in Microsoft 365 process all data as it flows through the system to provide cloud services, for example, anti-virus, e-discovery, data loss prevention, and data indexing. Microsoft 365 does not generate customer-visible logs for this activity. In addition, Microsoft personnel do not access your data as part of these normal system operations.

Exchange Online and Skype for Business availability key logging

When Exchange Online and Skype for Business accesses availability key to provide service, Microsoft 365 publishes customer-visible logs accessible from the Security and Compliance Center. An audit log record for the availability key operation is generated each time the service uses the availability key. A new record type called "Customer Key Service Encryption" with activity type "Fallback to Availability Key" allows admins to filter Unified Audit Log search results to view availability key records.

Log records include attributes such as date, time, activity, organization ID, and data encryption policy ID. The

record is available as part of Unified Audit Logs and is accessible from the Security & Compliance Center Audit Log Search tab.

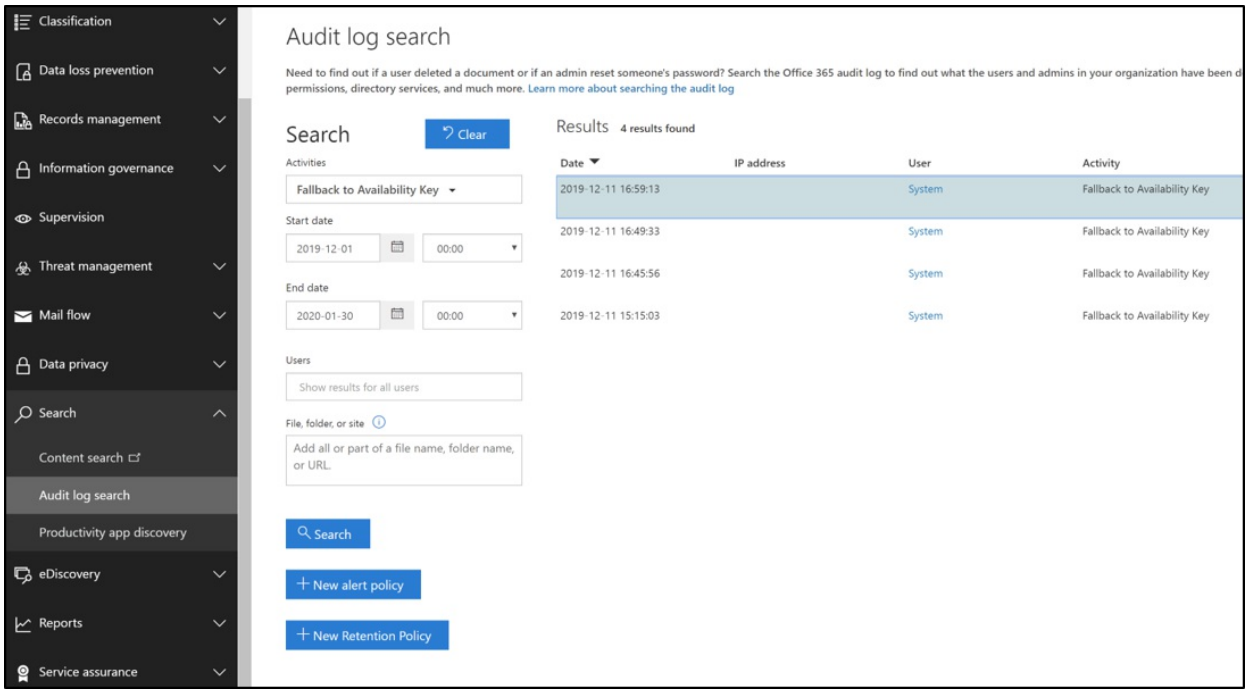

Exchange Online and Skype for Business availability key records use the Office 365 Management Activity common schema with added custom parameters: Policy Id, Scope Key Version Id, and Request Id.

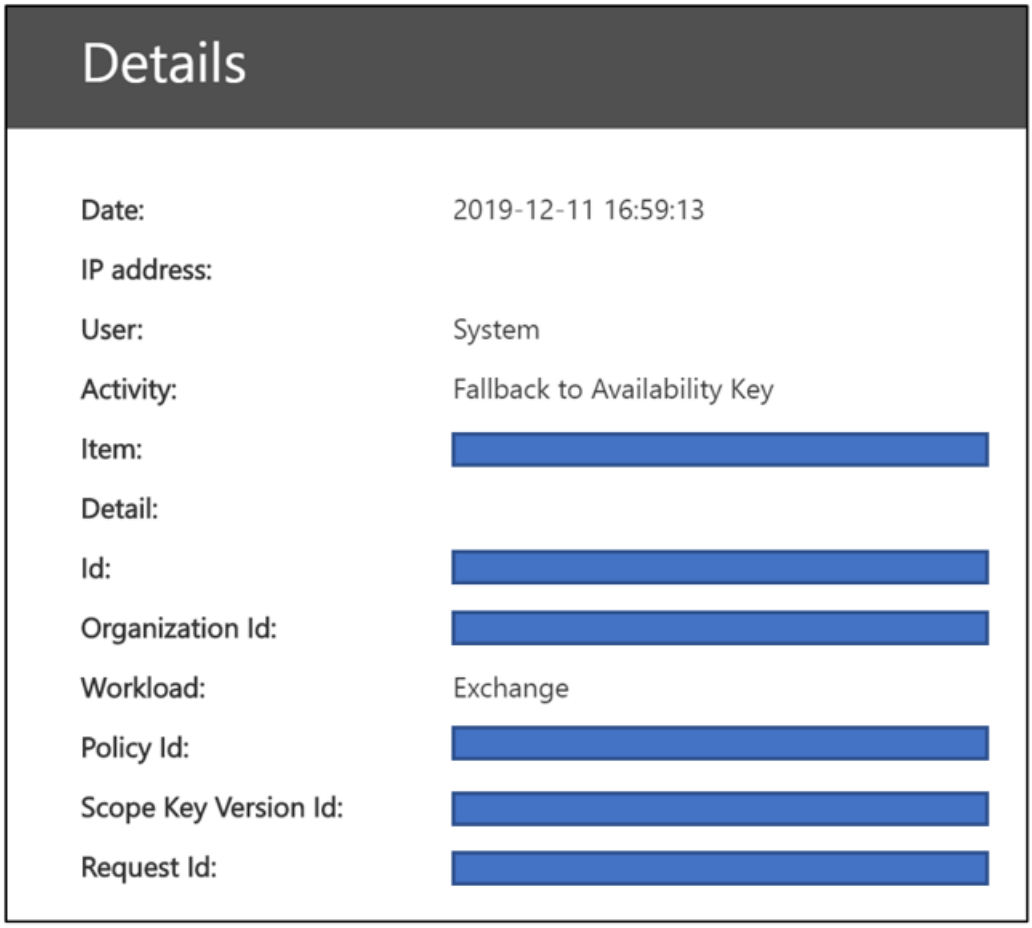

SharePoint Online, OneDrive for Business, and Teams files availability key logging

Availability key logging isn't available yet for these services. For SharePoint Online, OneDrive for Business, and Teams files, the availability key is only activated by Microsoft, when instructed by you, for recovery purposes. As a result, you already know every event in which the availability key is used for these services.

Availability key in the Customer Key hierarchy

Microsoft 365 uses the availability key to wrap the tier of keys lower in the key hierarchy established for Customer Key service encryption. Different key hierarchies exist between services. Key algorithms also differ between availability keys and other keys in the hierarchy of each applicable service. The availability key algorithms used by the different services are as follows:

- The Exchange Online and Skype for Business availability keys use AES-256.
- The SharePoint Online, OneDrive for Business, and Teams files availability keys use RSA-2048.

Encryption ciphers used to encrypt keys for Exchange Online and Skype for Business

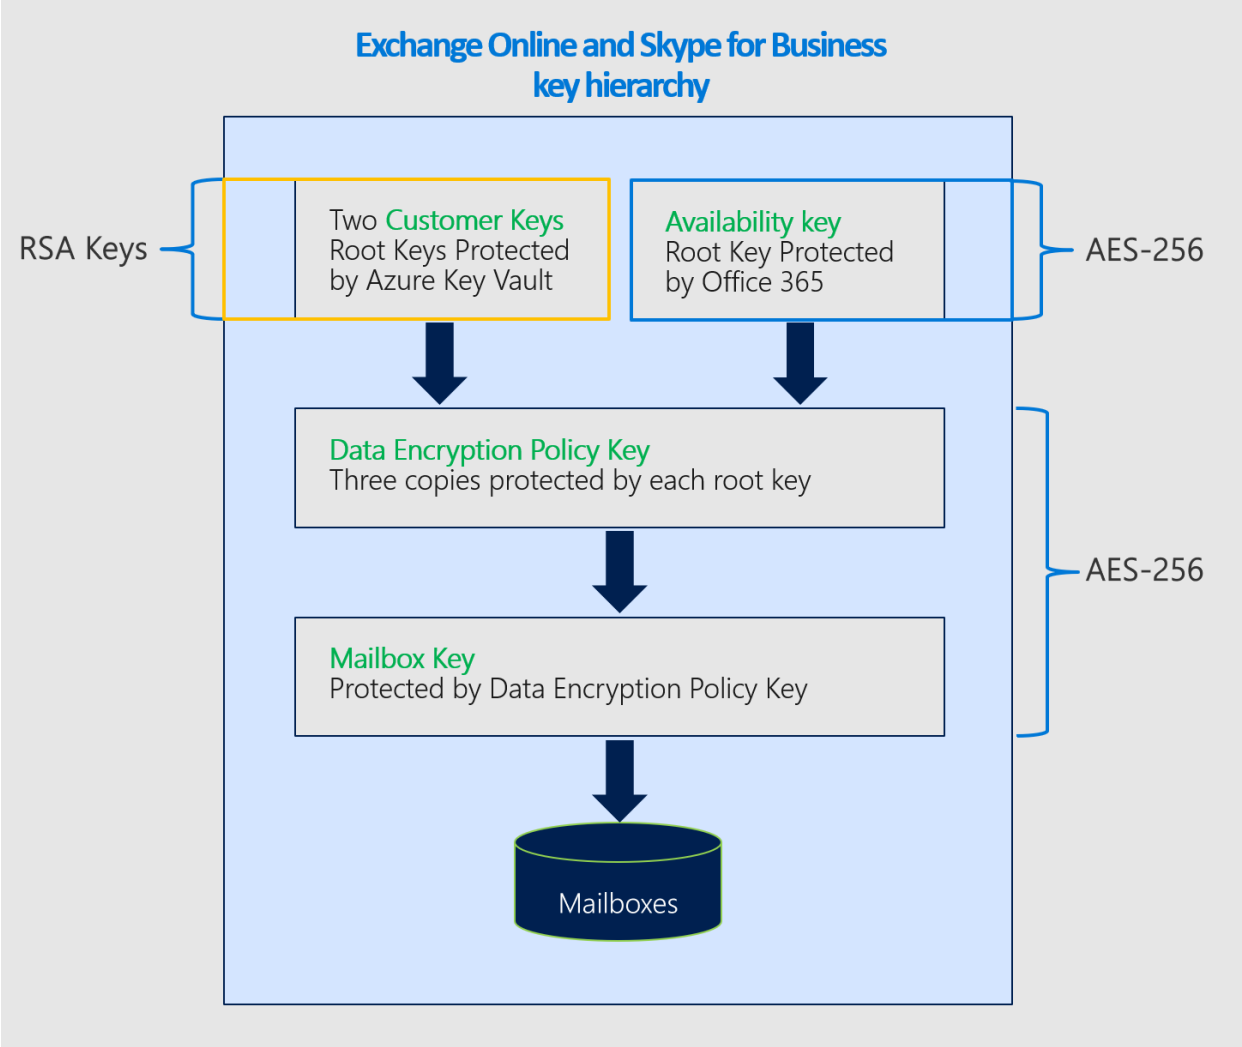

Encryption ciphers used to encrypt keys for SharePoint Online and OneDrive for Business

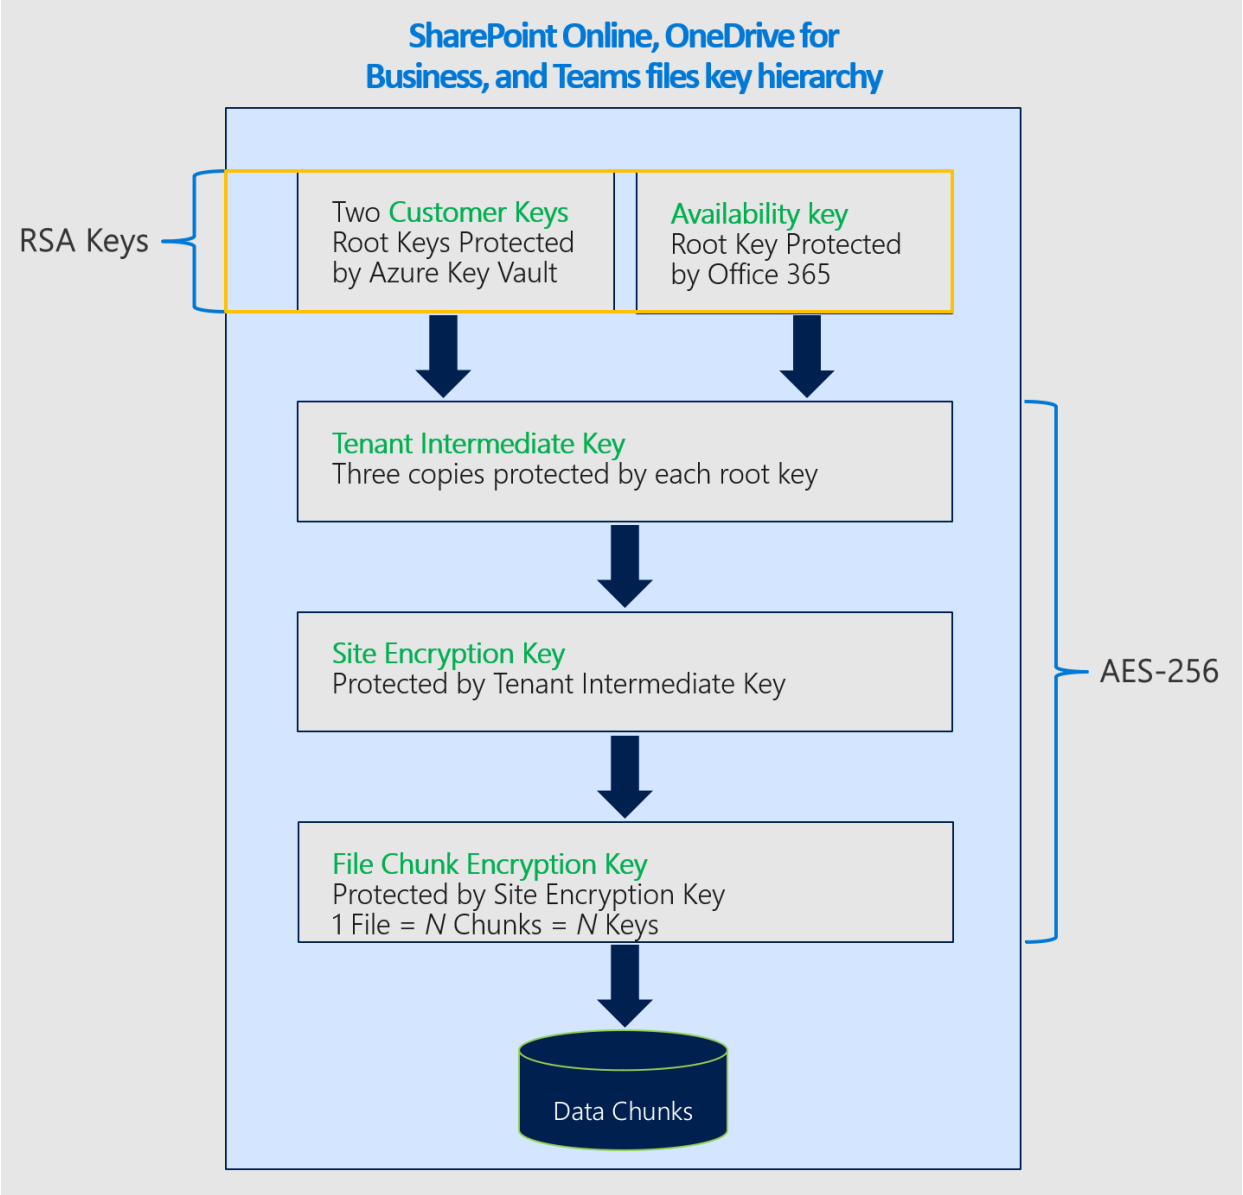

Related articles

- Service encryption with Customer Key
- Set up Customer Key
- Manage Customer Key
- Roll or rotate a Customer Key or an availability key

Email encryption

2/18/2021 • 6 minutes to read • Edit Online

This article compares encryption options in Microsoft 365 including Office 365 Message Encryption (OME), S/MIME, Information Rights Management (IRM), and introduces Transport Layer Security (TLS).

Microsoft 365 delivers multiple encryption options to help you meet your business needs for email security. This article presents three ways to encrypt email in Office 365. If you want to learn more about all security features in Office 365, visit the Office 365 Trust Center. This article introduces the three types of encryption available for Microsoft 365 administrators to help secure email in Office 365:

- Office Message Encryption (OME).
- Secure/Multipurpose Internet Mail Extensions (S/MIME).
- Information Rights Management (IRM).

How Microsoft 365 uses email encryption

Encryption is the process by which information is encoded so that only an authorized recipient can decode and consume the information. Microsoft 365 uses encryption in two ways: in the service, and as a customer control. In the service, encryption is used in Microsoft 365 by default; you don't have to configure anything. For example, Microsoft 365 uses Transport Layer Security (TLS) to encrypt the connection, or session, between two servers.

Here's how email encryption typically works:

- A message is encrypted, or transformed from plain text into unreadable ciphertext, either on the sender's machine, or by a central server while the message is in transit.
- The message remains in ciphertext while it's in transit in order to protect it from being read in case the message is intercepted.
- Once the message is received by the recipient, the message is transformed back into readable plain text in one of two ways:
    - The recipient's machine uses a key to decrypt the message, or
    - A central server decrypts the message on behalf of the recipient, after validating the recipient's identity.

For more information on how Microsoft 365 secures communication between servers, such as between organizations within Microsoft 365 or between Microsoft 365 and a trusted business partner outside of Microsoft 365, see How Exchange Online uses TLS to secure email connections in Office 365.

Watch this video for an introduction to Encryption in Office 365.

Comparing email encryption options available in Office 365

EMAIL ENCRYPTION TECHNOLOGY	OME 	IRM 	S/MIME 
What is it?	Office 365 Message Encryption (OME) is a service built on Azure Rights Management (Azure RMS) that lets you send encrypted email to people inside or outside your organization, regardless of the destination email address (Gmail, Yahoo! Mail, Outlook.com, etc.). As an admin, you can set up transport rules that define the conditions for encryption. When a user sends a message that matches a rule, encryption is applied automatically. To view encrypted messages, recipients can either get a one-time passcode, sign in with a Microsoft account, or sign in with a work or school account associated with Office 365. Recipients can also send encrypted replies. They don't need a Microsoft 365 subscription to view encrypted messages or send encrypted replies.	IRM is an encryption solution that also applies usage restrictions to email messages. It helps prevent sensitive information from being printed, forwarded, or copied by unauthorized people. IRM capabilities in Microsoft 365 use Azure Rights Management (Azure RMS).	S/MIME is a certificate-based encryption solution that allows you to both encrypt and digitally sign a message. The message encryption helps ensure that only the intended recipient can open and read the message. A digital signature helps the recipient validate the identity of the sender. Both digital signatures and message encryption are made possible through the use of unique digital certificates that contain the keys for verifying digital signatures and encrypting or decrypting messages. To use S/MIME, you must have public keys on file for each recipient. Recipients have to maintain their own private keys, which must remain secure. If a recipient's private keys are compromised, the recipient needs to get a new private key and redistribute public keys to all potential senders.
What does it do?	OME: Encrypts messages sent to internal or external recipients. Allows users to send encrypted messages to any email address, including Outlook.com, Yahoo! Mail, and Gmail. Allows you, as an admin, to customize the email viewing portal to reflect your organization's brand. Microsoft securely manages and stores the keys, so you don't have to. No special client side software is needed as long as the encrypted message (sent as an HTML attachment) can be opened in a browser.	IRM: Uses encryption and usage restrictions to provide online and offline protection for email messages and attachments. Gives you, as an admin, the ability to set up transport rules or Outlook protection rules to automatically apply IRM to select messages. Lets users manually apply templates in Outlook or Outlook on the web (formerly known as Outlook Web App).	S/MIME addresses sender authentication with digital signatures, and message confidentiality with encryption.

EMAIL ENCRYPTION TECHNOLOGY	OME 	IRM 	S/MIME 
What does it not do?	OME doesn't let you apply usage restrictions to messages. For example, you can't use it to stop a recipient from forwarding or printing an encrypted message.	Some applications may not support IRM emails on all devices. For more information about these and other products that support IRM email, see Client device capabilities.	S/MIME doesn't allow encrypted messages to be scanned for malware, spam, or policies.
Recommendations and example scenarios	We recommend using OME when you want to send sensitive business information to people outside your organization, whether they're consumers or other businesses. For example: A bank employee sending credit card statements to customers A doctor's office sending medical records to a patient An attorney sending confidential legal information to another attorney	We recommend using IRM when you want to apply usage restrictions as well as encryption. For example: A manager sending confidential details to her team about a new product applies the "Do Not Forward" option. An executive needs to share a bid proposal with another company, which includes an attachment from a partner who is using Office 365, and require both the email and the attachment to be protected.	We recommend using S/MIME when either your organization or the recipient's organization requires true peer-to-peer encryption. S/MIME is most commonly used in the following scenarios: Government agencies communicating with other government agencies A business communicating with a government agency

What encryption options are available for my Microsoft 365 subscription?

For information about email encryption options for your Microsoft 365 subscription see the Exchange Online service description. Here, you can find information about the following encryption features:

- Azure RMS, including both IRM capabilities and OME
- S/MIME
- TLS
- Encryption of data at rest (through BitLocker)

You can also use third-party encryption tools with Microsoft 365, for example, PGP (Pretty Good Privacy). Microsoft 365 does not support PGP/MIME and you can only use PGP/Inline to send and receive PGP-encrypted emails.

What about encryption for data at rest?

"Data at rest" refers to data that isn't actively in transit. In Microsoft 365, email data at rest is encrypted using BitLocker Drive Encryption. BitLocker encrypts the hard drives in Microsoft datacenters to provide enhanced protection against unauthorized access. To learn more, see BitLocker Overview.

More information about email encryption options

For more information about the email encryption options in this article as well as TLS, see these articles:

OME

Office 365 Message Encryption (OME)

IRM

Information Rights Management in Exchange Online

What is Azure Rights Management?

S/MIME

S/MIME for message signing and encryption

Understanding S/MIME

Understanding Public Key Cryptography

TLS

Configure custom mail flow by using connectors

Message Encryption

2/18/2021 • 7 minutes to read • Edit Online

People often use email to exchange sensitive information, such as financial data, legal contracts, confidential product information, sales reports and projections, patient health information, or customer and employee information. As a result, mailboxes can become repositories for large amounts of potentially sensitive information and information leakage can become a serious threat to your organization.

With Office 365 Message Encryption, your organization can send and receive encrypted email messages between people inside and outside your organization. Office 365 Message Encryption works with Outlook.com, Yahoo!, Gmail, and other email services. Email message encryption helps ensure that only intended recipients can view message content.

How Office 365 Message Encryption works

The rest of this article applies to the new OME capabilities.

Office 365 Message Encryption is an online service that's built on Microsoft Azure Rights Management (Azure RMS) which is part of Azure Information Protection. This includes encryption, identity, and authorization policies to help secure your email. You can encrypt messages by using rights management templates, the Do Not Forward option, and the encrypt-only option.

Users can then encrypt email messages and a variety of attachments by using these options. For a full list of supported attachment types, see "File types covered by IRM policies when they are attached to messages" in Introduction to IRM for email messages.

As an administrator, you can also define mail flow rules to apply this protection. For example, you can create a rule that requires the encryption of all messages addressed to a specific recipient, or that contains specific words in the subject line, and also specify that recipients can't copy or print the contents of the message.

Unlike the previous version of OME, the new capabilities provide a unified sender experience whether you're sending mail inside your organization or to recipients outside of Microsoft 365. In addition, recipients who receive a protected email message sent to a Microsoft 365 account in Outlook 2016 or Outlook on the web, don't have to take any additional action to view the message. It works seamlessly. Recipients using other email clients and email service providers also have an improved experience. For information, see Learn about protected messages in Office 365 and How do I open a protected message.

For a detailed list of the differences between the previous version of OME and the new OME capabilities, see Compare versions of OME.

When someone sends an email message that matches an encryption mail flow rule, the message is encrypted before it's sent. All Microsoft 365 end users that use Outlook clients to read mail receive native, first-class reading experiences for encrypted and rights-protected mail even if they're not in the same organization as the sender. Supported Outlook clients include Outlook desktop, Outlook Mac, Outlook mobile on iOS and Android, and Outlook on the web (formerly known as Outlook Web App).

Recipients of encrypted messages who receive encrypted or rights-protected mail sent to their Outlook.com, Gmail, and Yahoo accounts receive a wrapper mail that directs them to the OME Portal where they can easily authenticate using a Microsoft account, Gmail, or Yahoo credentials.

End users that read encrypted or rights-protected mail on clients other than Outlook also use the OME portal to view encrypted and rights-protected messages that they receive.

If the sender of the protected mail is in GCC High and the recipient is outside of GCC High, including commercial users, Outlook.com users, and users of other email providers such as Gmail, the recipient receives a wrapper mail. The wrapper mail directs the recipient to the OME Portal where the recipient is able to read and reply to the message. Otherwise, if the sender and recipient are both in the GCC High environment, even if they're not in the same organization, then recipients that use Outlook clients to read mail receive native, first-class reading experiences for encrypted and rights-protected mail. For more information about the different experience in GCC High, see Compare versions of OME.

For more information about size limits for messages and attachments that you can encrypt using OME, see Exchange Online Limits.

How Office 365 Advanced Message Encryption works on top of OME

Office 365 Advanced Message Encryption lets you create multiple branding templates so you can fine-tune control over recipient mail and create custom branding experiences to support a diverse organizational structure.

Advanced Message Encryption in Microsoft 365 helps you meet compliance obligations that require more flexible control over external recipient's access to encrypted emails. With Advanced Message Encryption in Office 365, as an administrator, you can control sensitive emails shared outside the organization with automatic policies that detect sensitive information types (for example, PII, Financial or Health IDs) or keywords to enhance protection by expiring access through a secure web portal to encrypted emails. As an admin you can further control encrypted emails accessed through a Microsoft 365 web portal by revoking access to an email anytime.

Message revocation and expiration only work for emails that your users send to recipients outside your organization. In addition, the recipients must access the email through the web portal. To ensure the recipient uses the portal to receive email, you set up a custom branding template that applies the wrapper. Then, you apply the branding template in a mail flow rule. For more information about Advanced Message Encryption, see Office 365 Advanced Message Encryption.

Defining rules for Office 365 Message Encryption

One way to enable the new capabilities for Office 365 Message Encryption is for Exchange Online and Exchange Online Protection administrators to define mail flow rules. These rules determine under what conditions email messages should be encrypted. When an encryption action is set for a rule, any messages that match the rule conditions are encrypted before they're sent.

Mail flow rules are flexible, letting you combine conditions so you can meet specific security requirements in a single rule. For example, you can create a rule to encrypt all messages that contain specified keywords and are addressed to external recipients. The new capabilities for Office 365 Message Encryption also encrypt replies from recipients of encrypted email.

For more information about how to create mail flow rules to take advantage of the new OME capabilities, see Define Rules for Office 365 Message Encryption.

Get started with the new OME capabilities

If you're ready to get started using the new OME capabilities within your organization, see Set up new Office 365 Message Encryption capabilities.

Sending, viewing, and replying to encrypted email messages

With Office 365 Message Encryption, users can send encrypted email from Outlook and Outlook on the web. Additionally, admins can set up mail flow rules in Microsoft 365 to automatically encrypt emails based on keyword matching or other conditions.

Recipients of encrypted messages who are in organizations will be able to read those messages seamlessly in any version Outlook, including Outlook for PC, Outlook for Mac, Outlook on the web, Outlook for iOS, and Outlook for Android. Users that receive encrypted messages on other email clients can view the messages in the OME portal.

For detailed guidance about how to send and view encrypted messages, take a look at these articles.

READ THIS ARTICLE...	IF YOU ARE...
Learn about protected messages in Office 365	An end user that wants to learn more about how encrypted messages work and what options are available to you.
How do I open a protected message?	An end user that wants to read a protected message that was sent to you. This article includes information about reading messages in several versions of Outlook and from different email accounts, including those accounts outside of Microsoft 365 such as gmail and Yahoo! accounts.
Send, view, and reply to encrypted messages in Outlook	An end user that wants to send, view, or reply to an encrypted message from Outlook. Even if you're not a member of an organization, you still receive notification of encrypted messages sent to you in Outlook. Use this article for instructions on how to view and reply to encrypted messages sent from Office 365.
Send a digitally signed or encrypted message	An end user that wants to send, view, or reply to encrypted messages using Outlook for Mac. This article also covers using encryption methods other than OME, such as S/MIME.
View encrypted messages on your Android device	An end user who has received a message encrypted with Office 365 Message Encryption on your Android device, you can use the free OME Viewer app to view the message and send an encrypted reply. This article explains how.
View encrypted messages on your iPhone or iPad	An end user who has received a message encrypted with Office 365 Message Encryption on your iPhone or iPad, you can use the free OME Viewer app to view the message and send an encrypted reply. This article explains how.

Set up new Message Encryption capabilities

11/2/2020 • 4 minutes to read • Edit Online

The new Office 365 Message Encryption (OME) capabilities allow organizations to share protected email with anyone on any device. Users can exchange protected messages with other Microsoft 365 organizations, as well as non-customers using Outlook.com, Gmail, and other email services.

Follow the steps below to ensure that the new OME capabilities are available in your organization.

Verify that Azure Rights Management is active

The new OME capabilities leverage the protection features in Azure Rights Management Services (Azure RMS), the technology used by Azure Information Protection to protect emails and documents via encryption and access controls.

The only prerequisite for using the new OME capabilities is that Azure Rights Management must be activated in your organization's tenant. If it is, Microsoft 365 activates the new OME capabilities automatically and you don't need to do anything.

Azure RMS is also activated automatically for most eligible plans, so you probably don't have to do anything in this regard either. See Activating Azure Rights Management for more information.

> **IMPORTANT**
>
> If you use Active Directory Rights Management service (AD RMS) with Exchange Online, you need to migrate to Azure Information Protection before you can use the new OME capabilities. OME is not compatible with AD RMS.

For more information, see:

- What subscriptions do I need to use the new OME capabilities? to check whether your subscription plan includes Azure Information Protection (which includes Azure RMS functionality).
- Azure Information Protection for information about purchasing an eligible subscription.

Manually activating Azure Rights Management

If you disabled Azure RMS, or if it was not automatically activated for any reason, you can activate it manually in the:

- **Microsoft 365 admin center**: See How to activate Azure Rights Management from the admin center for instructions.
- **Azure portal**: See How to activate Azure Rights Management from the Azure portal for instructions.

Configure management of your Azure Information Protection tenant key

This is an optional step. Allowing Microsoft to manage the root key for Azure Information Protection is the default setting and recommended best practice for most organizations. If this is the case, you don't need to do anything.

There are many reasons, for example compliance requirements, that may necessitate you generating and managing your own root key (also known as bring your own key (BYOK)). If this is the case, we recommend that you complete the required steps before setting up the new OME capabilities. See Planning and implementing

your Azure Information Protection tenant key for more.

Verify new OME configuration in Exchange Online PowerShell

You can verify that your Microsoft 365 tenant is properly configured to use the new OME capabilities in Exchange Online PowerShell.

1. Connect to Exchange Online PowerShell using an account with global administrator permissions in your Microsoft 365 tenant.
2. Run the Get-IRMConfiguration cmdlet.

   You should see a value of $True for the AzureRMSLicensingEnabled parameter, which indicates that OME is configured in your tenant. If it is not, use Set-IRMConfiguration to set the value of AzureRMSLicensingEnabled to $True to enable OME.
3. Run the Test-IRMConfiguration cmdlet using the following syntax:

   ```
   Test-IRMConfiguration [-Sender <email address >]
   ```

   Example:

   ```
   Test-IRMConfiguration -Sender securityadmin@contoso.com
   ```

   - Providing a sender email is optional, but forces the system to perform additional checks. Use the email address of any user in your Microsoft 365 tenant.

   Your results should be similar to:

   ```
   Results : Acquiring RMS Templates ...
             - PASS: RMS Templates acquired.  Templates available: Contoso  - Confidential View Only,
   Contoso  - Confidential, 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
   ```

   - Your organization name will replace *Contoso*.
   - The default template names may be different from those displayed above. See Configuring and managing templates for Azure Information Protection for more.
4. Run the Remove-PSSession cmdlet to disconnect from the Rights Management service.

   ```
   Remove-PSSession $session
   ```

Next steps: Define mail flow rules to use new OME capabilities

If there are previously configured mail flow rules to encrypt email in your organization, you need to update the existing rules to use the new OME capabilities. For new deployments, you need to create new mail flow rules.

IMPORTANT

If you do not update existing mail flow rules, your users will continue to receive encrypted mail that uses the previous HTML attachment format, instead of the new seamless OME experience.

Mail flow rules determine under what conditions email messages should be encrypted, as well as conditions for removing that encryption. When you set an action for a rule, any messages that match the rule conditions are encrypted when they're sent.

For steps on creating mail flow rules for OME, see Define mail flow rules to encrypt email messages in Office 365.

To update existing rules to use the new OME capabilities:

1. In the Microsoft 365 admin center, go to **Admin centers > Exchange**.
2. In the Exchange admin center, go to **Mail flow > Rules**.
3. For each rule, in **Do the following**:
   - Select **Modify the message security**.
   - Select **Apply Office 365 Message Encryption and rights protection**.
   - Select an RMS template from the list.
   - Select **Save**.
   - Select **OK**.

Define mail flow rules to encrypt email messages

11/2/2020 • 6 minutes to read • Edit Online

As a global administrator, you can create mail flow rules (also known as transport rules) to help protect email messages you send and receive. You can set up rules to encrypt any outgoing email messages and remove encryption from encrypted messages coming from inside your organization or from replies to encrypted messages sent from your organization. You can use the Exchange admin center (EAC) or Exchange Online PowerShell to create these rules. In addition to overall encryption rules, you can also choose to enable or disable individual message encryption options for end users.

You can't encrypt inbound mail from senders outside of your organization.

If you recently migrated from Active Directory RMS to Azure Information Protection, you'll need to review your existing mail flow rules to ensure that they continue to work in your new environment. Also, if you want to take advantage of the new Office 365 Message Encryption (OME) capabilities available to you through Azure Information Protection, you need to update your existing mail flow rules. Otherwise, your users will continue to receive encrypted mail that uses the previous HTML attachment format instead of the new, seamless OME experience. If you haven't set up OME yet, see Set up new Office 365 Message Encryption capabilities for information.

For information about the components that make up mail flow rules and how mail flow rules work, see Mail flow rules (transport rules) in Exchange Online. For additional information about how mail flow rules work with Azure Information Protection, see Configuring Exchange Online mail flow rules for Azure Information Protection labels.

> **IMPORTANT**
>
> For hybrid Exchange environments, on-premises users can send and receive encrypted mail using OME only if email is routed through Exchange Online. To configure OME in a hybrid Exchange environment, you need to first configure hybrid using the Hybrid Configuration wizard and then configure mail to flow from Office 365 to your email server and configure mail to flow from your email server to Office 365. Once you've configured mail to flow through Office 365, then you can configure mail flow rules for OME by using this guidance.

Create mail flow rules to encrypt email messages with the new OME capabilities

You can define mail flow rules for triggering message encryption with the new OME capabilities by using the EAC.

Use the EAC to create a rule for encrypting email messages with the new OME capabilities

1. In a web browser, using a work or school account that has been granted global administrator permissions, sign in to Office 365.
2. Choose the **Admin** tile.
3. In the Microsoft 365 admin center, choose **Admin centers** > **Exchange**.
4. In the EAC, go to **Mail flow** > **Rules** and select **New +** > **Create a new rule**. For more information about using the EAC, see Exchange admin center in Exchange Online.
5. In **Name**, type a name for the rule, such as Encrypt mail for DrToniRamos@hotmail.com.

6. In **Apply this rule if**, select a condition, and enter a value if necessary. For example, to encrypt messages going to DrToniRamos@hotmail.com:

   a. In **Apply this rule if**, select **the recipient is**.

   b. Select an existing name from the contact list or type a new email address in the **check names** box.

      - To select an existing name, select it from the list and then click **OK**.
      - To enter a new name, type an email address in the **check names** box and then select **check names** > **OK**.

7. To add more conditions, choose **More options** and then choose **add condition** and select from the list.

   For example, to apply the rule only if the recipient is outside your organization, select **add condition** and then select **The recipient is external/internal** > **Outside the organization** > **OK**.

8. To enable encryption using the new OME capabilities, from **Do the following**, select **Modify the message security** and then choose **Apply Office 365 Message Encryption and rights protection**. Select an RMS template from the list, choose **Save**, and then choose **OK**.

The list of templates includes all default templates and options as well as any custom templates you've created for use by Office 365. If the list is empty, ensure that you have set up Office 365 Message Encryption with the new capabilities as described in Set up new Office 365 Message Encryption capabilities. For information about the default templates, see Configuring and managing templates for Azure Information Protection. For information about the **Do Not Forward** option, see Do Not Forward option for emails. For information about the **encrypt only** option, see Encrypt Only option for emails.

You can choose **add action** if you want to specify another action.

Use the EAC to update an existing mail flow rule to use the new OME capabilities

1. In a web browser, using a work or school account that has been granted global administrator permissions, sign in to Office 365.
2. Choose the **Admin** tile.
3. In the Microsoft 365 admin center, choose **Admin centers** > **Exchange**.
4. In the EAC, go to **Mail flow** > **Rules**.
5. In the list of mail flow rules, select the rule you want to modify to use the new OME capabilities and then choose **Edit** ✎.
6. To enable encryption using the new OME capabilities, from **Do the following**, choose **Modify the message security** and then choose **Apply Office 365 Message Encryption and rights protection**. Select an RMS template from the list, choose **Save** and then choose **OK**.

   The list of templates includes all default templates and options as well as any custom templates you've created for use by Office 365. If the list is empty, ensure that you have set up Office 365 Message Encryption with the new capabilities as described in Set up new Office 365 Message Encryption capabilities built on top of Azure Information Protection. For information about the default templates, see Configuring and managing templates for Azure Information Protection. For information about the **Do Not Forward** option, see Do Not Forward option for emails. For information about the **encrypt only** option, see Encrypt Only option for emails.

   You can choose **add action** if you want to specify another action.

7. From the **Do the following** list, remove any actions that are assigned to **Modify the message security** > **Apply the previous version of OME**.

8. Choose **Save**.

Create mail flow rules to remove encryption for outgoing email messages with the new OME capabilities

You can define mail flow rules for triggering remove message encryption with the new OME capabilities by using the EAC.

Use the EAC to create a rule to remove encryption from email messages with the new OME capabilities

1. In a web browser, using a work or school account that has been granted global administrator permissions, sign in to Office 365.
2. Choose the **Admin** tile.
3. In the Microsoft 365 admin center, choose **Admin centers** > **Exchange**.
4. In the EAC, go to **Mail flow** > **Rules** and select **New +** > **Create a new rule**. For more information about using the EAC, see Exchange admin center in Exchange Online.
5. In **Name**, type a name for the rule, such as Remove encryption from outgoing mail.
6. In **Apply this rule if**, select the conditions where encryption should be removed from messages. Add **The sender is located** > **Inside the organization**. Now add additional conditions to target specific recipients, such as **The recipient is located** > **Outside the organization**.
7. In **Do the following**, select **Modify the message security** > **Remove Office 365 Message Encryption and rights protection**.
8. Select **Save**.

Create mail flow rules for Office 365 Message Encryption without the new capabilities

If you haven't yet moved your organization to the new OME capabilities, Microsoft recommends that you make a plan to move to the new OME capabilities as soon as it is reasonable for your organization. For instructions, see Set up new Office 365 Message Encryption capabilities built on top of Azure Information Protection. Otherwise, see Defining mail flow rules for Office 365 Message Encryption that don't use the new OME capabilities.

Related Topics

Encryption in Office 365

Set up new Office 365 Message Encryption capabilities

Add branding to encrypted messages

Mail flow rules (transport rules) in Exchange Online

Mail flow rules (transport rules) in Exchange Online Protection

Add your organization's brand to your Microsoft 365 for business Message Encryption encrypted messages

2/18/2021 • 10 minutes to read • Edit Online

You can apply your company branding to customize the look of your organization's email messages and the encryption portal. You'll need to apply global administrator permissions to your work or school account before you can get started. Once you have these permissions, use the Get-OMEConfiguration and Set-OMEConfiguration Windows PowerShell cmdlets to customize these parts of encrypted email messages:

- Introductory text
- Disclaimer text
- URL for Your organization's privacy statement
- Text in the OME portal
- Logo that appears in the email message and OME portal, or whether to use a logo at all
- Background color in the email message and OME portal

You can also revert back to the default look and feel at any time.

If you'd like more control, use Office 365 Advanced Message Encryption to create multiple templates for encrypted emails originating from your organization. Use these templates to control parts of the end-user experience. For example, specify whether recipients can use Google, Yahoo, and Microsoft Accounts to sign in to the encryption portal. Use templates to fulfill several use cases, such as:

- Individual departments, such as Finance, Sales, and so on.
- Different products
- Different geographical regions or countries
- Whether you want to allow emails to be revoked
- Whether you want emails sent to external recipients to expire after a specified number of days.

Once you've created the templates, you can apply them to encrypted emails by using Exchange mail flow rules. If you have Office 365 Advanced Message Encryption, you can revoke any email that you've branded by using these templates.

Work with OME branding templates

You can modify several features within a branding template. You can modify, but not remove, the default template. If you have Advanced Message Encryption, you can also create, modify, and remove custom templates. Use Windows PowerShell to work with one branding template at a time.

- Set-OMEConfiguration - Modify the default branding template or a custom branding template that you created.
- New-OMEConfiguration - Create a new branding template, Advanced Message Encryption only.
- Remove-OMEConfiguration - Remove a custom branding template, Advanced Message Encryption only. You

can't delete the default branding template.

Modify an OME branding template

Use Windows PowerShell to modify one branding template at a time. If you have Advanced Message Encryption, you can also create, modify, and remove custom templates.

1. Using a work or school account that has global administrator permissions in your organization, start a Windows PowerShell session and connect to Exchange Online. For instructions, see Connect to Exchange Online PowerShell.
2. Use the Set-OMEConfiguration cmdlet as described in Set-OMEConfiguration or use the following graphic and table for guidance.

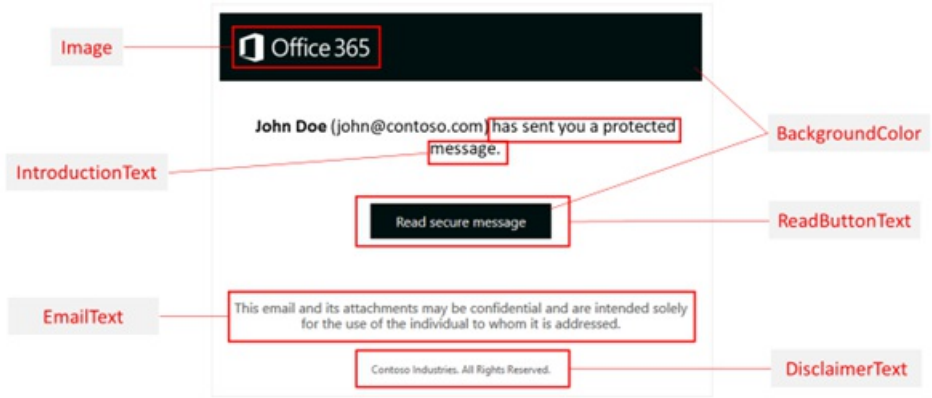

TO CUSTOMIZE THIS FEATURE OF THE ENCRYPTION EXPERIENCE	USE THESE COMMANDS

Background color	`Set-OMEConfiguration -Identity "<OMEConfigurationName>" -BackgroundColor "<#RRGGBB hexadecimal color code or name value>"`
	Example:
	`Set-OMEConfiguration -Identity "Branding Template 1" -BackgroundColor "#ffffff"`
	For more information about background colors, see the Background colors section later in this article.

Logo	`Set-OMEConfiguration -Identity "<OMEConfigurationName>" -Image <Byte[]>`
	Example:
	`Set-OMEConfiguration -Identity "Branding Template 1" -Image (Get-Content "C:\Temp\contosologo.png" -Encoding byte)`
	Supported file formats: .png, .jpg, .bmp, or .tiff
	Optimal size of logo file: less than 40 KB
	Optimal size of logo image: 170x70 pixels. If your image exceeds these dimensions, the service resizes your logo for display in the portal. The service doesn't modify the graphic file itself. For best results, use the optimal size.

Text next to the sender's name and email address	`Set-OMEConfiguration -Identity "<OMEConfigurationName>" -IntroductionText "<String up to 1024 characters>"`
	Example:
	`Set-OMEConfiguration -Identity "Branding Template 1" -IntroductionText "has sent you a secure message."`

TO CUSTOMIZE THIS FEATURE OF THE ENCRYPTION EXPERIENCE	USE THESE COMMANDS

Text that appears on the "Read Message" button	`Set-OMEConfiguration -Identity "<OMEConfigurationName>" -ReadButtonText "<String up to 1024 characters>"`
	Example:
	`Set-OMEConfiguration -Identity "OME Configuration" -ReadButtonText "Read Secure Message."`

Text that appears below the "Read Message" button	`Set-OMEConfiguration -Identity "<OMEConfigurationName>" -EmailText "<String up to 1024 characters>"`
	Example:
	`Set-OMEConfiguration -Identity "OME Configuration" -EmailText "Encrypted message from ContosoPharma secure messaging system."`

URL for the Privacy Statement link	`Set-OMEConfiguration -Identity "<OMEConfigurationName>" -PrivacyStatementURL "<URL>"`
	Example:
	`Set-OMEConfiguration -Identity "Branding Template 1" -PrivacyStatementURL "https://contoso.com/privacystatement.html"`

Disclaimer statement in the email that contains the encrypted message	`Set-OMEConfiguration -Identity "<OMEConfigurationName>" -DisclaimerText "<Disclaimer statement. String of up to 1024 characters.>"`
	Example:
	`Set-OMEConfiguration -Identity "Branding Template 1" -DisclaimerText "This message is confidential for the use of the addressee only."`

Text that appears at the top of the encrypted mail viewing portal	`Set-OMEConfiguration -Identity "<OMEConfigurationName>" -PortalText "<Text for your portal. String of up to 128 characters.>"`
	Example:
	`Set-OMEConfiguration -Identity "OME Configuration" -PortalText "ContosoPharma secure email portal."`

To enable or disable authentication with a one-time pass code for this custom template	`Set-OMEConfiguration -Identity "<OMEConfigurationName>" -OTPEnabled <$true\|$false>`
	Examples:
	To enable one-time passcodes for this custom template
	`Set-OMEConfiguration -Identity "Branding Template 1" -OTPEnabled $true`
	To disable one-time passcodes for this custom template
	`Set-OMEConfiguration -Identity "Branding Template 1" -OTPEnabled $false`

To enable or disable authentication with Microsoft, Google, or Yahoo identities for this custom template	`Set-OMEConfiguration -Identity "<OMEConfigurationName>" -SocialIdSignIn <$true\|$false>`
	Examples:
	To enable social IDs for this custom template
	`Set-OMEConfiguration -Identity "Branding Template 1" -SocialIdSignIn $true`
	To disable social IDs for this custom template
	`Set-OMEConfiguration -Identity "Branding Template 1" -SocialIdSignIn $false`

Create an OME branding template (Advanced Message Encryption)

If you have Office 365 Advanced Message Encryption, you can create custom branding templates for your organization by using the New-OMEConfiguration cmdlet. Once you've created the template, you modify the template by using the Set-OMEConfiguration cmdlet as described in Modify an OME branding template. You can create multiple templates.

To create a new custom branding template:

1. Using a work or school account that has global administrator permissions in your organization, start a Windows PowerShell session and connect to Exchange Online. For instructions, see Connect to Exchange Online PowerShell.
2. Use the New-OMEConfiguration cmdlet to create a new template.

   ```
   New-OMEConfiguration -Identity "<OMEConfigurationName>"
   ```

   For example,

   ```
   New-OMEConfiguration -Identity "Custom branding template"
   ```

Return the default branding template to its original values

To remove all modifications from the default template, including brand customizations, and so on, complete these steps:

1. Using a work or school account that has global administrator permissions in your organization, start a Windows PowerShell session and connect to Exchange Online. For instructions, see Connect to Exchange Online PowerShell.
2. Use the **Set-OMEConfiguration** cmdlet as described in Set-OMEConfiguration. To remove your organization's branded customizations from the DisclaimerText, EmailText, and PortalText values, set the value to an empty string, `""`. For all image values, such as Logo, set the value to `"$null"`.

   The following table describes the encryption customization option defaults.

USE THESE COMMANDS	

Default text that comes with encrypted email messages The default text appears above the instructions for viewing encrypted messages	`Set-OMEConfiguration -Identity "<OMEConfigurationName>" -EmailText "<empty string>"`
	Example:
	`Set-OMEConfiguration -Identity "OME Configuration" -EmailText ""`

Disclaimer statement in the email that contains the encrypted message	`Set-OMEConfiguration -Identity "<OMEConfigurationName>" DisclaimerText "<empty string>"`
	Example:
	`Set-OMEConfiguration -Identity "OME Configuration" -DisclaimerText ""`

Text that appears at the top of the encrypted mail viewing portal	`Set-OMEConfiguration -Identity "<OMEConfigurationName>" -PortalText "<empty string>"`
	Example reverting back to default:
	`Set-OMEConfiguration -Identity "OME Configuration" -PortalText ""`

USE THESE COMMANDS	

Logo	`Set-OMEConfiguration -Identity "<OMEConfigurationName>" -Image <"$null">`
	Example reverting back to default:
	`Set-OMEConfiguration -Identity "OME configuration" -Image $null`

Background color	`Set-OMEConfiguration -Identity "<OMEConfigurationName>" -BackgroundColor "$null">`
	Example reverting back to default:
	`Set-OMEConfiguration -Identity "OME configuration" -BackgroundColor $null`

Remove a custom branding template (Advanced Message Encryption)

You can only remove or delete branding templates that you've made. You can't remove the default branding template.

To remove a custom branding template:

1. Using a work or school account that has global administrator permissions in your organization, start a Windows PowerShell session and connect to Exchange Online. For instructions, see Connect to Exchange Online PowerShell.
2. Use the **Remove-OMEConfiguration** cmdlet as follows:

   ```
   Remove-OMEConfiguration -Identity ""<OMEConfigurationName>"
   ```

   For example,

   ```
   Remove-OMEConfiguration -Identity "Branding template 1"
   ```

   For more information, see Remove-OMEConfiguration.

Create an Exchange mail flow rule that applies your custom branding to encrypted emails

After you've either modified the default template or created new branding templates, you can create Exchange mail flow rules to apply your custom branding based on certain conditions. Such a rule will apply custom branding in the following scenarios:

- If the email was manually encrypted by the end user using Outlook or Outlook on the web, formerly Outlook Web App
- If the email was automatically encrypted by an Exchange mail flow rule or Data Loss Prevention policy

For information on how to create an Exchange mail flow rule that applies encryption, see Define mail flow rules to encrypt email messages in Office 365.

1. In a web browser, using a work or school account that has been granted global administrator permissions, sign in to Office 365.
2. Choose the **Admin** tile.

3. In the Microsoft 365 admin center, choose **Admin centers** > **Exchange**.
4. In the EAC, go to **Mail flow** > **Rules** and select **New +** > **Create a new rule**. For more information about using the EAC, see Exchange admin center in Exchange Online.
5. In **Name**, type a name for the rule, such as Branding for sales department.
6. In **Apply this rule if**, select the condition **The sender is located inside the organization** and other conditions you want from the list of available conditions. For example, you might want to apply a particular branding template to:
   - All encrypted emails sent from members of the finance department
   - Encrypted emails sent with a certain keyword such as "External" or "Partner"
   - Encrypted emails sent to a particular domain
7. From **Do the following**, select **Modify the message security** > **Apply custom branding to OME messages**. Next, from the drop-down, select a branding template.
8. (Optional) You can configure the mail flow rule to apply encryption and custom branding. From **Do the following**, select **Modify the message security**, and then choose **Apply Office 365 Message Encryption and rights protection**. Select an RMS template from the list, choose **Save**, and then choose **OK**.

   The list of templates includes default templates and options and any custom templates you create. If the list is empty, ensure that you have set up Office 365 Message Encryption with the new capabilities. For instructions, see Set up new Office 365 Message Encryption capabilities. For information about the default templates, see Configuring and managing templates for Azure Information Protection. For information about the **Do Not Forward** option, see Do Not Forward option for emails. For information about the **encrypt only** option, see Encrypt Only option for emails.

   Choose **add action** if you want to specify another action.

Background color reference

The color names that you can use for the background color are limited. Instead of a color name, you can use a hex code value (#RRGGBB). You can use a hex code value that corresponds to a color name, or you can use a custom hex code value. Be sure to enclose the hex code value in quotation marks (for example, `"#f0f8ff"`).

The available background color names and their corresponding hex code values are described in the following table.

COLOR NAME	COLOR CODE
`aliceblue`	#f0f8ff
`antiquewhite`	#faebd7
`aqua`	#00ffff
`aquamarine`	#7fffd4
`azure`	#f0ffff
`beige`	#f5f5dc
`bisque`	#ffe4c4

COLOR NAME	COLOR CODE
`black`	#000000
`blanchedalmond`	#ffebcd
`blue`	#0000ff
`blueviolet`	#8a2be2
`brown`	#a52a2a
`burlywood`	#deb887
`cadetblue`	#5f9ea0
`chartreuse`	#7fff00
`chocolate`	#d2691e
`coral`	#ff7f50
`cornflowerblue`	#6495ed
`cornsilk`	#fff8dc
`crimson`	#dc143c
`cyan`	#00ffff
`darkblue`	#00008b
`darkcyan`	#008b8b
`darkgoldenrod`	#b8860b
`darkgray`	#a9a9a9
`darkgreen`	#006400
`darkkhaki`	#bdb76b
`darkmagenta`	#8b008b
`darkolivegreen`	#556b2f
`darkorange`	#ff8c00

COLOR NAME	COLOR CODE
`darkorchid`	#9932cc
`darkred`	#8b0000
`darksalmon`	#e9967a
`darkseagreen`	#8fbc8f
`darkslateblue`	#483d8b
`darkslategray`	#2f4f4f
`darkturquoise`	#00ced1
`darkviolet`	#9400d3
`deeppink`	#ff1493
`deepskyblue`	#00bfff
`dimgray`	#696969
`dodgerblue`	#1e90ff
`firebrick`	#b22222
`floralwhite`	#fffaf0
`forestgreen`	#228b22
`fuchsia`	#ff00ff
`gainsboro`	#dcdcdc
`ghostwhite`	#f8f8ff
`gold`	#ffd700
`goldenrod`	#daa520
`gray`	#808080
`green`	#008000
`greenyellow`	#adff2f
`honeydew`	#f0fff0

COLOR NAME	COLOR CODE
`hotpink`	#ff69b4
`indianred`	#cd5c5c
`indigo`	#4b0082
`ivory`	#fffff0
`khaki`	#f0e68c
`lavender`	#e6e6fa
`lavenderblush`	#fff0f5
`lawngreen`	#7cfc00
`lemonchiffon`	#fffacd
`lightblue`	#add8e6
`lightcoral`	#f08080
`lightcyan`	#e0ffff
`lightgoldenrodyellow`	#fafad2
`lightgray`	#d3d3d3
`lightgrey`	#d3d3d3
`lightgreen`	#90ee90
`lightpink`	#ffb6c1
`lightsalmon`	#ffa07a
`lightseagreen`	#20b2aa
`lightskyblue`	#87cefa
`lightslategray`	#778899
`lightsteelblue`	#b0c4de
`lightyellow`	#ffffe0
`lime`	#00ff00

COLOR NAME	COLOR CODE
`limegreen`	#32cd32
`linen`	#faf0e6
`magenta`	#ff00ff
`maroon`	#800000
`mediumaquamarine`	#66cdaa
`mediumblue`	#0000cd
`mediumorchid`	#ba55d3
`mediumpurple`	#9370db
`mediumseagreen`	#3cb371
`mediumslateblue`	#7b68ee
`mediumspringgreen`	#00fa9a
`mediumturquoise`	#48d1cc
`mediumvioletred`	#c71585
`midnightblue`	#191970
`mintcream`	#f5fffa
`mistyrose`	#ffe4e1
`moccasin`	#ffe4b5
`navajowhite`	#ffdead
`navy`	#000080
`oldlace`	#fdf5e6
`olive`	#808000
`olivedrab`	#6b8e23
`orange`	#ffa500
`orangered`	#ff4500

COLOR NAME	COLOR CODE
`orchid`	#da70d6
`palegoldenrod`	#eee8aa
`palegreen`	#98fb98
`paleturquoise`	#afeeee
`palevioletred`	#db7093
`papayawhip`	#ffefd5
`peachpuff`	#ffdab9
`peru`	#cd853f
`pink`	#ffc0cb
`plum`	#dda0dd
`powderblue`	#b0e0e6
`purple`	#800080
`red`	#ff0000
`rosybrown`	#bc8f8f
`royalblue`	#4169e1
`saddlebrown`	#8b4513
`salmon`	#fa8072
`sandybrown`	#f4a460
`seagreen`	#00ff00
`seashell`	#fff5ee
`sienna`	#a0522d
`silver`	#c0c0c0
`skyblue`	#87ceeb
`slateblue`	#6a5acd

COLOR NAME	COLOR CODE
`slategray`	#708090
`snow`	#fffafa
`springgreen`	#00ff7f
`steelblue`	#4682b4
`tan`	#d2b48c
`teal`	#008080
`thistle`	#d8bfd8
`tomato`	#ff6347
`turquoise`	#40e0d0
`violet`	#ee82ee
`wheat`	#f5deb3
`white`	#ffffff
`whitesmoke`	#f5f5f5
`yellow`	#ffff00
`yellowgreen`	#9acd32

Create a sensitive information type policy for your organization using Message Encryption

11/2/2020 • 2 minutes to read • Edit Online

You can use either Exchange mail flow rules or Data Loss Prevention (DLP) to create a sensitive information type policy with Office 365 Message Encryption. To create an Exchange mail flow rule, you can use either the Exchange admin center (EAC) or PowerShell.

To create the policy by using mail flow rules in the EAC

Sign in to the Exchange admin center (EAC) and go to **Mail flow** > **Rules**. On the Rules page, create a rule that applies Office 365 Message Encryption. You can create a rule based on conditions such as the presence of certain keywords or sensitive information types in the message or attachment.

To create the policy by using mail flow rules in PowerShell

Use a work or school account that has global administrator permissions in your organization, start a Windows PowerShell session and connect to Exchange Online. For instructions, see Connect to Exchange Online PowerShell. Use the Set-IRMConfiguration and New-TransportRule cmdlets to create the policy.

Example mail flow rule created with PowerShell

Run the following commands in PowerShell to create an Exchange mail flow rule that automatically encrypts emails sent outside your organization with the *Encrypt-Only* policy if the emails or their attachments contain the following sensitive information types:

- ABA routing number
- Credit card Number
- Drug Enforcement Agency (DEA) number
- U.S. / U.K. passport number
- U.S. bank account number
- U.S. Individual Taxpayer Identification Number (ITIN)
- U.S. Social Security Number (SSN)

```
Set-IRMConfiguration -DecryptAttachmentForEncryptOnly $true
New-TransportRule -Name "Encrypt outbound sensitive emails (out of box rule)" -SentToScope
NotInOrganization  -ApplyRightsProtectionTemplate "Encrypt" -MessageContainsDataClassifications
@(@{Name="ABA Routing Number"; minCount="1"},@{Name="Credit Card Number"; minCount="1"},@{Name="Drug
Enforcement Agency (DEA) 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"}) -
SenderNotificationType "NotifyOnly"
```

For more information, see Set-IRMConfiguration and New-TransportRule.

How recipients access attachments

After Microsoft encrypts a message, recipients have unrestricted access to attachments when they access and open their encrypted email.

To prepare for this change

You may want to update any applicable end-user documentation and training materials to prepare people in your organization for this change. Share these Office 365 Message Encryption resources with your users as appropriate:

- Send, view, and reply to encrypted messages in Outlook for PC
- Microsoft 365 Essentials Video: Office Message Encryption

View these changes in the audit log

Microsoft 365 audits this activity and makes it available to administrators. The operation is 'New-TransportRule' and a snippet of a sample audit entry from the Audit Log Search in Security & Compliance Center is below:


```
*{"CreationTime":"2018-11-28T23:35:01","Id":"a1b2c3d4-daa0-4c4f-a019-03a1234a1b0c","Operation":"New-
TransportRule","OrganizationId":"123456-221d-12345
","RecordType":1,"ResultStatus":"True","UserKey":"Microsoft
Operator","UserType":3,"Version":1,"Workload":"Exchange","ClientIP":"123.456.147.68:17584","ObjectId":"","Us
erId":"Microsoft
Operator","ExternalAccess":true,"OrganizationName":"contoso.onmicrosoft.com","OriginatingServer":"CY4PR13MBX
XXX (15.20.1382.008)","Parameters": {"Name":"Organization","Value":"123456-221d-12346"
{"Name":"ApplyRightsProtectionTemplate","Value":"Encrypt"},{"Name":"Name","Value":"Encrypt outbound
sensitive emails (out of box rule)"},{"Name":"MessageContainsDataClassifications"…etc.*
```


To disable or customize the sensitive information types policy

Once you've created the Exchange mail flow rule, you can disable or edit the rule by going to **Mail flow** > **Rules** in the Exchange admin center (EAC) and disable the rule "*Encrypt outbound sensitive emails (out of box rule)*".

Manage Office 365 Message Encryption

11/2/2020 • 8 minutes to read • Edit Online

Once you've finished setting up Office 365 Message Encryption (OME), you can customize the configuration of your deployment in several ways. For example, you can configure whether to enable one-time pass codes, display the **Encrypt** button in Outlook on the web, and more. The tasks in this article describe how.

Manage whether Google, Yahoo, and Microsoft Account recipients can use these accounts to sign in to the Office 365 Message Encryption portal

When you set up the new Office 365 Message Encryption capabilities, users in your organization can send messages to recipients that are outside of your organization. If the recipient uses a *social ID* such as a Google account, Yahoo account, or Microsoft account, the recipient can sign in to the OME portal with a social ID. If you want, you can choose not to allow recipients to use social IDs to sign in to the OME portal.

To manage whether recipients can use social IDs to sign in to the OME portal

1. Connect to Exchange Online Using Remote PowerShell.
2. Run the Set-OMEConfiguration cmdlet with the SocialIdSignIn parameter as follows:

```
Set-OMEConfiguration -Identity <"OMEConfigurationIdParameter"> -SocialIdSignIn <$true|$false>
```

   For example, to disable social IDs:

```
Set-OMEConfiguration -Identity "OME Configuration" -SocialIdSignIn $false
```

   To enable social IDs:

```
Set-OMEConfiguration -Identity "OME Configuration" -SocialIdSignIn $true
```

Manage the use of one-time pass codes for the Office 365 Message Encryption portal

If the recipient of a message encrypted by OME doesn't use Outlook, regardless of the account used by the recipient, the recipient receives a limited-time web-view link that lets them read the message. This link includes a one-time pass code. As an administrator, you can decide if recipients can use one-time pass codes to sign in to the OME portal.

To manage whether OME generates one-time pass codes

1. Use a work or school account that has global administrator permissions in your organization and start a Windows PowerShell session and connect to Exchange Online. For instructions, see Connect to Exchange Online PowerShell.
2. Run the Set-OMEConfiguration cmdlet with the OTPEnabled parameter:

```
Set-OMEConfiguration -Identity <"OMEConfigurationIdParameter "> -OTPEnabled <$true|$false>
```

For example, to disable one-time pass codes:

```
Set-OMEConfiguration -Identity "OME Configuration" -OTPEnabled $false
```

To enable one-time pass codes:

```
Set-OMEConfiguration -Identity "OME Configuration" -OTPEnabled $true
```

Manage the display of the Encrypt button in Outlook on the web

As an administrator, you can manage whether to display this button to end users.

To manage whether the Encrypt button appears in Outlook on the web

1. Use a work or school account that has global administrator permissions in your organization and start a Windows PowerShell session and connect to Exchange Online. For instructions, see Connect to Exchange Online PowerShell.
2. Run the Set-IRMConfiguration cmdlet with the -SimplifiedClientAccessEnabled parameter:

   ```
   Set-IRMConfiguration -SimplifiedClientAccessEnabled <$true|$false>
   ```

   For example, to disable the **Encrypt** button:

   ```
   Set-IRMConfiguration -SimplifiedClientAccessEnabled $false
   ```

   To enable the **Encrypt** button:

   ```
   Set-IRMConfiguration -SimplifiedClientAccessEnabled $true
   ```

Enable service-side decryption of email messages for iOS mail app users

The iOS mail app can't decrypt messages protected with Office 365 Message Encryption. As a Microsoft 365 administrator, you can apply service-side decryption for messages delivered to the iOS mail app. When you choose to do use service-side decryption, the service sends a decrypted copy of the message to the iOS device. The client device stores a decrypted copy of the message. The message also retains information about usage rights even though the iOS mail app doesn't apply client-side usage rights to the user. The user can copy or print the message even if they didn't originally have the rights to do so. However, if the user attempts to complete an action that requires the Microsoft 365 mail server, such as forwarding the message, the server won't permit the action if the user didn't originally have the usage right to do so. However, end users can work around "Do Not Forward" usage restriction by forwarding the message from a different account within the iOS mail app. Regardless of whether you set up service-side decryption of mail, attachments to encrypted and rights protected mail can't be viewed in the iOS mail app.

If you choose not to allow decrypted messages to be sent to iOS mail app users, users receive a message that states that they don't have the rights to view the message. By default, service-side decryption of email messages is not enabled.

For more information, and for a view of the client experience, see View encrypted messages on your iPhone or iPad.

To manage whether iOS mail app users can view messages protected by Office 365 Message Encryption

1. Use a work or school account that has global administrator permissions in your organization and start a Windows PowerShell session and connect to Exchange Online. For instructions, see Connect to Exchange Online PowerShell.

2. Run the Set-ActiveSyncOrganizations cmdlet with the AllowRMSSupportForUnenlightenedApps parameter:

   ```
   Set-ActiveSyncOrganizationSettings -AllowRMSSupportForUnenlightenedApps <$true|$false>
   ```

   For example, to configure the service to decrypt messages before they're sent to unenlightened apps like the iOS mail app:

   ```
   Set-ActiveSyncOrganizationSettings -AllowRMSSupportForUnenlightenedApps $true
   ```

   Or, to configure the service not to send decrypted messages to unenlightened apps:

   ```
   Set-ActiveSyncOrganizationSettings -AllowRMSSupportForUnenlightenedApps $false
   ```

> **NOTE**
> Individual mailbox policies (OWA/ActiveSync) override these settings (i.e. if -IRMEnabled is set to False within the respective OWA Mailbox policy, or ActiveSync Mailbox policy, then these configurations would not apply).

Enable service-side decryption of email attachments for web browser mail clients

Normally, when you use Office 365 message encryption, attachments are automatically encrypted. As an administrator, you can apply service-side decryption for email attachments that users download from a web browser.

When you use service-side decryption, the service sends a decrypted copy of the file to the device. The message is still encrypted. The email attachment also keeps information about usage rights even though the browser doesn't apply client-side usage rights to the user. The user can copy or print the email attachment even if they didn't originally have the rights to do so. However, if the user tries to complete an action that requires the Microsoft 365 mail server, such as forwarding the attachment, the server won't permit the action if the user didn't originally have the usage right to do so.

Regardless of whether you set up service-side decryption of attachments, users can't view any attachments to encrypted and rights protected mail in the iOS mail app.

If you choose not to allow decrypted email attachments, which is the default, users receive a message that states that they don't have the rights to view the attachment.

For more information about how Microsoft 365 implements encryption for emails and email attachments with the Encrypt-Only option, see Encrypt-Only option for emails.

To manage whether email attachments are decrypted on download from a web browser

1. Use a work or school account that has global administrator permissions in your organization and start a

Windows PowerShell session and connect to Exchange Online. For instructions, see Connect to Exchange Online PowerShell.

2. Run the Set-IRMConfiguration cmdlet with the DecryptAttachmentForEncryptOnly parameter:

```
Set-IRMConfiguration -DecryptAttachmentForEncryptOnly <$true|$false>
```

For example, to configure the service to decrypt email attachments when a user downloads them from a web browser:

```
Set-IRMConfiguration -DecryptAttachmentForEncryptOnly $true
```

To configure the service to leave encrypted email attachments as they are upon download:

```
Set-IRMConfiguration -DecryptAttachmentForEncryptOnly $false
```

Ensure all external recipients use the OME Portal to read encrypted mail

You can use custom branding templates to force recipients to receive a wrapper mail that directs them to read encrypted email in the OME Portal instead of using Outlook or Outlook on the web. You might want to do this if you use want greater control over how recipients use the mail they receive. For example, if external recipients view email in the web portal, you can set an expiration date for the email, and you can revoke the email. These features are only supported through the OME Portal. You can use the Encrypt option and the Do Not Forward option when creating the mail flow rules.

Use a custom template to force all external recipients to use the OME Portal and for encrypted email

1. Use a work or school account that has global administrator permissions in your organization and start a Windows PowerShell session and connect to Exchange Online. For instructions, see Connect to Exchange Online PowerShell.

2. Run the New-TransportRule cmdlet:

```
New-TransportRule -name "<mail flow rule name>" -FromScope "InOrganization" -
ApplyRightsProtectionTemplate "<option name>" -ApplyRightsProtectionCustomizationTemplate "<template
name>"
```

where:

- `mail flow rule name` is the name you want to use for the new mail flow rule.
- `option name` is either `Encrypt` or `Do Not Forward`.
- `template name` is the name you gave the custom branding template, for example `OME Configuration`.

To encrypt all external email with the "OME Configuration" template and apply the Encrypt-Only option:

```
New-TransportRule -name "<All outgoing mail>" -FromScope "InOrganization" -
ApplyRightsProtectionTemplate "Encrypt" -ApplyRightsProtectionCustomizationTemplate "OME
Configuration"
```

To encrypt all external email with the "OME Configuration" template and apply the Do Not Forward

option:

```
New-TransportRule -name "<All outgoing mail>" -FromScope "InOrganization" -
ApplyRightsProtectionTemplate "Do Not Forward" -ApplyRightsProtectionCustomizationTemplate "OME
Configuration"
```

Customize the appearance of email messages and the OME portal

For detailed information about how you can customize OME for your organization, see Add your organization's brand to your encrypted messages.

Disable the new capabilities for OME

We hope it doesn't come to it, but if you need to, disabling the new capabilities for OME is very straightforward. First, you'll need to remove any mail flow rules you've created that use the new OME capabilities. For information about removing mail flow rules, see Manage mail flow rules. Then, complete these steps in Exchange Online PowerShell.

To disable the new capabilities for OME

1. Using a work or school account that has global administrator permissions in your organization, start a Windows PowerShell session and connect to Exchange Online. For instructions, see Connect to Exchange Online PowerShell.
2. If you enabled the **Encrypt** button in Outlook on the web, disable it by running the Set-IRMConfiguration cmdlet with the SimplifiedClientAccessEnabled parameter. Otherwise, skip this step.

   ```
   Set-IRMConfiguration -SimplifiedClientAccessEnabled $false
   ```

3. Disable the new capabilities for OME by running the Set-IRMConfiguration cmdlet with the AzureRMSLicensingEnabled parameter set to false:

   ```
   Set-IRMConfiguration -AzureRMSLicensingEnabled $false
   ```

Advanced Message Encryption

4/21/2020 • 2 minutes to read • Edit Online

Office 365 Advanced Message Encryption is included in Microsoft 365 Enterprise E5, Office 365 E5, Microsoft 365 E5 (Nonprofit Staff Pricing), Office 365 Enterprise E5 (Nonprofit Staff Pricing), and Office 365 Education A5. If your organization has a subscription that does not include Office 365 Advanced Message Encryption, you can purchase it with the Microsoft 365 E5 Compliance SKU add-on for Microsoft 365 E3, Microsoft 365 E3 (Nonprofit Staff Pricing), or the Office 365 Advanced Compliance SKU add-on for Microsoft 365 E3, Microsoft 365 E3 (Nonprofit Staff Pricing), Office 365 SKUs, or the Microsoft 365 E5/A5 Information Protection and Governance SKU add-on for Microsoft 365 A3/E3.

Advanced Message Encryption helps customers meet compliance obligations that require more flexible controls over external recipients and their access to encrypted emails. With Advanced Message Encryption in Office 365, you can control sensitive emails shared outside the organization with automatic policies. You configure these policies to identify sensitive information types such as PII, Financial, or Health IDs, or you can use keywords to enhance protection. Once you've configured the policies, you pair policies with custom branded email templates and then add an expiration date for extra control of emails that fit the policy. Also, admins can further control encrypted emails accessed externally through a secure web portal by revoking access to the mail at any time.

You can only revoke and set an expiration date for emails sent to external recipients.

Get started with Office 365 Advanced Message Encryption

The following articles describe how you set up and use Advanced Message Encryption.

Your organization must have a subscription that includes Office 365 Advanced Message Encryption. For detailed information about supported subscriptions, see the Message policy and compliance service description.

If you do not have Office 365 Message Encryption set up already, see Set up new Office 365 Message Encryption capabilities.

With Advanced Message Encryption you're not limited to a single branding template. Instead, you can create and use multiple branding templates. For information, see Add your organization's brand to your encrypted messages.

Set an expiration date for email encrypted by Office 365 Advanced Message Encryption. Control sensitive emails shared outside the organization with automatic policies that enhance protection by expiring access through a secure web portal to encrypted emails.

Revoke email encrypted by Office 365 Advanced Message Encryption. Control sensitive emails shared outside the organization and enhance protection by revoking access through a secure web portal to encrypted emails.

With Office 365 Advanced Message Encryption, anytime you apply a custom branding template, Microsoft applies a wrapper to email that fits the mail flow rule to which you apply the template. You can only revoke messages and apply expiration dates to messages that users receive through the portal. In other words, email that has a custom branding template applied. For more information and an example, see the guidance in Ensure all external recipients use the OME Portal to read encrypted mail.

Set an expiration date for email encrypted by Office 365 Advanced Message Encryption

11/2/2020 • 2 minutes to read • Edit Online

Office 365 Advanced Message Encryption is included in Microsoft 365 Enterprise E5, Office 365 E5, Microsoft 365 E5 (Nonprofit Staff Pricing), Office 365 Enterprise E5 (Nonprofit Staff Pricing), and Office 365 Education A5. If your organization has a subscription that does not include Office 365 Advanced Message Encryption, you can purchase it with the Microsoft 365 E5 Compliance SKU add-on for Microsoft 365 E3, Microsoft 365 E3 (Nonprofit Staff Pricing), or the Office 365 Advanced Compliance SKU add-on for Microsoft 365 E3, Microsoft 365 E3 (Nonprofit Staff Pricing), or Office 365 SKUs.

You can use message expiration on emails that your users send to external recipients who use the OME Portal to access encrypted emails. You force recipients to use the OME portal to view and reply to encrypted emails sent by your organization by using a custom branded template that specifies an expiration date in Windows PowerShell.

As an Office 365 global administrator, when you apply your company brand to customize the look of your organization's email messages, you can also specify an expiration for these email messages. With Office 365 Advanced Message Encryption, you can create multiple templates for encrypted emails that originate from your organization. Using a template, you can control how long recipients have access to mail sent by your users.

When an end user receives mail that has an expiration date set, the user sees the expiration date in the wrapper email. If a user tries to open an expired mail, an error appears in the OME portal.

You can only set expiration dates for emails to external recipients.

With Office 365 Advanced Message Encryption, anytime you apply custom branding, the Office 365 applies the wrapper to email that fits the mail flow rule to which you apply the template. In addition, you can only use expiration if you use custom branding.

Create a custom branding template to force mail expiration by using PowerShell

1. Connect to Exchange Online PowerShell with an account that has global administrator permissions in your organization.
2. Run the New-OMEConfiguration cmdlet.

```
New-OMEConfiguration -Identity "Expire in 7 days" -ExternalMailExpiryInDays 7
```

Where:

- `Identity` is the name of the custom template.
- `ExternalMailExpiryInDays` identifies the number of days that recipients can keep mail before it expires. You can use any value between 1–730 days.

More information about Office 365 Advanced Message Encryption

- Office 365 Advanced Message Encryption

- Revoke email encrypted by Office 365 Advanced Message Encryption
- Message policy and compliance service description

Revoke email encrypted by Advanced Message Encryption

2/18/2021 • 5 minutes to read • Edit Online

Email revocation is offered as part of Office 365 Advanced Message Encryption. Office 365 Advanced Message Encryption is included in Microsoft 365 Enterprise E5, Office 365 E5, Microsoft 365 E5 (Nonprofit Staff Pricing), Office 365 Enterprise E5 (Nonprofit Staff Pricing), and Office 365 Education A5. If your organization has a subscription that does not include Office 365 Advanced Message Encryption, you can purchase it with the Microsoft 365 E5 Compliance SKU add-on for Microsoft 365 E3, Microsoft 365 E3 (Nonprofit Staff Pricing), or the Office 365 Advanced Compliance SKU add-on for Microsoft 365 E3, Microsoft 365 E3 (Nonprofit Staff Pricing), or Office 365 SKUs.

This article is part of a larger series of articles about Office 365 Message Encryption.

If a message was encrypted using Office 365 Advanced Message Encryption, and you are a Microsoft 365 admin or you are the sender of the message, you can revoke the message under certain conditions. Admins revoke messages using PowerShell. As a sender, you revoke a message that you sent directly from Outlook on the web. This article describes the circumstances under which revocation is possible and how to do it.

Encrypted emails that you can revoke

Admins and message senders can revoke encrypted emails if the recipient received a link-based, branded encrypted email. If the recipient received a native inline experience in a supported Outlook client, then you can't revoke the message.

Whether a recipient receives a link-based experience or an inline experience depends on the recipient identity type: Office 365 and Microsoft account recipients (for example, outlook.com users) get an inline experience in supported Outlook clients. All other recipient types, such as Gmail and Yahoo recipients, get a link-based experience.

Admins and message senders can revoke messages that are encrypted using encryption applied directly from Outlook on the web. For example, messages encrypted with the Encrypt Only option.

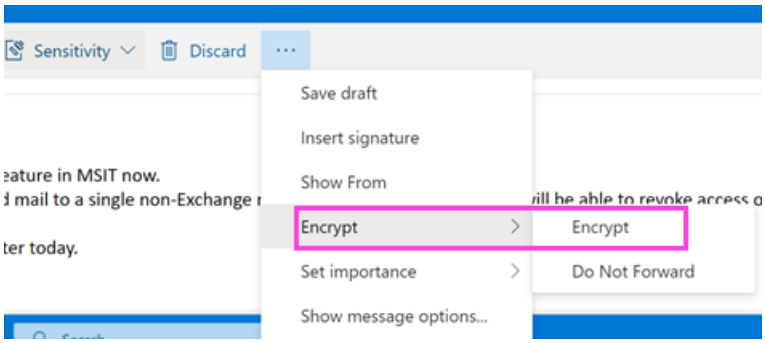

Recipient experience for revoked encrypted emails

Once an email has been revoked, the recipient receives an error when they access the encrypted email through the Office 365 Message Encryption portal: "The message has been revoked by the sender".

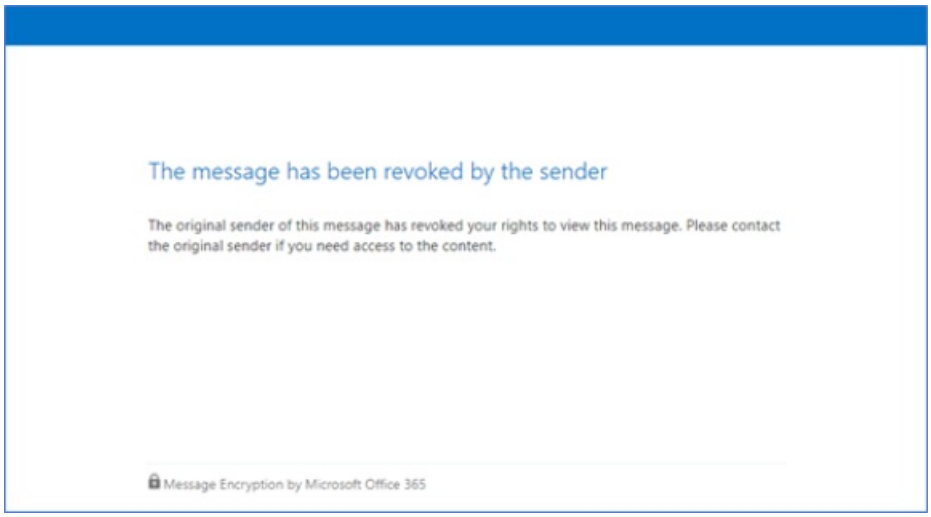

How to revoke an encrypted message that you sent

You can revoke a mail that you sent to a single recipient that uses a social account such as gmail.com or yahoo.com. In other words, you can revoke an email sent to a single recipient that received the link-based experience.

You cannot revoke a mail that you sent to a recipient that uses a work or school account from Office 365 or Microsoft 365 or a user that uses a Microsoft account, for example, an outlook.com account.

To revoke an encrypted message that you sent, complete these steps

1. In Outlook on the web, in your **Sent** folder, browse to the message you want to revoke.

   If the mail is revocable, you'll see the "Remove external access" link at the top of the message.

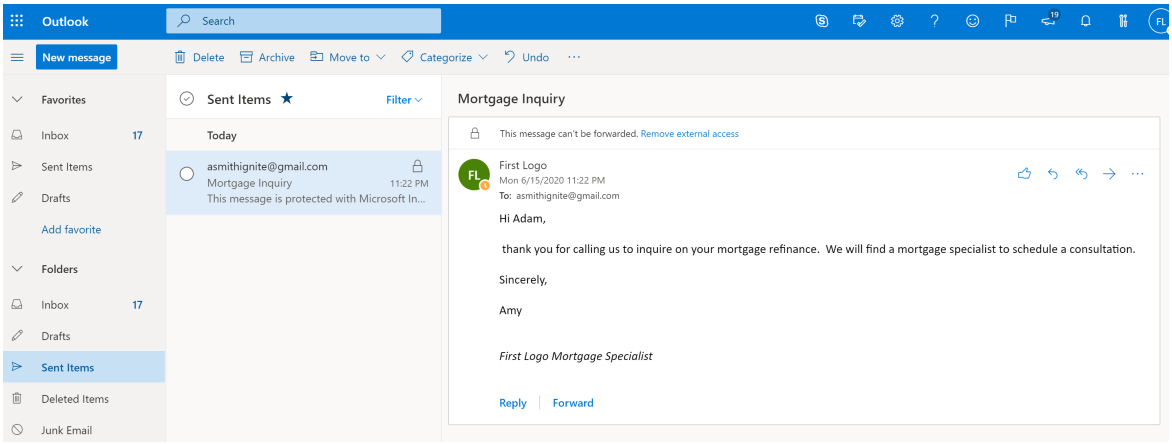

2. Click **Remove external access** to revoke the message.

   The message shows that its status is revoked.

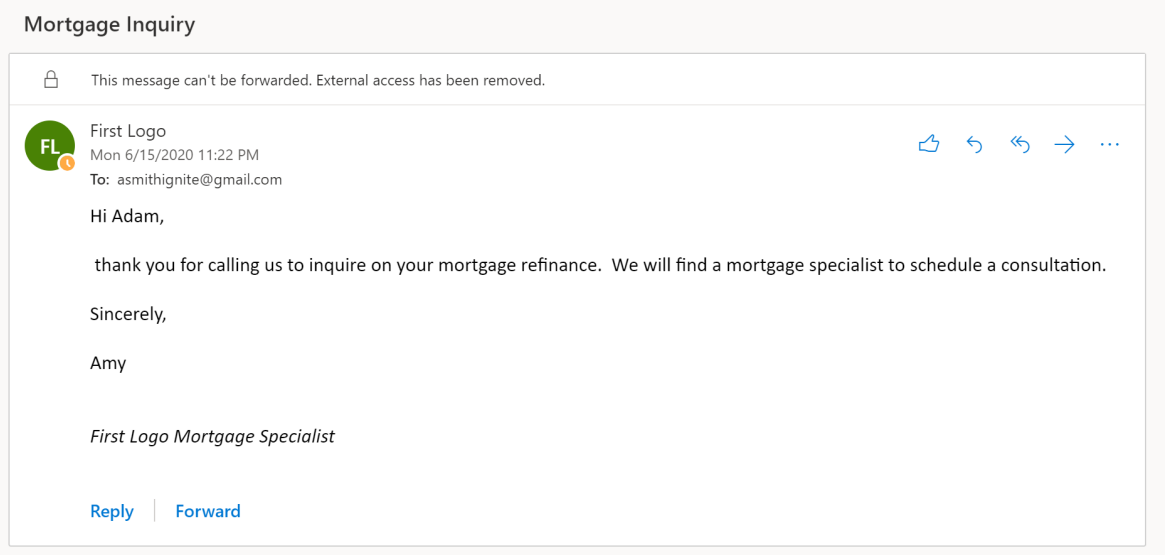

How to revoke an encrypted message as an administrator

Microsoft 365 administrators follow these general steps to revoke an eligible encrypted email:

- Get the Message ID of the email.
- Verify that you can revoke the message.
- Revoke the mail.

Step 1. Obtain the Message ID of the email

Before you can revoke an encrypted mail, gather the Message ID of the mail. The MessageId is usually of the format:

```
<xxxxxxxxxxxxxxxxxxxxxxx@xxxxxx.xxxx.prod.outlook.com>
```

There are multiple ways to find the Message ID of the email that you want to revoke. This section describes a couple of options, but you can use any method that provides the ID.

To identify the Message ID of the email you want to revoke by using Message Trace in the Security & Compliance Center

1. Search for the email by sender or recipient using New Message Trace in Security & Compliance Center.
2. Once you've located the email, select it to bring up the **Message trace details** pane. Expand **More Information** to locate the Message ID.

To identify the Message ID of the email you want to revoke by using Office Message Encryption reports in the Security & Compliance Center

1. In the Security & Compliance Center, navigate to the **Message encryption report**. For information on this report, see View email security reports in the Security & Compliance Center.
2. Choose the **View details** table and identify the message that you want to revoke.
3. Double-click the message to view details that include the Message ID.

Step 2. Verify that the mail is revocable

To verify whether you can revoke a message, check whether the Revocation Status field is visible in the Encryption report, in the **Details** table in the Security & Compliance Center.

To verify whether you can revoke a particular email message by using Windows PowerShell, complete these steps.

1. Using a work or school account that has global administrator permissions in your organization, start a Windows PowerShell session and connect to Exchange Online. For instructions, see Connect to Exchange Online PowerShell.

2. Run the Get-OMEMessageStatus cmdlet as follows:

```
Get-OMEMessageStatus -MessageId "<message id>" | ft -a  Subject, IsRevocable
```

   This command returns the subject of the message and whether the message is revocable. For example,

```
Subject        IsRevocable
-------        -----------
"Test message"True
```

Step 3. Revoke the mail

Once you know the Message ID of the email you want to revoke, and you have verified that the message is revocable, you can revoke the email using the Security & Compliance Center or Windows PowerShell.

To revoke the message using the Security & Compliance Center

1. Using a work or school account that has global administrator permissions in your organization, connect to the Security & Compliance Center.
2. In the **Encryption report**, in the **Details** table for the message, choose **Revoke message**.

To revoke an email by using Windows PowerShell, use the Set-OMEMessageRevocation cmdlet.

1. Using a work or school account that has global administrator permissions in your organization, Connect to Exchange Online PowerShell.
2. Run the Set-OMEMessageRevocation cmdlet as follows:

```
Set-OMEMessageRevocation -Revoke $true -MessageId "<messageId>"
```

3. To check whether the email was revoked, run the Get-OMEMessageStatus cmdlet as follows:

```
Get-OMEMessageStatus -MessageId "<messageId>" | ft -a  Subject, Revoked
```

   If revocation was successful, the cmdlet returns the following result:

```
Revoked: True
```

More information about Office 365 Advanced Message Encryption

- Office 365 Advanced Message Encryption
- Office 365 Advanced Message Encryption - email expiration
- Message policy and compliance service description

Compare versions of OME

2/18/2021 • 7 minutes to read • Edit Online

> **IMPORTANT**
> On February 28, 2021, Microsoft will deprecate support for AD RMS in Exchange Online. If you've deployed a hybrid environment where your Exchange mailboxes are online and you're using IRM with Active Directory RMS on-premises, you'll need to migrate to Azure. Organizations that have deployed into the GCC Moderate environment are also affected. See "Overview of AD RMS deprecation in Exchange Online" in this article for information.

The rest of this article compares legacy Office 365 Message Encryption (OME) to the new OME capabilities and Office 365 Advanced Message Encryption. The new capabilities are a merger and newer version of both OME and Information Rights Management (IRM). Unique characteristics of deploying into GCC High are also outlined. The two can coexist in your organization. For information on how the new capabilities work, see Office 365 Message Encryption (OME).

This article is part of a larger series of articles about Office 365 Message Encryption. This article is intended for administrators and ITPros. If you're just looking for information on sending or receiving an encrypted message, see the list of articles in Office 365 Message Encryption (OME) and locate the article that best fits your needs.

Overview of AD RMS deprecation in Exchange Online

Exchange Online includes Information Rights Management (IRM) functionality that provides online and offline protection of email messages and attachments. By default, Exchange Online uses Azure Azure Information Protection. However, your organization may have configured Exchange Online IRM to use on-premises Active Directory Rights Management Service (AD RMS). AD RMS support in Exchange Online is retiring. Instead, Azure Information Protection will replace AD RMS entirely.

Before you begin, review and assess the impact for your organization. If your organization is already using Azure Information Protection to encrypt email in Exchange Online, there is nothing for you to do. If you encrypt your email using Exchange mail flow rules, for example using Office 365 Message Encryption, you won't have to change your secure email. Otherwise, you'll need to prepare for AD RMS deprecation by switching to Azure Information Protection.

Prepare for AD RMS deprecation

If you've already set up Azure Information Protection but you're not using it, enable the service using Exchange Online PowerShell. On your local computer, using a work or school account that has global administrator permissions in your organization, connect to Exchange Online PowerShell in a Windows PowerShell window.

To enable Azure Information Protection, use the Set-IrmConfiguration cmdlet by typing the following command.

```
Set-IrmConfiguration -AzureRMSLicensingEnabled $true
```

If your organization has not yet set up Azure Information Protection, you'll need to migrate from AD RMS to Azure Information Protection. For instructions, see Migrating from AD RMS to Azure Information Protection.

Side-by-side comparison of features and capabilities

SITUATION	LEGACY OME	IRM IN AD RMS	NEW OME CAPABILITIES
Sending an encrypted mail	Through Exchange mail flow rules	End-user initiated from Outlook desktop or Outlook on the Web; or through Exchange mail flow rules	End-user initiated from Outlook desktop, Outlook for Mac, or Outlook on the Web; through Exchange mail flow rules (also known as transport rules) and Data Loss Prevention (DLP)
Rights management template	N/A	Do Not Forward option and custom templates	Do Not Forward option, Encrypt-Only option, and custom templates
Recipient type	Internal and external recipients	Internal recipients only	Internal and external recipients
Experience for internal recipient	Recipients receive an HTML message, which they download and open in a web browser or mobile app	Native inline experience in Outlook clients	Native inline experience for recipients in the same organization using Outlook clients. Recipients can read message from OME portal using clients other than Outlook (no download or app required).
Experience for external recipient	Recipients receive an HTML message, which they download and open in a web browser or mobile app	N/A	Native inline experience for Microsoft 365 recipients. All other recipients can read message from OME portal (no download or app required).
Attachment permissions	No restrictions on attachments	Attachments are protected	Attachments are protected for the Do Not Forward option and custom templates. Admins can choose whether attachments for the Encrypt-Only option are protected or not.
Bring your own key (BYOK) support	None	None	BYOK supported

Advantages of the new OME capabilities over legacy OME

The new capabilities provide the following advantages:

- Ability to use Encrypt-Only (which enables secure collaboration), Do Not Forward, and custom restrictions.
- Senders can send mail encrypted with the new capabilities manually from Outlook Desktop, Outlook for Mac and Outlook on the web clients.
- Microsoft 365 recipients get to use an inline experience in supported Outlook clients. Alternatively, admins can choose to show Microsoft 365 recipients a branded experience.
- Accounts outside of Microsoft 365, such as Gmail, Yahoo, and Microsoft accounts, are federated with the OME portal, which provides a better user experience for these recipients. All other identities use a one-time

pass code to access encrypted messages.

- Admins can customize branding, and create multiple branding templates.
- Admins can revoke emails encrypted with the new capabilities.
- The new capabilities provide detailed usage reports through the Security & Compliance Center.

Office 365 Advanced Message Encryption capabilities

Office 365 Advanced Message Encryption offers additional capabilities on top of the new OME capabilities. You must have the new Office 365 Message Encryption capabilities set up in your organization in order to use the Advanced Message Encryption capabilities. Also, in order to use these capabilities, recipients must view and reply to secure mail through the OME Portal. The advanced capabilities include:

- Message revocation
- Message expiration
- Multiple branding templates

Office 365 Advanced Message Encryption is not supported in GCC High.

For information on using Advanced Message Encryption, see Office 365 Advanced Message Encryption.

Unique characteristics of Office 365 Message Encryption in a GCC High deployment

If you plan to use Office 365 Message Encryption in a GCC High environment, there are some unique characteristics regarding the recipient experience.

Encrypted email between GCC High and GCC High recipients

Senders can manually encrypt emails in Outlook for PC and Mac and Outlook on the web, or organizations can set up a policy to encrypt emails using Exchange mail flow rules.

Recipients inside GCC High receive the same inline reading experience in Outlook for PC and Mac and Outlook on the web as all other users.

Encrypted email between GCC High and Non-GCC High recipients

Senders inside GCC High can send encrypted email outside of the GCC High boundary and vice versa.

All recipients outside GCC High, including commercial Microsoft 365 users, Outlook.com users, and other users of other email providers such as Gmail and Yahoo, receive a wrapper mail. This wrapper mail redirects the recipient to the OME Portal where the recipient can read and reply to the message. This is also true for senders outside GCC High sending OME encrypted mail to GCC High.

Coexistence of legacy OME and the new capabilities in the same tenant

You can use both legacy OME and the new capabilities in the same tenant. As an administrator, you do this by choosing which version of OME you want to use when you create your mail flow rules.

- To specify the legacy version of OME, use the Exchange mail flow rule action **Apply the previous version of OME**.
- To specify the new capabilities, use the Exchange mail flow rule action **Apply Office 365 Message Encryption and rights protection**.

Users can manually send mail that is encrypted with the new capabilities from Outlook Desktop, Outlook for

Mac, and Outlook on the web.

Migrate from legacy OME to the new capabilities

Even though both versions of OME can coexist, we highly recommend that you edit your old mail flow rules that use the rule action **Apply the previous version of OME** to use the new capabilities. Update these rules to use the mail flow rule action **Apply Office 365 Message Encryption and rights protection**. For instructions, see Define mail flow rules to encrypt email messages in Office 365.

Get started with OME

Typically, the new OME capabilities are automatically enabled for your organization. For more information about the new OME capabilities within your organization, see Set up new Office 365 Message Encryption capabilities.

The legacy version of OME is automatically enabled for your organization if you have enabled Azure Information Protection. In the past, legacy OME worked even if Azure Information Protection wasn't enabled. This is no longer the case.

To start using legacy OME, if you have enabled Azure Information Protection, configure mail flow rules that use the rule action **Apply the previous version of OME**. For instructions, see Define mail flow rules to encrypt email messages.

Message Encryption FAQ

2/18/2021 • 11 minutes to read • Edit Online

Have a question about how the new message protection capabilities work? Check for an answer here. Also, take a look at Frequently asked questions about data protection in Azure Information Protection for answers to questions about the data protection service, Azure Rights Management, in Azure Information Protection.

What is Office 365 Message Encryption (OME)?

OME combines email encryption and rights management capabilities. Rights management capabilities are powered by Azure Information Protection.

Who can use OME?

You can use the new capabilities for OME under the following conditions:

- If you have never set up OME or IRM for Exchange Online in Office 365.
- If you have set up OME and IRM, you can use these steps if you are using the Azure Rights Management service from Azure Information Protection.
- If you are using Exchange Online with Active Directory Rights Management service (AD RMS), you can't enable these new capabilities right away. Instead, you need to migrate AD RMS to Azure Information Protection first. When you've finished the migration, you can successfully set up OME.

  If you choose to continue to use on-premises AD RMS with Exchange Online instead of migrating to Azure Information Protection, you will not be able to use these new capabilities.

What subscriptions do I need to use the new OME capabilities?

To use the new OME capabilities, you need one of the following plans:

- Office 365 Message Encryption is offered as part of Office 365 Enterprise E3 and E5, Microsoft Enterprise E3 and E5, Microsoft 365 Business Premium, Office 365 A1, A3, and A5, and Office 365 Government G3 and G5. Customers do not need additional licenses to receive the new protection capabilities powered by Azure Information Protection.
- You can also add Azure Information Protection Plan 1 to the following plans to receive the new Office 365 Message Encryption capabilities: Exchange Online Plan 1, Exchange Online Plan 2, Office 365 F1, Microsoft 365 Business Basic, Microsoft 365 Business Standard, or Office 365 Enterprise E1.
- Each user benefiting from Office 365 Message Encryption needs to be licensed to be covered by the feature.
- For the full list see the Exchange Online service descriptions for Office 365 Message Encryption.

Can I use Exchange Online with bring your own key (BYOK) in Azure Information Protection?

Yes! Microsoft recommends that you complete the steps to set up BYOK before you set up OME.

For more information about BYOK, see Planning and implementing your Azure Information Protection tenant key.

Do OME and BYOK with Azure Information Protection change Microsoft's approach to third-party data requests such as subpoenas?

No. OME and the option to provide and control your own encryption keys, called BYOK, from Azure Information Protection were not designed to respond to law enforcement subpoenas. OME, with BYOK for Azure Information Protection, was designed for compliance-focused customers. Microsoft takes third-party requests for customer data very seriously. As a cloud service provider, we always advocate for the privacy of customer data. In the event we get a subpoena, we always attempt to redirect the third party to the customer to obtain the information. (Please read Brad Smith's blog: Protecting customer data from government snooping). We periodically publish detailed information of the request we receive. For more information regarding third-party data requests, see Responding to government and law enforcement requests to access customer data on the Microsoft Trust Center. Also, see "Disclosure of Customer Data" in the Online Services Terms (OST).

How is this feature related to legacy Office 365 Message Encryption (OME) and Information Rights Management (IRM) features?

The new capabilities for Office 365 Message Encryption are an evolution of the existing IRM and legacy OME solutions. The following table provides more details.

Comparison of legacy OME, IRM, and new OME capabilities

CAPABILITY	PREVIOUS VERSIONS OF OME	IRM	NEW OME CAPABILITIES
Sending an encrypted email	Only through Exchange mail flow rules	End-user initiated from Outlook for Windows, Outlook for Mac, or Outlook on the web; or through Exchange mail flow rules	End-user initiated from Outlook for Windows, Outlook for Mac, or Outlook on the web; or through mail flow rules
Rights management	-	Do Not Forward option and custom templates	Do Not Forward option, encrypt-only option, default and custom templates
Supported recipient type	External recipients only	Internal recipients only	Internal and external recipients
Experience for recipient	External recipients received an HTML message that they downloaded and opened in a browser or downloaded mobile app.	Internal recipients only received encrypted email in Outlook for Windows, Outlook for Mac, and Outlook on the web.	Internal and external recipients receive email in Outlook for Windows, Outlook for Mac, Outlook on the web, Outlook for Android, and Outlook for iOS, or through a web portal, regardless of whether or not they are in the same organization or in any organization. The OME portal requires no separate download.
Bring Your Own Key support	Not available	Not available	BYOK supported

How do I enable the new OME capabilities for my organization?

See Set up new Office 365 Message Encryption capabilities.

Will the previous version of OME be deprecated?

You can still use the previous version of OME, it will not be deprecated at this time. However, we highly encourage organizations to use the new and improved OME solution. Customers that have not already deployed OME cannot set up a new deployment of the previous version of OME.

My organization uses Active Directory Rights Management, can I use this functionality?

No. If you are using Exchange Online with Active Directory Rights Management service (AD RMS), you can't enable these new capabilities right away. Instead, you need to migrate AD RMS to Azure Information Protection first.

My organization has an Exchange Hybrid deployment. Can I use this feature?

On-premises users can send encrypted mail using Exchange Online mail flow rules. In order to do this, you need to route email through Exchange Online. For more information, see Part 2: Configure mail to flow from your email server to Microsoft 365.

What email client do I need to use in order to create an OME encrypted message? What applications are supported for sending protected messages?

You can create protected messages from Outlook 2016, Outlook 2013 for Windows and Mac, and from Outlook on the web. For more information on sending encrypted messages, see Send, view, and reply to encrypted messages in Outlook for PC.

What email clients are supported to read and reply to protected emails?

Microsoft 365 users can read and respond from Outlook for Windows and Mac (2013 and 2016), Outlook on the web, and Outlook mobile (Android and iOS). You can also use the iOS native mail client if your organization allows it. If you are not a Microsoft 365 user, you can read and reply to encrypted messages on the web through your web browser.

What email clients support the encrypt-only protected emails?

Microsoft 365 users can use Outlook for PC versions 2019 and Microsoft 365 to create mail protected with the encrypt-only policy. That means messages that have the new encrypt-only policy applied can be read directly in Outlook on the web, in Outlook for iOS and Android, and now Outlook for PC versions 2019 and Microsoft 365.

Is there a size limit for messages you can send with OME?

Yes. The maximum message size you can send with OME, including attachments, is 25 MB. For more information, see Message limits.

What file types are supported as attachments in protected emails? Do attachments inherit the protection policies associated with protected

emails?

You can attach any file type to a protected mail. With one exception, protection policies are applied only on the file formats mentioned in File types supported by the Azure Information Protection client. OME does not support the 97-2003 versions of the following Office programs: Word (.doc), Excel (.xls), and PowerPoint (.ppt).

If a file format is supported, such as a Word, Excel, or PowerPoint file, the file is always protected, even after the attachment has been downloaded by the recipient. For example, say an attachment is protected by Do Not Forward. The original recipient downloads the file, creates a message to a new recipient and attaches the file. When the new recipient receives the file, the recipient will not be able to open the protected file.

Are PDF file attachments supported?

The short answer is yes! PDF encryption allows you to protect sensitive PDF documents through secure communication or secure collaboration. When you send email, the Office 365 service encrypts PDF file attachments not the Outlook client.

For Outlook on the web, Outlook for iOS, and Outlook for Android, you can encrypt PDFs you send without any additional steps. These clients natively support PDF encryption.

Outlook desktop does not natively support encryption of PDF file attachments. Instead, you'll need to set up Exchange mail flow rules or DLP to apply encryption to PDF attachments first. When you send mail from Outlook Desktop with a PDF attachment, the client sends the message with the attachment to the service first. When the service receives the file, the service applies the OME protection of the data loss prevention (DLP) policy or mail flow rule in Exchange Online. Next, Exchange Online sends the message with the protected PDF file attachment.

To enable encryption for PDF attachments, run the following command in Exchange Online PowerShell:

```
Set-IRMConfiguration -EnablePdfEncryption $true
```

PDF encryption allows you to protect sensitive PDF documents through secure communication or secure collaboration. For all Outlook clients, messages and unprotected PDF attachments inherit the OME protection of the data loss prevention (DLP) policy or mail flow rule in Exchange Online. Also, if an Outlook on the web user attaches an unprotected PDF document and applies protection to message, the message inherits the protection of the message. Users can only open the encrypted attachments in applications that support protected PDFs (for example, the OME Portal and the Azure Information Protection Viewer).

> **IMPORTANT**
> Outlook desktop client does not support PDF encryption.

Are OneDrive for Business attachments supported?

Not yet. OneDrive for Business attachments are not supported and end-users can't encrypt a mail that contains a cloud OneDrive for Business attachment.

What email clients support preview of encrypted attachments in protected emails?

When attachments are protected with a protected mail, Outlook clients provide tha ability to preview the document directly. Outlook supports preview of Office documents (docx, xlsx, pptx, doc, xls, ppt). Outlook on the web supports preview of Office documents (docx, xlsx, pptx) and PDF.

What email clients support revocation of protected emails?

Outlook on the web supports revocation of protected mail. See How to revoke an encrypted message that you sent for details.

Can I automatically encrypt messages by setting up policies?

Yes. Use mail flow rules in Exchange Online to automatically encrypt a message based on certain conditions. For example, you can create policies that are based on recipient ID, recipient domain, or on the content in the body or subject of the message. See Define mail flow rules to encrypt email messages in Office 365.

Can I automatically remove encryption on incoming and outgoing mail?

Admins can set up a mail flow rule to remove encryption for outgoing mail. You can't set up a rule to remove encryption for incoming mail.

Can I automatically encrypt messages by setting up policies in Data Loss Prevention (DLP) through the Security & Compliance Center?

Yes! You can set up mail flow rules in Exchange Online or by using DLP in the Security & Compliance Center.

Can I customize encrypted messages with my company branding?

Yes! For information on customizing email messages and the OME portal, see Add your organization's brand to your encrypted messages. See Add your organization's brand to your encrypted messages.

Are there any reporting capabilities or insights for encrypted emails?

There is an Encryption report in the Security and Compliance Center. See View email security reports in the Security & Compliance Center.

Can I use message encryption with compliance features such as eDiscovery?

Yes. All encrypted email messages are discoverable by Microsoft 365 compliance features.

Can I remove encryption from email?

Admins can set up a mail flow rule to remove encryption from outgoing mail. You can't remove encryption using a mail flow rule from incoming messages.

Is delegated access supported?

Not at this time.

Can I send as a shared mailbox and encrypt emails?

When someone sends an email message that matches an encryption mail flow rule, the message is encrypted before it's sent.

Can I open encrypted messages sent to a shared mailbox?

Yes! Encrypted messages are supported for a shared mailbox.

- Users can open protected mails in a shared mailbox where the shared mailbox received a protected mail as part of a distribution group.
- Users can view attachments that inherit protection from email when they use Outlook for Windows, Outlook for Mac, and Outlook on the web.

The following table lists the supported clients for shared mailboxes.

PLATFORM	READ MAIL	VIEW EMAIL ATTACHMENTS
Outlook on the web	Yes	Yes
Outlook for Windows	Yes	Yes
Outlook for Mac	Yes	Yes
Outlook for Android	Yes	No
Outlook for iOS	Yes	No

There are currently two known limitations:

- You can't open attachments to emails that you receive on mobile devices by using Outlook mobile.
- We don't support assignment through an email enabled security group. We only support access provided by direct user assignment to the shared mailbox and that automapping is enabled for Exchange Online. Automapping is enabled by default for Exchange Online.

To assign a user to the shared mailbox

1. Connect to Exchange Online Using Remote PowerShell.
2. Run the Add-MailboxPermission cmdlet with the Automapping parameter. This example gives Ayla full access permissions to a support mailbox.

```
Add-MailboxPermission -Identity support@contoso.onmicrosoft.com -User ayla@contoso.com -AccessRights
FullAccess -AutoMapping $true
```

Can I open encrypted messages sent to another user's mailbox with Fullaccess?

Users can open encrypted messages as long as they are given direct access and automapping is turned ON. Access is not allowed if the access is granted via an email-enabled security group.

What do I do if I don’t receive the one-time pass code after I requested it?

First, check the junk or spam folder in your email client. DKIM and DMARC settings for your organization may cause these emails to end up filtered as spam.

Next, check quarantine in the Security & Compliance Center. Often, messages containing a one-time pass code, especially the first ones your organization receives, end up in quarantine.

How Exchange Online secures your email secrets

4/21/2020 • 2 minutes to read • Edit Online

This article describes how Microsoft secures your email secrets in its datacenters.

How do we secure secret information provided by you?

In addition to the Office 365 Trust Center which provides Security, Privacy and Compliance Information for Office 365, you might want to know how Microsoft helps protects secrets you provide in its datacenters. We use a technology called Distributed Key Manager (DKM).

Distributed Key Manager (DKM) is a client-side functionality that uses a set of secret keys to encrypt and decrypt information. Only members of a specific security group in Active Directory Domain Services can access those keys in order 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 datacenter, 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.

For debugging, troubleshooting, or auditing purposes, a datacenter administrator must request elevated access to gain temporary credentials that are part of the security group. This process requires multiple levels of legal approval. If access is granted, all activity is logged and audited. In addition access is only granted for a set interval of time after which it automatically expires.

For extra protection, DKM technology includes automated key rollover and archiving. This also ensures that you can continue to access your older content without having to rely on the same key indefinitely.

Where does Exchange Online make use of DKM?

Microsoft uses Distributed Key Manager to encrypt your secrets in Exchange Online datacenters. For example:

- Email account credentials for connected accounts. Connected accounts are third-party accounts such as Hotmail, Gmail, and Yahoo! mail accounts.
- Customer key. If you are using Service encryption with Customer Key, you'll use Azure Key Vault to safeguard your secrets.

Related topics

Encryption in Office 365

Technical reference details about encryption

Service assurance in the Security & Compliance Center

How Exchange Online uses TLS to secure email connections

2/18/2021 • 6 minutes to read • Edit Online

Learn how Exchange Online and Microsoft 365 use Transport Layer Security (TLS) and Forward Secrecy (FS) to secure email communications. Also provides information about the certificate issued by Microsoft for Exchange Online.

TLS basics for Microsoft 365 and Exchange Online

Transport Layer Security (TLS), and SSL that came before TLS, are cryptographic protocols that secure communication over a network by using security certificates to encrypt a connection between computers. TLS supersedes Secure Sockets Layer (SSL) and is often referred to as SSL 3.1. For Exchange Online, we use TLS to encrypt the connections between our Exchange servers and the connections between our Exchange servers and other servers such as your on-premises Exchange servers or your recipients' mail servers. Once the connection is encrypted, all data sent through that connection is sent through the encrypted channel. However, if you forward a message that was sent through a TLS-encrypted connection, that message isn't necessarily encrypted. This is because, in simple terms, TLS doesn't encrypt the message, just the connection.

If you want to encrypt the message you need to use an encryption technology that encrypts the message contents, for example, something like Office Message Encryption. See Email encryption in Office 365 and Office 365 Message Encryption (OME) for information on message encryption options in Office 365.

We recommend using TLS in situations where you want to set up a secure channel of correspondence between Microsoft and your on-premises organization or another organization, such as a partner. Exchange Online always attempts to use TLS first to secure your email but cannot always do this if the other party does not offer TLS security. Keep reading to find out how you can secure all mail to your on-premises servers or important partners by using *connectors*.

To provide the best-in-class encryption to our customers, Microsoft has deprecated Transport Layer Security (TLS) versions 1.0 and 1.1 in Office 365 and Office 365 GCC. However, you can continue to use an unencrypted SMPT connection without any TLS. We don't recommend email transmission without any encryption.

How Exchange Online uses TLS between Exchange Online customers

Exchange Online servers always encrypt connections to other Exchange Online servers in our datacenters with TLS 1.2. When you send mail to a recipient that is within your organization, that email is automatically sent over a connection that is encrypted using TLS. Also, all email that you send to other customers is sent over connections that are encrypted using TLS and are secured using Forward Secrecy.

How Microsoft 365 uses TLS between Microsoft 365 and external, trusted partners

By default, Exchange Online always uses opportunistic TLS. This means Exchange Online always tries to encrypt connections with the most secure version of TLS first, then works its way down the list of TLS ciphers until it finds one on which both parties can agree. Unless you have configured Exchange Online to ensure that messages to that recipient are only sent through secure connections, then by default the message will be sent unencrypted if the recipient organization doesn't support TLS encryption. Opportunistic TLS is sufficient for most businesses. However, for business that have compliance requirements such as medical, banking, or

government organizations, you can configure Exchange Online to require, or force, TLS. For instructions, see Configure mail flow using connectors in Office 365.

If you decide to configure TLS between your organization and a trusted partner organization, Exchange Online can use forced TLS to create trusted channels of communication. Forced TLS requires your partner organization to authenticate to Exchange Online with a security certificate in order to send mail to you. Your partner will need to manage their own certificates in order to do this. In Exchange Online, we use connectors to protect messages that you send from unauthorized access before they arrive at the recipient's email provider. For information on using connectors to configure mail flow, see Configure mail flow using connectors in Office 365.

TLS and hybrid Exchange Server deployments

If you are managing a hybrid Exchange deployment, your on-premises Exchange server needs to authenticate to Microsoft 365 using a security certificate in order to send mail to recipients whose mailboxes are only in Office 365. As a result, you need to manage your own security certificates for your on-premises Exchange servers. You must also securely store and maintain these server certificates. For more information about managing certificates in hybrid deployments, see Certificate requirements for hybrid deployments.

How to set up forced TLS for Exchange Online in Office 365

For Exchange Online customers, in order for forced TLS to work to secure all of your sent and received email, you need to set up more than one connector that requires TLS. You'll need one connector for email sent to your user mailboxes and another connector for email sent from your user mailboxes. Create these connectors in the Exchange admin center in Office 365. For instructions, see Configure mail flow using connectors in Office 365.

TLS certificate information for Exchange Online

The certificate information used by Exchange Online is described in the following table. If your business partner is setting up forced TLS on their email server, you will need to provide this information to them. Be aware that for security reasons, our certificates do change from time to time. We have rolled out an update to our certificate within our datacenters. The new certificate is valid from September 3, 2018.

Current certificate information valid from September 3, 2018

ATTRIBUTE	VALUE
Certificate authority root issuer	GlobalSign Root CA – R1
Certificate name	mail.protection.outlook.com
Organization	Microsoft Corporation
Organization unit	
Certificate key strength	2048

Deprecated certificate information valid until September 3, 2018

To help ensure a smooth transition, we will continue to provide the old certificate information for your reference for some time, however, you should use the current certificate information from now on.

ATTRIBUTE	VALUE
Certificate authority root issuer	Baltimore CyberTrust Root
Certificate name	mail.protection.outlook.com
Organization	Microsoft Corporation
Organization unit	Microsoft Corporation
Certificate key strength	2048

Prepare for the new Exchange Online certificate

The new certificate is issued by a different certificate authority (CA) from the previous certificate used by Exchange Online. As a result, you may need to perform some actions in order to use the new certificate.

The new certificate requires connecting to the endpoints of the new CA as part of validating the certificate. Failure to do so can result in mail flow being negatively affected. If you protect your mail servers with firewalls that only let the mail servers connect with certain destinations you need to check if your server is able to validate the new certificate. To confirm that your server can use the new certificate, complete these steps:

1. Connect to your local Exchange Server using Windows PowerShell and then run the following command: `certutil -URL https://crl.globalsign.com/gsorganizationvalsha2g3.crl`
2. On the window that appears, choose **Retrieve**.
3. When the utility completes its check it returns a status. If the status displays **OK**, then your mail server can successfully validate the new certificate. If not, you need to determine what is causing the connections to fail. Most likely, you need to update the settings of a firewall. The full list of endpoints that need to be accessed include:
   - ocsp.globalsign.com
   - crl.globalsign.com
   - secure.globalsign.com

Normally, you receive updates to your root certificates automatically through Windows Update. However some deployments have additional security in place that prevents these updates from occurring automatically. In these locked-down deployments where Windows Update can't automatically update root certificates, you need to ensure that the correct root CA certificate is installed by completing these steps:

1. Connect to your local Exchange Server using Windows PowerShell and then run the following command: `certmgr.msc`
2. Under **Trusted Root Certification Authority/Certificates**, confirm that the new certificate is listed.

Get more information about TLS and Microsoft 365

For a list of supported cipher suites, see Technical reference details about encryption.

Set up connectors for secure mail flow with a partner organization

Connectors with enhanced email security

Encryption in Microsoft 365

BitLocker and Distributed Key Manager (DKM) for Encryption

5/5/2020 • 2 minutes to read • Edit Online

Microsoft servers use BitLocker to encrypt the disk drives containing customer data at rest at the volume-level. BitLocker encryption is a data protection feature that is built into Windows. BitLocker is one of the technologies used to safeguard against threats in case there are lapses in other processes or controls (e.g., access control or recycling of hardware) that could lead to someone gaining physical access to disks containing customer data. In this case, BitLocker eliminates the potential for data theft or exposure because of lost, stolen, or inappropriately decommissioned computers and disks.

BitLocker is deployed with Advanced Encryption Standard (AES) 256-bit encryption on disks containing customer data in Exchange Online, SharePoint Online, and Skype for Business. Disk sectors are encrypted with a Full Volume Encryption Key (FVEK), which is encrypted with the Volume Master Key (VMK), which in turn is bound to the Trusted Platform Module (TPM) in the server. The VMK directly protects the FVEK and therefore, protecting the VMK becomes critical. The following figure illustrates an example of the BitLocker key protection chain for a given server (in this case, using an Exchange Online server).

The following table describes the BitLocker key protection chain for a given server (in this case, an Exchange Online server).

KEY PROTECTOR	GRANULARITY	HOW GENERATED?	WHERE IS IT STORED?	PROTECTION
AES 256-bit External Key	Per Server	BitLocker APIs	TPM or Secret Safe	Lockbox / Access Control
			Mailbox Server Registry	TPM encrypted
48-digit Numerical Password	Per Disk	BitLocker APIs	Active Directory	Lockbox / Access Control
X.509 Certificate as Data Recovery Agent (DRA) also called Public Key Protector	Environment (e.g., Exchange Online multitenant)	Microsoft CA	Build System	No one user has the full password to the private key. The password is under physical protection.

BitLocker key management involves the management of recovery keys that are used to unlock/recover encrypted disks in a Microsoft datacenter. Microsoft 365 stores the master keys in a secured share, only accessible by individuals who have been screened and approved. The credentials for the keys are stored in a secured repository for access control data (what we call a "secret store"), which requires a high level of elevation and management approvals to access using a just-in-time access elevation tool.

BitLocker supports keys which fall into two management categories:

- BitLocker-managed keys, which are generally short-lived and tied to the lifetime of an operating system instance installed on a server or to a given disk. These keys are deleted and reset during server reinstallation or disk formatting.
- BitLocker recovery keys, which are managed outside of BitLocker but used for disk decryption. BitLocker

uses recovery keys for the scenario in which an operating system is reinstalled, and encrypted data disks already exist. Recovery keys are also used by Managed Availability monitoring probes in Exchange Online where a responder may need to unlock a disk.

BitLocker-protected volumes are encrypted with a full volume encryption key, which in turn is encrypted with a volume master key. BitLocker uses FIPS-compliant algorithms to ensure that encryption keys are never stored or sent over the wire in the clear. The Microsoft 365 implementation of customer data-at-rest-protection does not deviate from the default BitLocker implementation.

Legacy information for Office 365 Message Encryption

2/18/2021 • 17 minutes to read • Edit Online

If you haven't yet moved your organization to the new OME capabilities, but you have already deployed OME, then the information in this article applies to your organization. Microsoft recommends that you make a plan to move to the new OME capabilities as soon as it is reasonable for your organization. For instructions, see Set up new Office 365 Message Encryption capabilities built on top of Azure Information Protection. If you want to find out more about how the new capabilities work first, see Office 365 Message Encryption. The rest of this article refers to OME behavior before the release of the new OME capabilities.

With Office 365 Message Encryption, your organization can send and receive encrypted email messages between people inside and outside your organization. Office 365 Message Encryption works with Outlook.com, Yahoo, Gmail, and other email services. Email message encryption helps ensure that only intended recipients can view message content.

Here are some examples:

- A bank employee sends credit card statements to customers
- An insurance company representative provides policy details to customers
- A mortgage broker requests financial information from a customer for a loan application
- A health care provider sends health care information to patients
- An attorney sends confidential information to a customer or another attorney

How Office 365 Message Encryption works without the new capabilities

Office 365 Message Encryption is an online service that's built on Microsoft Azure Rights Management (Azure RMS). With Azure RMS, administrators can define mail flow rules to determine the conditions for encryption. For example, a rule can require the encryption of all messages addressed to a specific recipient.

When someone sends an email message in Exchange Online that matches an encryption rule, the message is sent with an HTML attachment. The recipient opens the HTML attachment and follows instructions to view the encrypted message on the Office 365 Message Encryption portal. The recipient can choose to view the message by signing in with a Microsoft account or a work or school associated with Office 365, or by using a one-time pass code. Both options help ensure that only the intended recipient can view the encrypted message. This process is very different for the new OME capabilities.

The following diagram summarizes the passage of an email message through the encryption and decryption process.

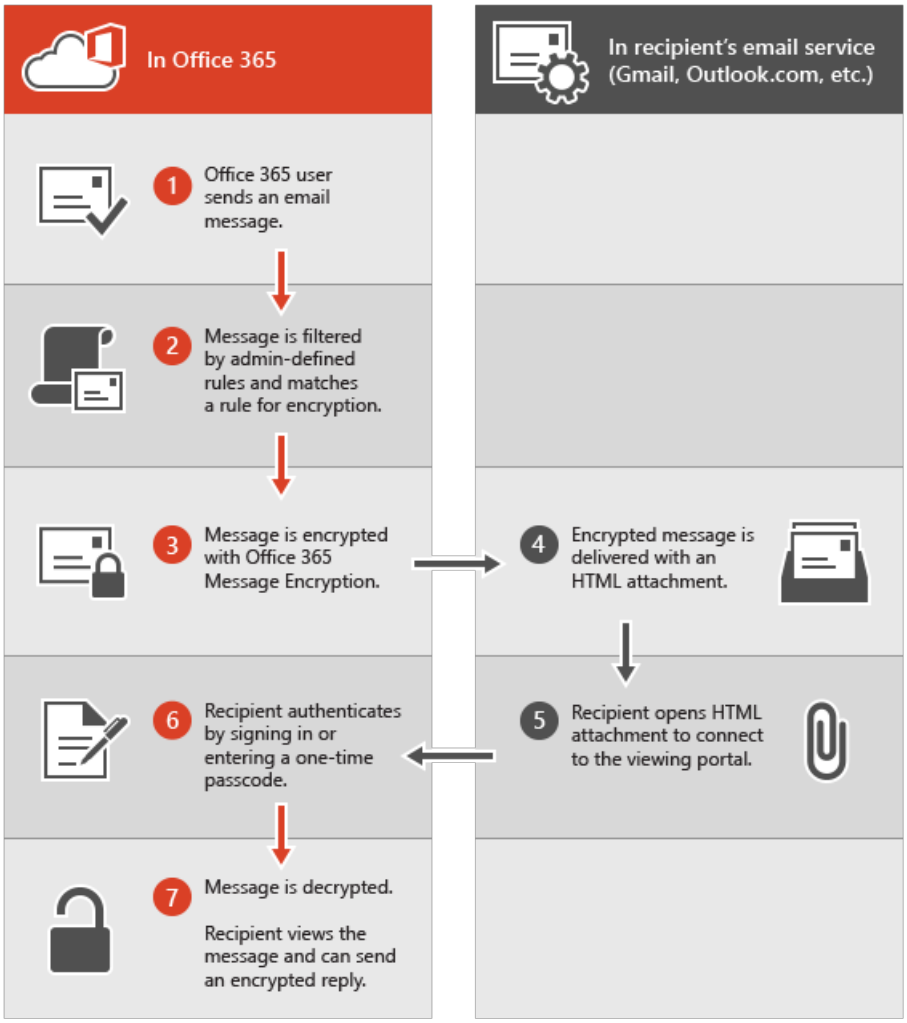

For more information, see Service information for legacy Office 365 Message Encryption prior to the release of the new OME capabilities.

Defining mail flow rules for Office 365 Message Encryption that don't use the new OME capabilities

To enable Office 365 Message Encryption without the new capabilities, Exchange Online and Exchange Online Protection administrators define Exchange mail flow rules. These rules determine under what conditions email messages should be encrypted, as well as conditions for removing message encryption. When an encryption action is set for a rule, the service performs the action on any messages that match the rule conditions before sending the messages.

Mail flow rules are flexible, letting you combine conditions so you can meet specific security requirements in a single rule. For example, you can create a rule to encrypt all messages that contain specified keywords and are addressed to external recipients. Office 365 Message Encryption also encrypts replies from recipients of encrypted email, and you can create a rule that decrypts those replies as a convenience for your email users. That way, users in your organization won't have to sign in to the encryption portal to view replies.

For more information about how to create Exchange mail flow rules, see Define Rules for Office 365 Message Encryption.

Use the EAC to create a mail flow rule for encrypting email messages without the new OME capabilities

1. In a web browser, using a work or school account that has been granted global administrator permissions, sign in to Office 365.
2. Choose the **Admin** tile.

3. In the Microsoft 365 admin center, choose **Admin centers** > **Exchange**.
4. In the EAC, go to **Mail flow** > **Rules** and select **New +** > **Create a new rule**. For more information about using the EAC, see Exchange admin center in Exchange Online.
5. In **Name**, type a name for the rule, such as Encrypt mail for DrToniRamos@hotmail.com.
6. In **Apply this rule if** select a condition, and enter a value if necessary. For example, to encrypt messages going to DrToniRamos@hotmail.com:
   a. In **Apply this rule if**, select **the recipient is**.
   b. Select an existing name from the contact list or type a new email address in the **check names** box.
      - To select an existing name, select it from the list and then click **OK**.
      - To enter a new name, type an email address in the **check names** box and then select **check names** > **OK**.
7. To add more conditions, choose **More options** and then select **add condition** and select from the list.

   For example, to apply the rule only if the recipient is outside your organization, select **add condition** and then select **The recipient is external/internal** > **Outside the organization** > **OK**.
8. To enable encryption without using the new OME capabilities, in **Do the following**, select **Modify the message security** > **Apply the previous version of OME**, and then choose **Save**.

   If you receive an error that IRM licensing isn't enabled, then you're not using legacy OME.
9. (Optional) Choose **add action** to specify another action.

Use Exchange Online PowerShell to create a mail flow rule for encrypting email messages without the new OME capabilities

1. Connect to Exchange Online PowerShell. For more information, see Connect to Exchange Online PowerShell.
2. Create a rule by using the **New-TransportRule** cmdlet and set the *ApplyOME* parameter to `$true`.

   This example requires that all email messages sent to DrToniRamos@hotmail.com must be encrypted.

   ```
   New-TransportRule -Name "Encrypt rule for Dr Toni Ramos" -SentTo "DrToniRamos@hotmail.com" -
   SentToScope "NotinOrganization" -ApplyOME $true
   ```

   Where,

   - The unique name of the new rule is "Encrypt rule for Dr Toni Ramos".
   - The *SentTo* parameter specifies the message recipients (identified by name, email address, distinguished name, etc.). In this example, the recipient is identified by the email address "DrToniRamos@hotmail.com".
   - The *SentToScope* parameter specifies the location of the message recipients. In this example, the recipient's mailbox is in hotmail and is not part of the organization, so the value `NotInOrganization` is used.

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

Remove encryption from email replies encrypted without the new OME capabilities

When your email users send encrypted messages, recipients of those messages can respond with encrypted replies. You can create mail flow rules to automatically remove encryption from replies so email users in your organization don't have to sign in to the encryption portal to view them. You can use the EAC or Windows PowerShell cmdlets to define these rules. You can decrypt messages that are sent from within your organization

or messages that are replies to messages sent from within your organization. You cannot decrypt encrypted messages originating from outside of your organization.

Use the EAC to create a rule for removing encryption from email replies encrypted without the new OME capabilities

1. In a web browser, using a work or school account that has been granted admin permissions, sign in to Office 365.
2. Choose the **Admin** tile.
3. In the Microsoft 365 admin center, choose **Admin centers** > **Exchange**.
4. In the EAC, go to **Mail flow** > **Rules** and select **New +** > **Create a new rule**. For more information about using the EAC, see Exchange admin center in Exchange Online.
5. In **Name**, type a name for the rule, such as Remove encryption from incoming mail.
6. In **Apply this rule if** select the conditions where encryption should be removed from messages, such as **The recipient is located** > **Inside the organization**.
7. In **Do the following**, select **Modify the message security** > **Remove the previous version of OME**.
8. Select **Save**.

Use Exchange Online PowerShell to create a rule to remove encryption from email replies encrypted without the new OME capabilities

1. Connect to Exchange Online PowerShell. For more information, see Connect to Exchange Online PowerShell.
2. Create a rule by using the **New-TransportRule** cmdlet and set the *RemoveOME* parameter to `$true`.

   This example removes the encryption from all mail sent to recipients in the organization.

   ```
   New-TransportRule -Name "Remove encryption from incoming mail" -SentToScope "InOrganization" -
   RemoveOME $true
   ```

   Where,

   - The unique name of the new rule is "Remove encryption from incoming mail".
   - The *SentToScope* parameter specifies the location of the message recipients. In this example, the value `InOrganization` value is used, which indicates one of the following:
     - The recipient is a mailbox, mail user, group, or mail-enabled public folder in your organization.
     - The recipient's email address is in an accepted domain that's configured as an authoritative domain or an internal relay domain in your organization, *and* the message was sent or received over an authenticated connection.

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

Sending, viewing, and replying to messages encrypted without the new capabilities

With Office 365 Message Encryption, email messages are encrypted automatically, based on administrator-defined rules. An email that bears an encrypted message arrives in the recipient's Inbox with an attached HTML file.

Recipients follow instructions in the message to open the attachment and authenticate by using a Microsoft account or a work or school associated with Office 365. If recipients don't have either account, they're directed to create a Microsoft account that will let them sign in to view the encrypted message. Alternatively, recipients can

choose to get a one-time pass code to view the message. After signing in or using a one-time pass code, recipients can view the decrypted message and send an encrypted reply.

Customize encrypted messages with Office 365 Message Encryption

As an Exchange Online and Exchange Online Protection administrator, you can customize your encrypted messages. For example, you can add your company's brand and logo, specify an introduction, and add disclaimer text in encrypted messages and in the portal where recipients view your encrypted messages. Using Windows PowerShell cmdlets, you can customize the following aspects of the viewing experience for recipients of encrypted email messages:

- Introductory text of the email that contains the encrypted message
- Disclaimer text of the email that contains the encrypted message
- Portal text that will appear in the message viewing portal
- Logo that will appear in the email message and viewing portal

You can also revert back to the default look and feel at any time.

The following example shows a custom logo for ContosoPharma in the email attachment:

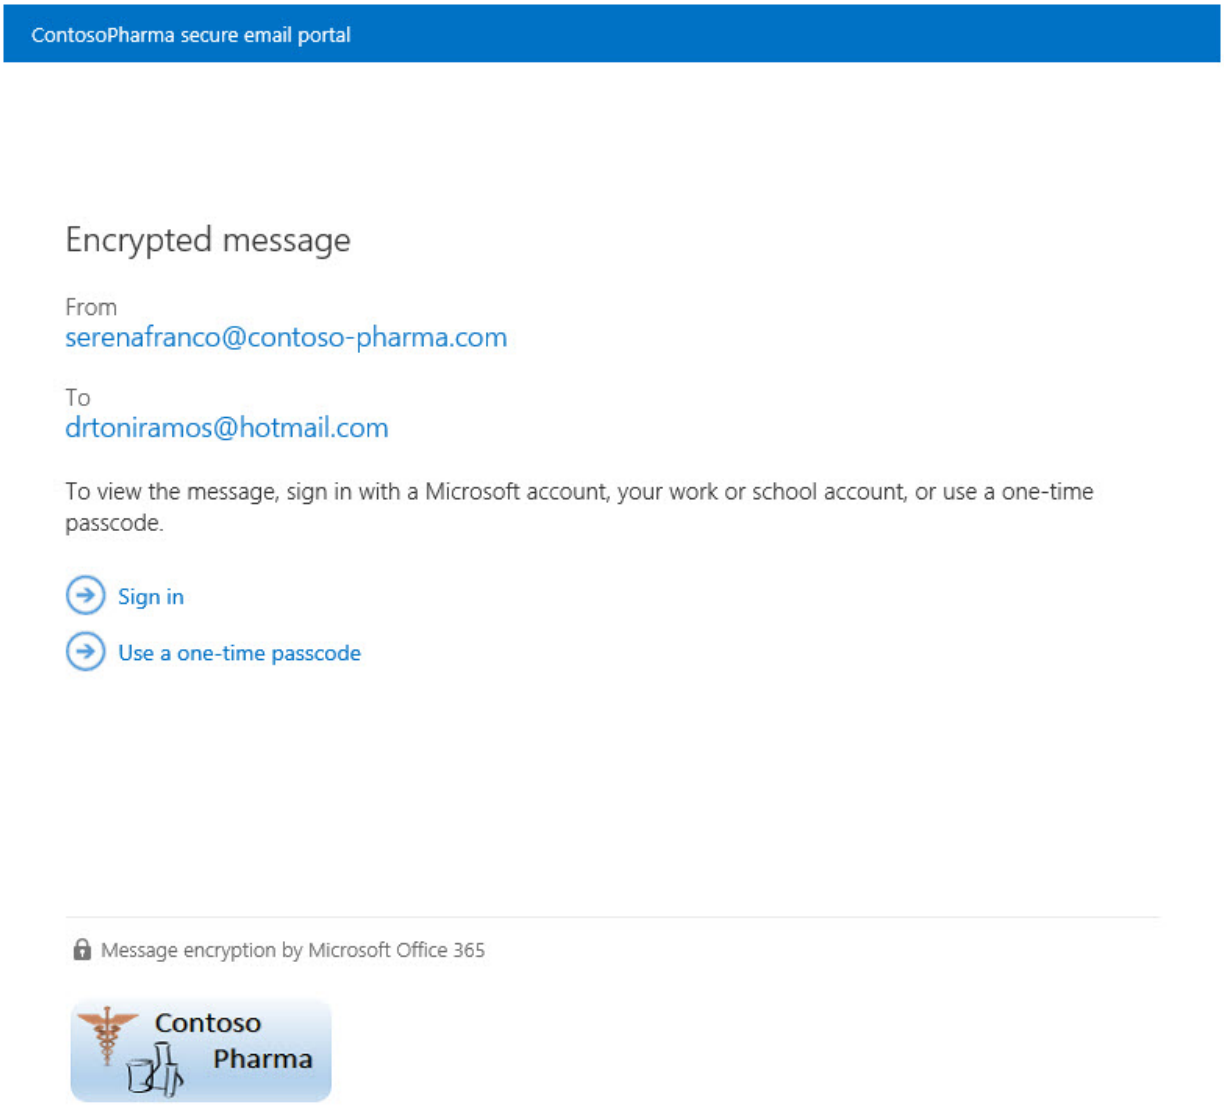

To customize encryption email messages and the encryption portal with your organization's brand

1. Connect to Exchange Online using Remote PowerShell, as described in Connect to Exchange Online Using

Remote PowerShell.

2. Use the Set-OMEConfiguration cmdlet as described here: Set-OMEConfiguration or use the following table for guidance.

Encryption customization options

TO CUSTOMIZE THIS FEATURE OF THE ENCRYPTION EXPERIENCE	USE THESE WINDOWS POWERSHELL COMMANDS

Default text that accompanies encrypted email messages The default text appears above the instructions for viewing encrypted messages	`Set-OMEConfiguration -Identity <OMEConfigurationIdParameter> -EmailText " <string of up to 1024 characters>"`
	Example:
	`Set-OMEConfiguration -Identity "OME Configuration" -EmailText "Encrypted message from ContosoPharma secure messaging system"`

Disclaimer statement in the email that contains the encrypted message	`Set-OMEConfiguration -Identity <OMEConfigurationIdParameter> DisclaimerText " <your disclaimer statement, string of up to 1024 characters>"`
	Example:
	`Set-OMEConfiguration -Identity "OME Configuration" -DisclaimerText "This message is confidential for the use of the addressee only"`

Text that appears at the top of the encrypted mail viewing portal	`Set-OMEConfiguration -Identity <OMEConfigurationIdParameter> -PortalText "<text for your portal, string of up to 128 characters>"`
	Example:
	`Set-OMEConfiguration -Identity "OME Configuration" -PortalText "ContosoPharma secure email portal"`

Logo	`Set-OMEConfiguration -Identity <OMEConfigurationIdParameter> -Image <Byte[]>`
	Example:
	`Set-OMEConfiguration -Identity "OME configuration" -Image (Get-Content "C:\Temp\contosologo.png" -Encoding byte)`
	Supported file formats: .png, .jpg, .bmp, or .tiff
	Optimal size of logo file: less than 40 KB
	Optimal size of logo image: 170x70 pixels

To remove brand customizations from encryption email messages and the encryption portal

1. Connect to Exchange Online using Remote PowerShell, as described in Connect to Exchange Online Using Remote PowerShell.

2. Use the Set-OMEConfiguration cmdlet as described here: Set-OMEConfiguration. To remove your organization's branded customizations from the DisclaimerText, EmailText, and PortalText values, set the value to an empty string, `""`. For all image values, such as Logo, set the value to `"$null"`.

Encryption customization options

TO REVERT THIS FEATURE OF THE ENCRYPTION EXPERIENCE BACK TO THE DEFAULT TEXT AND IMAGE	USE THESE WINDOWS POWERSHELL COMMANDS

TO REVERT THIS FEATURE OF THE ENCRYPTION EXPERIENCE BACK TO THE DEFAULT TEXT AND IMAGE	USE THESE WINDOWS POWERSHELL COMMANDS

Default text that accompanies encrypted email messages The default text appears above the instructions for viewing encrypted messages	`Set-OMEConfiguration -Identity <OMEConfigurationIdParameter> -EmailText "<empty string>"`
	Example:
	`Set-OMEConfiguration -Identity "OME Configuration" -EmailText ""`

Disclaimer statement in the email that contains the encrypted message	`Set-OMEConfiguration -Identity <OMEConfigurationIdParameter> DisclaimerText "<empty string>"`
	Example:
	`Set-OMEConfiguration -Identity "OME Configuration" -DisclaimerText ""`

Text that appears at the top of the encrypted mail viewing portal	`Set-OMEConfiguration -Identity <OMEConfigurationIdParameter> -PortalText "<empty string>"`
	Example reverting back to default:
	`Set-OMEConfiguration -Identity "OME Configuration" -PortalText ""`

Logo	`Set-OMEConfiguration -Identity <OMEConfigurationIdParameter> -Image <"$null">`
	Example reverting back to default:
	`Set-OMEConfiguration -Identity "OME configuration" -Image $null`

Service information for legacy Office 365 Message Encryption prior to the release of the new OME capabilities

The following table provides technical details for the Office 365 Message Encryption service prior to the release of the new OME capabilities.

SERVICE DETAILS	DESCRIPTION
Client device requirements	Encrypted messages can be viewed on any client device, as long as the HTML attachment can be opened in a modern browser that supports Form Post.
Encryption algorithm and Federal Information Processing Standards (FIPS) compliance	Office 365 Message Encryption uses the same encryption keys as Windows Azure Information Rights Management (IRM) and supports Cryptographic Mode 2 (2K key for RSA and 256 bits key for SHA-1 systems). For more information about the underlying IRM cryptographic modes, see AD RMS Cryptographic Modes.
Supported message types	Office 365 Message Encryption is only supported for items that have a message class ID of **IPM.Note**. For more information, see Item types and message classes.
Message size limits	Office 365 Message Encryption can encrypt messages of up to 25 megabytes. For more details about message size limits, see Exchange Online Limits.
Exchange Online email retention policies	Exchange Online doesn't store the encrypted messages.

SERVICE DETAILS	DESCRIPTION
Language support for Office 365 Message Encryption	Office 365 Message encryption supports Microsoft 365 languages, as follows: Incoming email messages and attached HTML files are localized based on the sender's language settings. The viewing portal is localized based on the recipient's browser settings. The body (content) of the encrypted message isn't localized.
Privacy information for OME Portal and OME Viewer App	The Office 365 Messaging Encryption Portal privacy statement provides detailed information about what Microsoft does and doesn't do with your private information.

Frequently Asked Questions about legacy OME

Got questions about Office 365 Message Encryption? Here are some answers. If you can't find what you need, check the Microsoft Tech Community forums for Office 365.

Q. My users send encrypted email messages to recipients outside our organization. Is there anything that external recipients have to do in order to read and reply to email messages that are encrypted with Office 365 Message Encryption?

Recipients outside your organization who receive Microsoft 365 encrypted messages can view them in one of two ways:

- By signing in with a Microsoft account or a work or school account associated with Office 365.
- By using a one-time pass code.

Q. Are Microsoft 365 encrypted messages stored in the cloud or on Microsoft servers?

No, the encrypted messages are kept on the recipient's email system, and when the recipient opens the message, it is temporarily posted for viewing on Microsoft servers. The messages are not stored there.

Q. Can I customize encrypted email messages with my brand?

Yes. You can use Windows PowerShell cmdlets to customize the default text that appears at the top of encrypted email messages, the disclaimer text, and the logo that you want to use for the email message and the encryption portal. This feature is now available in OMEv2. For details, see Add branding to encrypted messages.

Q. Does the service require a license for every user in my organization?

A license is required for every user in the organization who sends encrypted email.

Q. Do external recipients require subscriptions?

No, external recipients do not require a subscription to read or reply to encrypted messages.

Q. How is Office 365 Message Encryption different from Rights Management Services (RMS)?

RMS provides Information Rights Protection capabilities for an organization's internal emails by providing built-in templates, such as: Do not forward and Company Confidential. Office 365 Message Encryption supports email message encryption for messages that are sent to external recipients as well as internal recipients.

Q. How is Office 365 Message Encryption different from S/MIME?

S/MIME is essentially a client-side encryption technology, and requires complicated certificate management and publishing infrastructure. Office 365 Message Encryption uses mail flow rules (also known as transport rules) and does not depend on certificate publishing.

Q. Can I read the encrypted messages over mobile devices?

Yes, you can view messages on Android and iOS by downloading the OME Viewer apps from the Google Play store and the Apple App store. Open the HTML attachment in the OME Viewer app and then follow the instructions to open your encrypted message. For other mobile devices, you can open the HTML attachment as long as your mail client supports Form Post.

Q. Are replies and forwarded messages encrypted?

Yes. Responses continue to be encrypted throughout the duration of the thread.

Q. Does Office 365 Message Encryption provide localization?

Incoming email and HTML content is localized based on sender email settings. The viewing portal is localized based on recipient's browser settings. However, the actual body (content) of encrypted message isn't localized.

Q. What encryption method is used for Office 365 Message Encryption?

Office 365 Message Encryption uses Rights Management Services (RMS) as its encryption infrastructure. The encryption method used depends on where you obtain the RMS keys used to encrypt and decrypt messages.

- If you use Microsoft Azure RMS to obtain the keys, Cryptographic Mode 2 is used. Cryptographic Mode 2 is an updated and enhanced AD RMS cryptographic implementation. It supports RSA 2048 for signature and encryption, and supports SHA-256 for signature.
- If you use Active Directory (AD) RMS to obtain the keys, either Cryptographic Mode 1 or Cryptographic Mode 2 is used. The method used depends on your on-premises AD RMS deployment. Cryptographic Mode 1 is the original AD RMS cryptographic implementation. It supports RSA 1024 for signature and encryption, and supports SHA-1 for signature. This mode continues to be supported by all current versions of RMS.

For more information, see AD RMS Cryptographic Modes.

Q. Why do some encrypted messages say they come from Office365@messaging.microsoft.com?

When an encrypted reply is sent from the encryption portal or through the OME Viewer app, the sending email address is set to Office365@messaging.microsoft.com because the encrypted message is sent through a Microsoft endpoint. This helps to prevent encrypted messages from being marked as spam. The displayed name on the email and the address within the encryption portal aren't changed because of this labeling. Also, this labeling only applies to messages sent through the portal, not through any other email client.

Q. I am an Exchange Hosted Encryption (EHE) subscriber. Where can I learn more about the upgrade to Office 365 Message Encryption?

All EHE customers have been upgraded to Office 365 Message Encryption. For more information, visit the Exchange Hosted Encryption Upgrade Center.

Q. Do I need to open any URLs, IP addresses, or ports in my organization's firewall to support Office 365 Message Encryption?

Yes. You have to add URLs for Exchange Online to the allow list for your organization to enable authentication for messages encrypted by Office 365 Message Encryption. For a list of Exchange Online URLs, see Microsoft 365 URLs and IP address ranges.

Q. How many recipients can I send a Microsoft 365 encrypted message to?

The recipient limit is 500 recipients per message, or, when combined after distribution list expansion, 11,980 characters in the message's **To** field, whichever comes first.

Q. Is it possible to revoke a message sent to a particular recipient?

No. You can't revoke a message to a particular person after it's sent.

Q. Can I view a report of encrypted messages that have been received and read?

There isn't a report that shows if an encrypted message has been viewed, but there are Microsoft 365 reports available that you can leverage to determine the number of messages that matched a specific mail flow rule (also known as a transport rule), for instance.

Q. What does Microsoft do with the information I provide through the OME Portal and the OME Viewer App?

The Office 365 Messaging Encryption Portal privacy statement provides detailed information about what Microsoft does and doesn't do with your private information.

Q. What do I do if I don’t receive the one-time pass code after I requested it?

First, check the junk or spam folder in your email client. DKIM and DMARC settings for your organization may cause these emails to end up filtered as spam.

Next, check quarantine in the Security & Compliance Center. Often, messages containing a one-time pass code, especially the first ones your organization receives, end up in quarantine.

Set up Azure Rights Management for the previous version of Message Encryption

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

This topic describes the steps you need to follow in order to activate and then set up Azure Rights Management (RMS), part of Azure Information Protection, for use with the previous version of Office 365 Message Encryption (OME).

This article only applies to the previous version of OME

If you haven't yet moved your organization to the new OME capabilities, but you have already deployed OME, then the information in this article applies to your organization. Microsoft recommends that you make a plan to move to the new OME capabilities as soon as it is reasonable for your organization. For instructions, see Set up new Office 365 Message Encryption capabilities. If you want to find out more about how the new capabilities work first, see Office 365 Message Encryption. The rest of this article refers to OME behavior before the release of the new OME capabilities.

Prerequisites for using the previous version of Office 365 Message Encryption

Office 365 Message Encryption (OME), including IRM, depends on Azure Rights Management (Azure RMS). Azure RMS is the protection technology used by Azure Information Protection. To use OME, your organization must include an Exchange Online or Exchange Online Protection subscription that, in turn, includes an Azure Rights Management subscription.

- If you're not sure of what your subscription includes, see the Exchange Online service descriptions for Message Policy, Recovery, and Compliance.
- If you have Azure Rights Management but it's not set up for Exchange Online or Exchange Online Protection, this article explains how to activate Azure Rights Management and then the describes the best way to set up OME to work with Azure Rights Management.
- If you've already set up OME to work with Azure Rights Management for Exchange Online or Exchange Online Protection, depending on how you set it up, you may be ready to start using OME and its new capabilities right away. This article explains how to determine if you've set OME up correctly, what to do if you need to change your setup, and what happens if you choose not to change your setup. For example, in order to use the new capabilities, you must use Azure RMS with OME. You can't use the new capabilities with an on-premises Active Directory RMS.

Activate Azure Rights Management for the previous version of OME in Office 365

You need to activate Azure Rights Management so that the users in your organization can apply information protection to messages that they send, and open messages and files that have been protected by the Azure Rights Management service. For instructions, see Activating Azure Rights Management. Once you've completed the activation, return here and continue with the tasks in this article.

Set up the previous version of OME to use Azure RMS by importing

trusted publishing domains (TPDs)

A TPD is an XML file that contains information about your organization's rights management settings. For example, the TPD contains information about the server licensor certificate (SLC) used for signing and encrypting certificates and licenses, the URLs used for licensing and publishing, and so on. You import the TPD into your organization by using Windows PowerShell.

> **IMPORTANT**
> Previously, you could choose to import TPDs from the Active Directory Rights Management service (AD RMS) into your organization. However, doing so will prevent you from using the new OME capabilities and is not recommended. If your organization is currently configured this way, Microsoft recommends that you create a plan to migrate from your on-premises Active Directory RMS to cloud-based Azure Information Protection. For more information, see Migrating from AD RMS to Azure Information Protection. You will not be able to use the new OME capabilities until you have completed the migration to Azure Information Protection.

To import TPDs from Azure RMS

1. Connect to Exchange Online Using Remote PowerShell.
2. Choose the key-sharing URL that corresponds to your organization's geographic location:

LOCATION	KEY SHARING LOCATION URL
North America	https://sp-rms.na.aadrm.com/TenantManagement/ServicePartner.svc
European Union	https://sp-rms.eu.aadrm.com/TenantManagement/ServicePartner.svc
Asia	https://sp-rms.ap.aadrm.com/TenantManagement/ServicePartner.svc
South America	https://sp-rms.sa.aadrm.com/TenantManagement/ServicePartner.svc
Office 365 for Government (Government Community Cloud) This RMS key-sharing location is reserved for customers who have purchased Office 365 for Government SKUs.	https://sp-rms.govus.aadrm.com/TenantManagement/ServicePartner.svc

3. Configure the key-sharing location by running the Set-IRMConfiguration cmdlet as follows:

```
Set-IRMConfiguration -RMSOnlineKeySharingLocation "<RMSKeySharingURL >"
```

   For example, to configure the key sharing location if your organization is located in North America:

```
Set-IRMConfiguration -RMSOnlineKeySharingLocation "https://sp-rms.na.aadrm.com/TenantManagement/ServicePartner.svc"
```

4. Run the Import-RMSTrustedPublishingDomain cmdlet with the -RMSOnline switch to import the TPD from Azure Rights Management:

```
Import-RMSTrustedPublishingDomain -RMSOnline -Name "<TPDName> "
```

Where *TPDName* is the name you want to use for the TPD. For example, "Contoso North American TPD".

5. To verify that you successfully configured your organization to use the Azure Rights Management service, run the Test-IRMConfiguration cmdlet with the -RMSOnline switch as follows:

   ```
   Test-IRMConfiguration -RMSOnline
   ```

   Among other things, this cmdlet checks connectivity with the Azure Rights Management service, downloads the TPD, and checks its validity.

6. Run the Set-IRMConfiguration cmdlet as follows to disable Azure Rights Management templates from being available in Outlook on the web and Outlook:

   ```
   Set-IRMConfiguration -ClientAccessServerEnabled $false
   ```

7. Run the Set-IRMConfiguration cmdlet as follows to enable Azure Rights Management for your cloud-based email organization and configure it to use Azure Rights Management for Office 365 Message Encryption:

   ```
   Set-IRMConfiguration -InternalLicensingEnabled $true
   ```

8. To verify that you have successfully imported the TPD and enabled Azure Rights Management, use the Test-IRMConfiguration cmdlet to test Azure Rights Management functionality. For details, see "Example 1" in Test-IRMConfiguration.

I have the previous version of OME set up with Active Directory Rights Management not Azure Information Protection, what do I do?

You can continue to use your existing Office 365 Message Encryption mail flow rules with Active Directory Rights Management, but you can't configure or use the new OME capabilities. Instead, you need to migrate to Azure Information Protection. For information about migration and what this means for your organization, see Migrating from AD RMS to Azure Information Protection.

Next steps

Once you've completed Azure Rights Management setup, if you want to enable the new OME capabilities, see Set up new Office 365 Message Encryption capabilities built on top of Azure Information Protection.

After you've set up your organization to use the new OME capabilities, you're ready to Define mail flow rules to protect email messages with new OME capabilities.

Related topics

Encryption in Office 365

Technical reference details about encryption in Office 365

What is Azure Rights Management?

Set up Information Rights Management (IRM) in SharePoint admin center

11/2/2020 • 3 minutes to read • Edit Online

Within SharePoint Online, IRM protection is applied to files at the list and library level. Before your organization can use IRM protection, you must first set up Rights Management. IRM relies on the Azure Rights Management service from Azure Information Protection to encrypt and assign usage restrictions. Some Microsoft 365 plans include Azure Rights Management, but not all. To learn more, read How Office applications and services support Azure Rights Management.

Turn on IRM service using SharePoint admin center

Before your organization can IRM-protect SharePoint lists and libraries, you must first activate the Rights Management service for your organization. To learn how see Activating Azure Rights Management. You must use a work or school account that has global administrator privileges to enable the Rights Management service. Otherwise, you won't be able to use IRM features with SharePoint Online.

After activating the Rights Management service, sign in to the SharePoint admin center to turn on IRM.

1. Sign in as a global admin or SharePoint admin.
2. Select the app launcher icon in the upper-left and choose **Admin** to open the Microsoft 365 admin center. (If you don't see the Admin tile, you don't have administrator permissions in your organization.)
3. In the left pane, choose **Admin centers** > **SharePoint**.
4. In the left pane, choose **settings**, and then choose **classic settings page**.
5. In the **Information Rights Management (IRM)** section, choose **Use the IRM service specified in your configuration**, and then choose **Refresh IRM Settings**. After you refresh IRM settings, people in your organization can begin using IRM in their SharePoint lists and document libraries. However, the options to do so may take up to an hour to appear in Library Settings and List Settings.

IRM-enable SharePoint document libraries and lists

After refreshing IRM settings, site owners can IRM-protect their SharePoint lists and document libraries. For more information, see Apply Information Rights Management to a list or library.

When site owners enable IRM for a list or library, they can protect any supported file types in that list or library. When IRM is enabled for a library, rights management applies to all of the files in that library. When you enable IRM for a list, rights management applies only to files that are attached to list items, not the actual list items.

When people download files in an IRM-enabled list or library, the files are encrypted so that only authorized people can view them. Each rights-managed file also contains an issuance license that imposes restrictions on the people who view the file. Typical restrictions include making a file read-only, disabling the copying of text, preventing people from saving a local copy, and preventing people from printing the file. Client programs that can read IRM-supported file types use the issuance license within the rights-managed file to enforce these restrictions. This is how a rights-managed file retains its protection even after it is downloaded. To enable IRM on a list or library, see Apply Information Rights Management to a list or library.

You cannot create or edit documents in an IRM-enabled library using Office in a browser. Instead, one person at a time can download and edit IRM-encrypted files. Use check-in and check-out to manage *co-authoring* , or

authoring across multiple users.

When you download a PDF file from an IRM-protected library, Microsoft 365 creates a protected PDF file. The file's extension won't change, but the file is protected. To view this file you'll need the Azure Information Protection viewer, the full Azure Information Protection client, or another application that supports viewing protected PDF files.

SharePoint Online supports encryption of the following file types:

- PDF
- The 97-2003 file formats for the following Microsoft Office programs: Word, Excel, and PowerPoint
- The Office Open XML formats for the following Microsoft Office programs: Word, Excel, and PowerPoint
- The XML Paper Specification (XPS) format

> **NOTE**
> IRM protection cannot be applied to protected documents (like digitally signed PDF files) as SharePoint needs to open the document on upload.

Next steps

Once you've enabled IRM for SharePoint Online, you can start applying rights management to lists and libraries. For information, see Apply Information Rights Management to a list or library.

The new OneDrive sync client for Windows now supports synchronizing IRM-protected SharePoint document libraries and OneDrive locations (as long as the IRM setting for the library isn't set to expire document access rights). For more information, or to get started deploying the new sync client, see Deploy the new OneDrive sync client for Windows.

Top of page

Technical reference details about encryption

2/18/2021 • 4 minutes to read • Edit Online

Refer to this article to learn about certificates, technologies, and TLS cipher suites used for encryption in Office 365. This article also provides details about planned deprecations.

- If you're looking for overview information, see Encryption in Office 365.
- If you're looking for setup information, see Set up encryption in Office 365 Enterprise.
- For information about cipher suites supported by specific versions of Windows, see Cipher Suites in TLS/SSL (Schannel SSP).

Microsoft Office 365 certificate ownership and management

You don't need to purchase or maintain certificates for Office 365. Instead, Office 365 uses its own certificates.

Current encryption standards and planned deprecations

To provide best-in-class encryption, Office 365 regularly reviews supported encryption standards. Sometimes, old standards are deprecated as they become out of date and less secure. This article describes currently supported cipher suites and other standards and details about planned deprecations.

FIPS compliance for Office 365

All cipher suites supported by Office 365 use algorithms acceptable under FIPS 140-2. Office 365 inherits FIPS validations from Windows (through Schannel). For information about Schannel, see Cipher Suites in TLS/SSL (Schannel SSP).

Versions of TLS supported by Office 365

TLS, and SSL that came before TLS, are cryptographic protocols that secure communication over a network by using security certificates to encrypt a connection between computers. Office 365 supports TLS version 1.2 (TLS 1.2).

TLS version 1.3 (TLS 1.3) is currently not supported.

Support for TLS 1.0 and 1.1 deprecation

Office 365 stopped supporting TLS 1.0 and 1.1 on October 31, 2018. We have completed disabling TLS 1.0 and 1.1 in GCC High and DoD environments. We began disabling TLS 1.0 and 1.1 for Worldwide and GCC environments beginning on October 15, 2020 and will continue with roll-out over the next weeks and months.

To maintain a secure connection to Office 365 and Microsoft 365 services, all client-server and browser-server combinations use TLS 1.2 and modern cipher suites. You might have to update certain client-server and browser-server combinations. For information about how this change impacts you, see Preparing for the mandatory use of TLS 1.2 in Office 365.

Deprecating support for 3DES

Since October 31, 2018, Office 365 no longer supports the use of 3DES cipher suites for communication to Office 365. More specifically, Office 365 no longer supports the TLS_RSA_WITH_3DES_EDE_CBC_SHA cipher suite. Since February 28, 2019, this cipher suite has been disabled in Office 365. Clients and servers that

communicate with Office 365 must support one or more of the supported ciphers. For a list of supported ciphers, see TLS cipher suites supported by Office 365.

Deprecating SHA-1 certificate support in Office 365

Since June 2016, Office 365 no longer accepts an SHA-1 certificate for outbound or inbound connections. Use SHA-2 (Secure Hash Algorithm 2) or a stronger hashing algorithm in the certificate chain.

TLS cipher suites supported by Office 365

TLS uses *cipher suites*, collections of encryption algorithms, to establish secure connections. Office 365 supports the cipher suites listed in the following table. The table lists the cipher suites in order of strength, with the strongest cipher suite listed first.

Office 365 responds to a connection request by first attempting to connect using the most secure cipher suite. If the connection doesn't work, Office 365 tries the second most secure cipher suite in the list, and so on. The service continues down the list until the connection is accepted. Likewise, when Office 365 requests a connection, the receiving service chooses whether TLS will be used and which cipher suite to use.

> **IMPORTANT**
> Be aware that TLS versions deprecate, and that deprecated versions *should not be used* where newer versions are available. TLS 1.3 is currently not supported. If your legacy services do not require TLS 1.0 or 1.1 you should disable them.

CIPHER SUITE	KEY EXCHANGE ALGORITHM/STRENGTH	FORWARD SECRECY	CIPHER/STRENGTH	AUTHENTICATION ALGORITHM
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384	ECDH/192	Yes	AES/256	RSA/112
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256	ECDH/128	Yes	AES/128	RSA/112
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384	ECDH/192	Yes	AES/256	RSA/112
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256	ECDH/128	Yes	AES/128	RSA/112
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA	ECDH/192	Yes	AES/256	RSA/112
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA	ECDH/128	Yes	AES/128	RSA/112
TLS_RSA_WITH_AES_256_GCM_SHA384	RSA/112	No	AES/256	RSA/112

CIPHER SUITE	KEY EXCHANGE ALGORITHM/STRENGTH	FORWARD SECRECY	CIPHER/STRENGTH	AUTHENTICATION ALGORITHM
TLS_RSA_WITH_AES_128_GCM_SHA256	RSA/112	No	AES/256	RSA/112

These cipher suites supported TLS 1.0 and 1.1 protocols until their deprecation date. For GCC High and DoD environments that deprecation date was January 15, 2020, and for Worldwide and GCC environments that date was October 15, 2020.

PROTOCOLS	CIPHER SUITE NAME	KEY EXCHANGE ALGORITHM/STRENGTH	FORWARD SECRECY SUPPORT	AUTHENTICATION ALGORITHM/STRENGTH	CIPHER/STRENGTH
TLS 1.0, 1.1, 1.2	TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA	ECDH/192	Yes	RSA/112	AES/256
TLS 1.0, 1.1, 1.2	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA	ECDH/128	Yes	RSA/112	AES/128
TLS 1.0, 1.1, 1.2	TLS_RSA_WITH_AES_256_CBC_SHA	RSA/112	No	RSA/112	AES/256
TLS 1.0, 1.1, 1.2	TLS_RSA_WITH_AES_128_CBC_SHA	RSA/112	No	RSA/112	AES/128
TLS 1.0, 1.1, 1.2	TLS_RSA_WITH_AES_256_CBC_SHA256	RSA/112	No	RSA/112	AES/256
TLS 1.0, 1.1, 1.2	TLS_RSA_WITH_AES_128_CBC_SHA256	RSA/112	No	RSA/112	AES/256

Related articles

TLS Cipher Suites in Windows 10 v1903

Encryption in Office 365

Set up encryption in Office 365 Enterprise

Schannel implementation of TLS 1.0 in Windows security status update: November 24, 2015

TLS/SSL Cryptographic Enhancements (Windows IT Center)

Preparing for TLS 1.2 in Office 365 and Office 365 GCC

Disabling TLS 1.0 and 1.1 for Microsoft 365

2/18/2021 • 3 minutes to read • Edit Online

> **IMPORTANT**
>
> We temporarily halted disablement of TLS 1.0 and 1.1 for commercial customers due to COVID-19. As supply chains have adjusted and certain countries open back up, we restarted the TLS 1.2 enforcement rollout on October 15, 2020. Rollout will continue over the following weeks and months.

As of October 31, 2018, the Transport Layer Security (TLS) 1.0 and 1.1 protocols are deprecated for the Microsoft 365 service. The effect for end-users is minimal. This change has been publicized for over two years, with the first public announcement made in December 2017. This article is only intended to cover the Office 365 local client in relation to the Office 365 service but can also apply to on-premises TLS issues with Office and Office Online Server/Office Web Apps.

For SharePoint and OneDrive, you'll need to update and configure .NET to support TLS 1.2. For information, see How to enable TLS 1.2 on clients.

Office 365 and TLS overview

The Office client relies on the Windows web service (WINHTTP) to send and receive traffic over TLS protocols. The Office client can use TLS 1.2 if the web service of the local computer can use TLS 1.2. All Office clients can use TLS protocols, as TLS and SSL protocols are part of the operating system and not specific to the Office client.

On Windows 8 and later versions

By default, the TLS 1.2 and 1.1 protocols are available if no network devices are configured to reject TLS 1.2 traffic.

On Windows 7

TLS 1.1 and 1.2 protocols are not available without the KB 3140245 update. The update addresses this issue and adds the following registry sub key:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Internet Settings\WinHttp
```

> **NOTE**
>
> Windows 7 users who do not have this update are affected as of October 31, 2018. KB 3140245 has details about how to change WINHTTP settings to enable TLS protocols.

More information

The value of the **DefaultSecureProtocols** registry key that the KB article describes determines which network protocols can be used:

DEFAULTSECUREPROTOCOLS VALUE	PROTOCOL ENABLED
0x00000008	Enable SSL 2.0 by default
0x00000020	Enable SSL 3.0 by default

DEFAULTSECUREPROTOCOLS VALUE	PROTOCOL ENABLED
0x00000080	Enable TLS 1.0 by default
0x00000200	Enable TLS 1.1 by default
0x00000800	Enable TLS 1.2 by default

Office clients and TLS registry keys

You can refer to KB 4057306 Preparing for the mandatory use of TLS 1.2 in Office 365. This is a general article for IT administrators, and it's official documentation about the TLS 1.2 change.

The following table shows the appropriate registry key values in Office 365 clients after October 31, 2018.

ENABLED PROTOCOLS FOR OFFICE 365 SERVICE AFTER OCTOBER 31, 2018	HEXADECIMAL VALUE
TLS 1.0 + 1.1 + 1.2	0x00000A80
TLS 1.1 + 1.2	0x00000A00
TLS 1.0 + 1.2	0x00000880
TLS 1.2	0x00000800

IMPORTANT

Don't use the SSL 2.0 and 3.0 protocols, which can also be set by using the **DefaultSecureProtocols** key. SSL 2.0 and 3.0 are considered outdated and insecure protocols. The best practice is to end the use of SSL 2.0 and SSL 3.0, although the decision to do this ultimately depends on what best meets your product needs. For more information about SSL 3.0 vulnerabilities, refer to KB 3009008.

You can use the default Windows Calculator in Programmer mode to set up the same reference registry key values. For more information, see KB 3140245 Update to enable TLS 1.1 and TLS 1.2 as a default secure protocols in WinHTTP in Windows.

Regardless if the Windows 7 update (KB 3140245) is installed or not, the DefaultSecureProtocols registry sub key isn't present and must be added manually or through a group policy object (GPO). That is, unless you have to customize what secure protocols are enabled or restricted, this key is not required. You only need the Windows 7 SP1 (KB 3140245) update.

Update and configure the .NET Framework to support TLS 1.2

You'll need to update applications that call Microsoft 365 APIs over TLS 1.0 or TLS 1.1 to use TLS 1.2. .NET 4.5 defaults to TLS 1.1. To update your .NET configuration, see How to enable Transport Layer Security (TLS) 1.2 on clients.

More information

For more information, see Preparing for the mandatory use of TLS 1.2 in Office 365.

Disabling TLS 1.0 and 1.1 in Office 365 GCC High and DoD

2/18/2021 • 2 minutes to read • Edit Online

Summary

In order to comply with the latest compliance standards for the Federal Risk and Authorization Management Program (FedRAMP), we are disabling Transport Layer Security (TLS) versions 1.1 and 1.0 in Microsoft 365 for GCC High and DoD environments. This change was previously announced through Microsoft Support in Preparing for the mandatory use of TLS 1.2 in Office 365.

The security of your data is important, and we are committed to transparency about changes that could affect your use of the service.

Although the Microsoft TLS 1.0 implementation has no known security vulnerabilities, we remain committed to the FedRAMP compliance standards. Therefore, we disabled TLS 1.1 and 1.0 in Office 365 in GCC High and DoD environments on January 15, 2020. For information about how to remove TLS 1.1 and 1.0 dependencies, see the following white paper:

Solving the TLS 1.0 problem

More information

Starting on January 15, 2020, Office 365 in the GCC High and DoD environments will deprecate TLS 1.1 and 1.0.

By January 15, 2020, all combinations of client servers and browser servers should use TLS version 1.2 (or a later version) to make sure that all connections can be made without issues to Office 365 services. This may require updates to certain combinations of client servers and browser servers.

If you do not update to TLS version 1.2 (or a later version) by January 15, 2020, you will experience issues when you try to connect to Office 365. Additionally, you will be required to update to TLS 1.2 (or a later version) as part of the resolution.

You must update your client computers to make sure that you maintain uninterrupted access to Office 365 GCC High and DoD.

You'll need to update applications that call Microsoft 365 APIs over TLS 1.0 or TLS 1.1 to use TLS 1.2. .NET 4.5 defaults to TLS 1.1. To update your .NET configuration, see How to enable Transport Layer Security (TLS) 1.2 on clients. For more information, see Preparing for the mandatory use of TLS 1.2 in Office 365.

We know that the following client applications cannot use TLS 1.2:

- Android 4.3 and earlier versions
- Firefox version 5.0 and earlier versions
- Internet Explorer 8–10 on Windows 7 and earlier versions
- Internet Explorer 10 on Windows Phone 8.0
- Safari 6.0.4/OS X 10.8.4 and earlier versions

Although current analysis of connections to Microsoft Online services shows that most services and endpoints see very little TLS 1.1 and 1.0 usage, we're providing notice of this change so that you can update any affected clients or servers as necessary before support for TLS 1.1 and 1.0 ends. If you are using any on-premises infrastructure for hybrid scenarios or Active Directory Federation Services (AD FS), make sure that the

infrastructure can support both inbound and outbound connections that use TLS 1.2 (or a later version).

In addition to the outages that you might experience if you use the listed clients that cannot use TLS 1.2, removing TLS 1.1 and 1.0 will prevent you from being able to use the following Microsoft product:

- Lync phone

References

The following support article describes guidance and references to help make sure that your clients are using TLS 1.2:

Preparing for the mandatory use of TLS 1.2 in Office 365

Preparing for TLS 1.2 in Office 365 and Office 365 GCC

2/18/2021 • 3 minutes to read • Edit Online

Summary

To provide the best-in-class encryption to our customers, Microsoft plans to deprecate Transport Layer Security (TLS) versions 1.0 and 1.1 in Office 365 and Office 365 GCC. We understand that the security of your data is important, and we're committed to transparency about changes that may affect your use of the TLS service.

The Microsoft TLS 1.0 implementation has no known security vulnerabilities. But because of the potential for future protocol downgrade attacks and other TLS vulnerabilities, we are discontinuing support for TLS 1.0 and 1.1 in Microsoft Office 365 and Office 365 GCC.

For information about how to remove TLS 1.0 and 1.1 dependencies, see the following white paper: Solving the TLS 1.0 problem.

More information

We have already begun deprecation of TLS 1.0 and 1.1 as of January 2020. Any clients, devices, or services that connect to Office 365 through TLS 1.0 or 1.1 in our DoD or GCC High instances are unsupported. For our commercial customers of Office 365, deprecation of TLS 1.0 and 1.1 will begin October 15, 2020 and rollout will continue over the following weeks and months.

We recommend that all client-server and browser-server combinations use TLS 1.2 (or a later version) in order to maintain connection to Office 365 services. You might have to update certain client-server and browser-server combinations.

You'll need to update applications that call Microsoft 365 APIs over TLS 1.0 or TLS 1.1 to use TLS 1.2. .NET 4.5 defaults to TLS 1.1. To update your .NET configuration, see How to enable Transport Layer Security (TLS) 1.2 on clients.

The following clients are known to be unable to use TLS 1.2. Update these clients to ensure uninterrupted access to the service.

- Android 4.3 and earlier versions
- Firefox version 5.0 and earlier versions
- Internet Explorer 8-10 on Windows 7 and earlier versions
- Internet Explorer 10 on Windows Phone 8
- Safari 6.0.4/OS X10.8.4 and earlier versions

TLS 1.2 for Microsoft Teams Rooms and Surface Hub

Microsoft Teams Rooms (previously known as Skype Room System V2 SRS V2) have supported TLS 1.2 since December 2018. We recommend that Rooms devices have Microsoft Teams Rooms app version 4.0.64.0 or later installed. For more information, see the Release notes. The changes are backward and forward compatible.

Surface Hub released TLS 1.2 support in May 2019.

TLS 1.2 support for Microsoft Teams Rooms and Surface Hub products also requires the following server-side code changes:

- Skype for Business Online server changes were made live in April 2019. Now, Skype for Business Online

supports connecting Microsoft Teams Rooms and Surface Hub devices by using TLS 1.2.

- Skype for Business Server customers must install a cumulative update (CU) to use TLS 1.2 for Teams Rooms Systems and Surface Hub.
  - For Skype for Business Server 2015, CU9 is already released in May 2019.
  - For Skype for Business Server 2019, CU1 was previously planned for April 2019 but is delayed to June 2019.

> **NOTE**
> Skype for Business on-premises customers should not disable TLS 1.0/1.1 before installing specific CUs for Skype for Business Server.

If you are using any on-premises infrastructure for hybrid scenarios or Active Directory Federation Services, make sure that the infrastructure can support both inbound and outbound connections that use TLS 1.2.

References

The following resources provide guidance to help make sure that your clients are using TLS 1.2 or a later version and to disable TLS 1.0 and 1.1.

- For Windows 7 clients that connect to Office 365, make sure that TLS 1.2 is the default secure protocol in WinHTTP in Windows. For more information see KB 3140245 - Update to enable TLS 1.1 and TLS 1.2 as a default secure protocols in WinHTTP in Windows.
- To start addressing weak TLS use by removing TLS 1.0 and 1.1 dependencies, see TLS 1.2 support at Microsoft.
- New IIS functionality makes it easier to find clients on Windows Server 2012 R2 and Windows Server 2016 that connect to the service by using weak security protocols.
- Get more information about how to solve the TLS 1.0 problem.
- For general information about our approach to security, go to the Office 365 Trust Center.
- Preparing for TLS 1.0/1.1 Deprecation - Office 365 Skype for Business
- Exchange Server TLS guidance, part 1: Getting Ready for TLS 1.2
- Exchange Server TLS guidance Part 2: Enabling TLS 1.2 and Identifying Clients Not Using It
- Exchange Server TLS guidance Part 3: Turning Off TLS 1.0/1.1
- Enable TLS 1.1 and TLS 1.2 support in Office Online Server

Overview of data loss prevention

2/18/2021 • 29 minutes to read • Edit Online

NOTE

Data loss prevention capabilities were recently added to Microsoft Teams chat and channel messages for users licensed for Office 365 Advanced Compliance, which is available as a standalone option and is included in Office 365 E5 and Microsoft 365 E5 Compliance. To learn more about licensing requirements, see Microsoft 365 Tenant-Level Services Licensing Guidance.

To comply with business standards and industry regulations, organizations must protect sensitive information and prevent its inadvertent disclosure. Sensitive information can include financial data or personally identifiable information (PII) such as credit card numbers, social security numbers, or health records. With a data loss prevention (DLP) policy in the Office 365 Security & Compliance Center, you can identify, monitor, and automatically protect sensitive information across Office 365.

With a DLP policy, you can:

- **Identify sensitive information across many locations, such as Exchange Online, SharePoint Online, OneDrive for Business, and Microsoft Teams.**

  For example, you can identify any document containing a credit card number that's stored in any OneDrive for Business site, or you can monitor just the OneDrive sites of specific people.

- **Prevent the accidental sharing of sensitive information.**

  For example, you can identify any document or email containing a health record that's shared with people outside your organization, and then automatically block access to that document or block the email from being sent.

- **Monitor and protect sensitive information in the desktop versions of Excel, PowerPoint, and Word.**

  Just like in Exchange Online, SharePoint Online, and OneDrive for Business, these Office desktop programs include the same capabilities to identify sensitive information and apply DLP policies. DLP provides continuous monitoring when people share content in these Office programs.

- **Help users learn how to stay compliant without interrupting their workflow.**

  You can educate your users about DLP policies and help them remain compliant without blocking their work. For example, if a user tries to share a document containing sensitive information, a DLP policy can both send them an email notification and show them a policy tip in the context of the document library that allows them to override the policy if they have a business justification. The same policy tips also appear in Outlook on the web, Outlook, Excel, PowerPoint, and Word.

- **View DLP alerts and reports showing content that matches your organization's DLP policies.**

  To view alerts and metadata related to your DLP policies you can use the DLP Alerts Management Dashboard. You can also view policy match reports to assess how your organization is complying with a DLP policy. If a DLP policy allows users to override a policy tip and report a false positive, you can also view what users have reported

You create and manage DLP policies on the Data loss prevention page in the Microsoft 365 Compliance center.

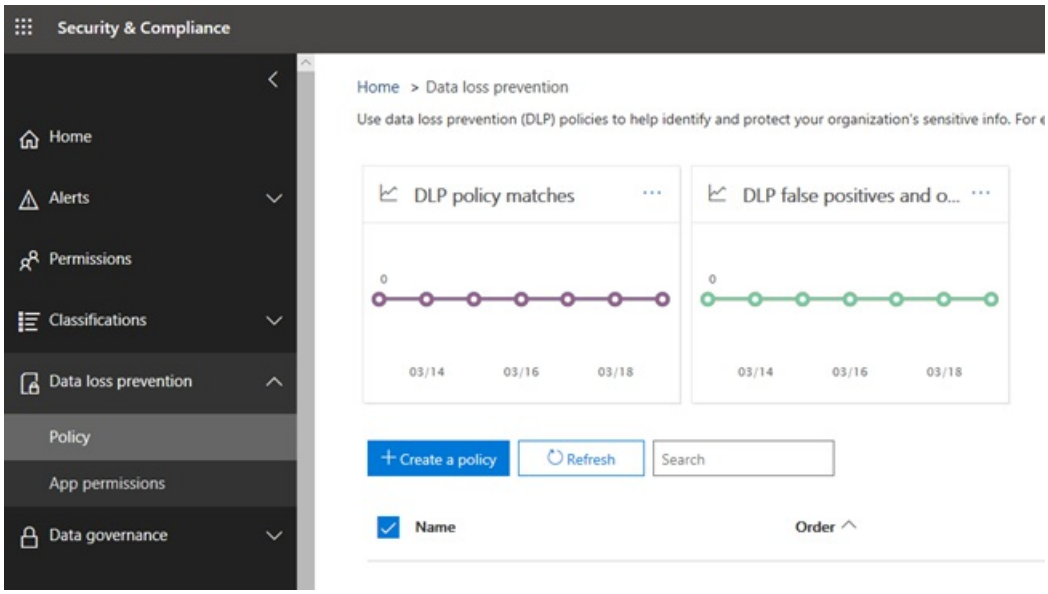

What a DLP policy contains

A DLP policy contains a few basic things:

- Where to protect the content: **locations** such as Exchange Online, SharePoint Online, and OneDrive for Business sites, as well as Microsoft Teams chat and channel messages.
- When and how to protect the content by enforcing **rules** comprised of:
  - **Conditions** the content must match before the rule is enforced. For example, a rule might be configured to look only for content containing Social Security numbers that's been shared with people outside your organization.
  - **Actions** that you want the rule to take automatically when content matching the conditions is found. For example, a rule might be configured to block access to a document and send both the user and compliance officer an email notification.

You can use a rule to meet a specific protection requirement, and then use a DLP policy to group together common protection requirements, such as all of the rules needed to comply with a specific regulation.

For example, you might have a DLP policy that helps you detect the presence of information subject to the Health Insurance Portability and Accountability Act (HIPAA). This DLP policy could help protect HIPAA data (the what) across all SharePoint Online sites and all OneDrive for Business sites (the where) by finding any document containing this sensitive information that's shared with people outside your organization (the conditions) and then blocking access to the document and sending a notification (the actions). These requirements are stored as individual rules and grouped together as a DLP policy to simplify management and reporting.

Locations

DLP policies are applied to sensitive items across Microsoft 365 locations and can be further scoped as detailed in this table.

LOCATION	INCLUDE/EXCLUDE BY
Exchange email	distribution groups
SharePoint sites	sites
OneDrive accounts	accounts
Teams chat and channel messages	accounts
Windows 10 devices	user or group
Microsoft Cloud App Security	instance

If you choose to include specific distribution groups in Exchange, the DLP policy will be scoped only to the members of that group. Similarly excluding a distribution group will exclude all the members of that distribution group from policy evaluation. You can choose to scope a policy to the members of distribution lists, dynamic distribution groups, and security groups. A DLP policy can contain no more than 50 such inclusions and exclusions.

If you choose to include or exclude specific SharePoint sites or OneDrive accounts, a DLP policy can contain no more than 100 such inclusions and exclusions. Although this limit exists, you can exceed this limit by applying either an org-wide policy or a policy that applies to entire locations.

Rules

> **NOTE**
> The default behavior of a DLP policy, when there is no alert configured, is not to alert or trigger. This applies only to default information types. For custom information types, the system will alert even if there is no action defined in the policy.

Rules are what enforce your business requirements on your organization's content. A policy contains one or more rules, and each rule consists of conditions and actions. For each rule, when the conditions are met, the actions are taken automatically. Rules are executed sequentially, starting with the highest-priority rule in each

policy.

A rule also provides options to notify users (with policy tips and email notifications) and admins (with email incident reports) that content has matched the rule.

Here are the components of a rule, each explained below.

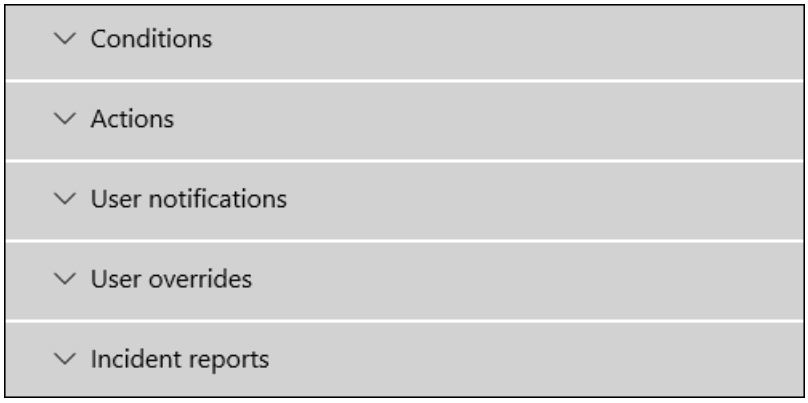

Conditions

Conditions are important because they determine what types of information you're looking for, and when to take an action. For example, you might choose to ignore content containing passport numbers unless the content contains more than 10 such numbers and is shared with people outside your organization.

Conditions focus on the **content**, such as what types of sensitive information you're looking for, and also on the **context**, such as who the document is shared with. You can use conditions to assign different actions to different risk levels. For example, sensitive content shared internally might be lower risk and require fewer actions than sensitive content shared with people outside the organization.

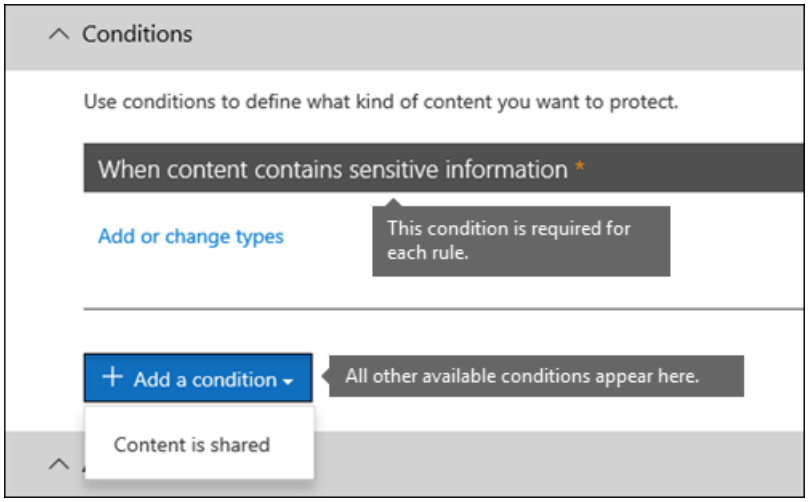

The conditions now available can determine if:

- Content contains a type of sensitive information.
- Content contains a label. For more information, see the below section Using a retention label as a condition in a DLP policy.
- Content is shared with people outside or inside your organization.

> **NOTE**
> Users who have non-guest accounts in a host organization's Active Directory or Azure Active Directory tenant are considered as people inside the organization.

Types of sensitive information

A DLP policy can help protect sensitive information, which is defined as a **sensitive information type**. Microsoft 365 includes definitions for many common sensitive information types across many different regions that are ready for you to use, such as a credit card number, bank account numbers, national ID numbers, and passport numbers.

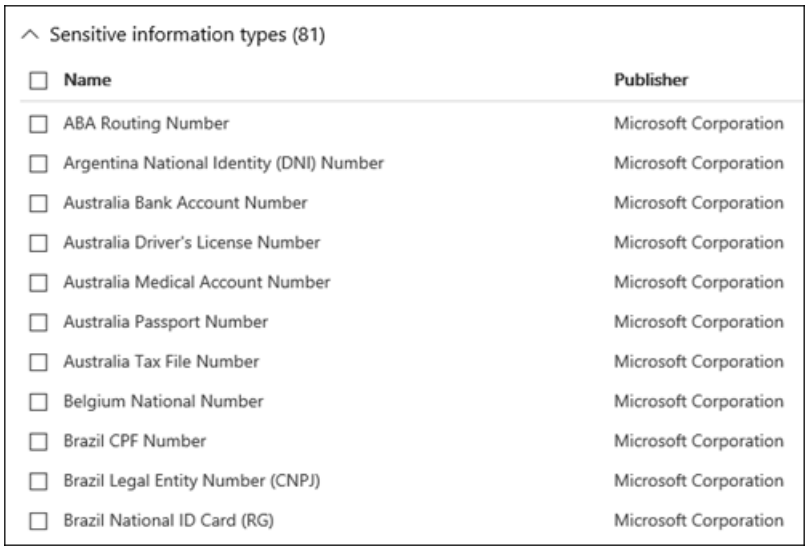

When a DLP policy looks for a sensitive information type such as a credit card number, it doesn't simply look for a 16-digit number. Each sensitive information type is defined and detected by using a combination of:

- Keywords.
- Internal functions to validate checksums or composition.
- Evaluation of regular expressions to find pattern matches.
- Other content examination.

This helps DLP detection achieve a high degree of accuracy while reducing the number of false positives that can interrupt peoples' work.

Actions

When content matches a condition in a rule, you can apply actions to automatically protect the content.

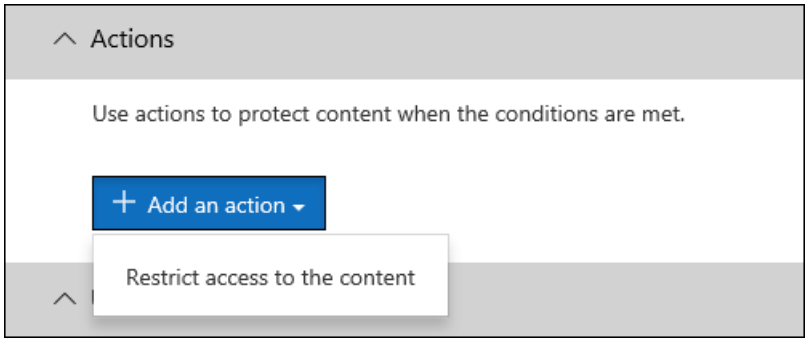

With the actions now available, you can:

- **Restrict access to the content** Depending on your need, you can restrict access to content in three ways:

  1. Restrict access to content for everyone.
  2. Restrict access to content for people outside the organization.
  3. Restrict access to "Anyone with the link."

  For site content, this means that permissions for the document are restricted for everyone except the primary site collection administrator, document owner, and person who last modified the document. These people can remove the sensitive information from the document or take other remedial action.

When the document is in compliance, the original permissions are automatically restored. When access to a document is blocked, the document appears with a special policy tip icon in the library on the site.

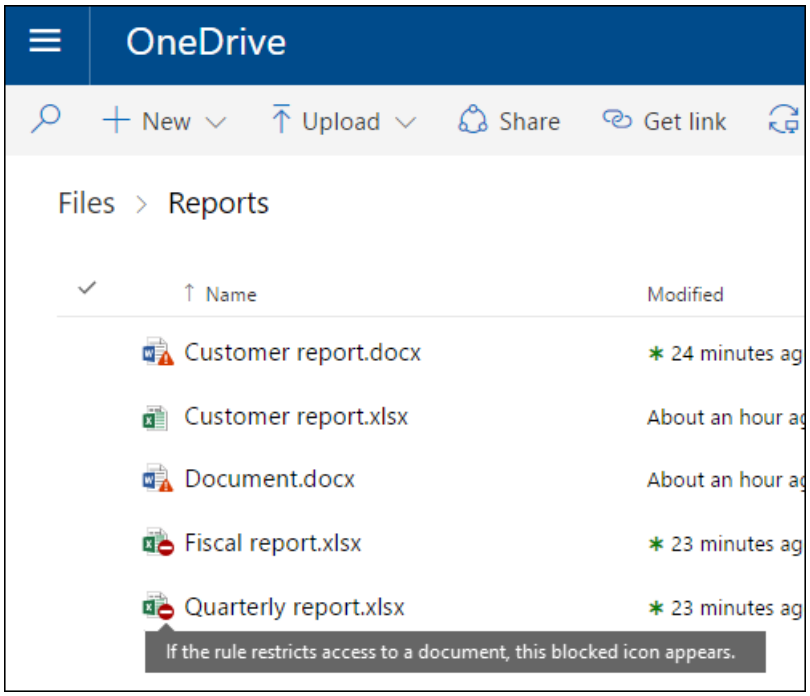

For email content, this action blocks the message from being sent. Depending on how the DLP rule is configured, the sender sees an NDR or (if the rule uses a notification) a policy tip and/or email notification.

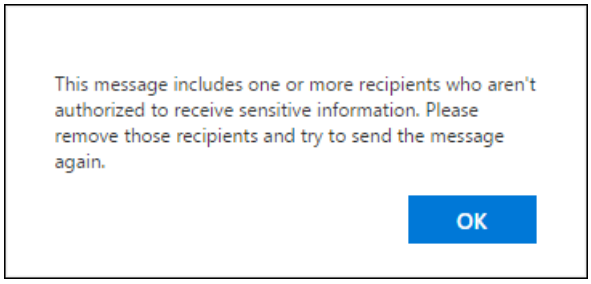

User notifications and user overrides

You can use notifications and overrides to educate your users about DLP policies and help them remain compliant without blocking their work. For example, if a user tries to share a document containing sensitive information, a DLP policy can both send them an email notification and show them a policy tip in the context of the document library that allows them to override the policy if they have a business justification.

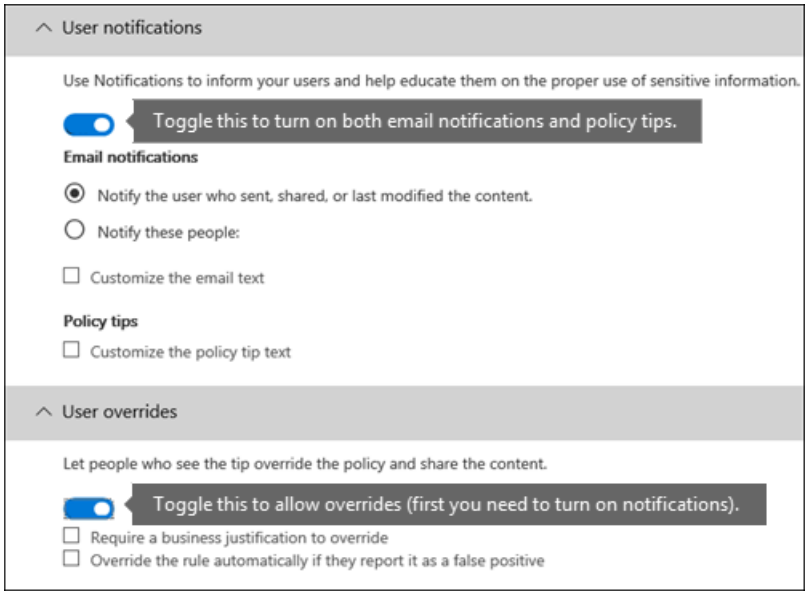

The email can notify the person who sent, shared, or last modified the content and, for site content, the primary site collection administrator and document owner. In addition, you can add or remove whomever you choose from the email notification.

In addition to sending an email notification, a user notification displays a policy tip:

- In Outlook and Outlook on the web.
- For the document on a SharePoint Online or OneDrive for Business site.
- In Excel, PowerPoint, and Word, when the document is stored on a site included in a DLP policy.

The email notification and policy tip explain why content conflicts with a DLP policy. If you choose, the email notification and policy tip can allow users to override a rule by reporting a false positive or providing a business justification. This can help you educate users about your DLP policies and enforce them without preventing people from doing their work. Information about overrides and false positives is also logged for reporting (see below about the DLP reports) and included in the incident reports (next section), so that the compliance officer can regularly review this information.

Here's what a policy tip looks like in a OneDrive for Business account.

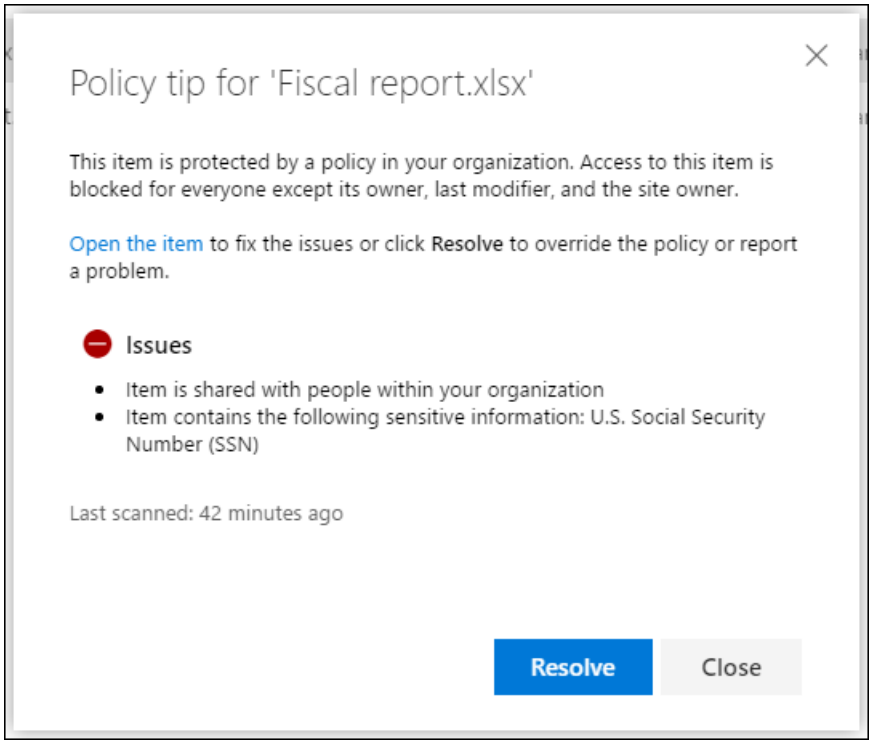

To learn more about user notifications and policy tips in DLP policies, see Use notifications and policy tips.

Alerts and Incident reports

When a rule is matched, you can send an alert email to your compliance officer (or any person(s) you choose) with details of the alert. This alert email will carry a link of the DLP Alerts Management Dashboard which the compliance officer can go to view the details of alert and events. The dashboard contains details of the event that triggered the alert along with details of the DLP policy matched and the sensitive content detected.

In addition, you can also send an incident report with details of the event. This report includes information about the item that was matched, the actual content that matched the rule, and the name of the person who last modified the content. For email messages, the report also includes as an attachment the original message that matches a DLP policy.

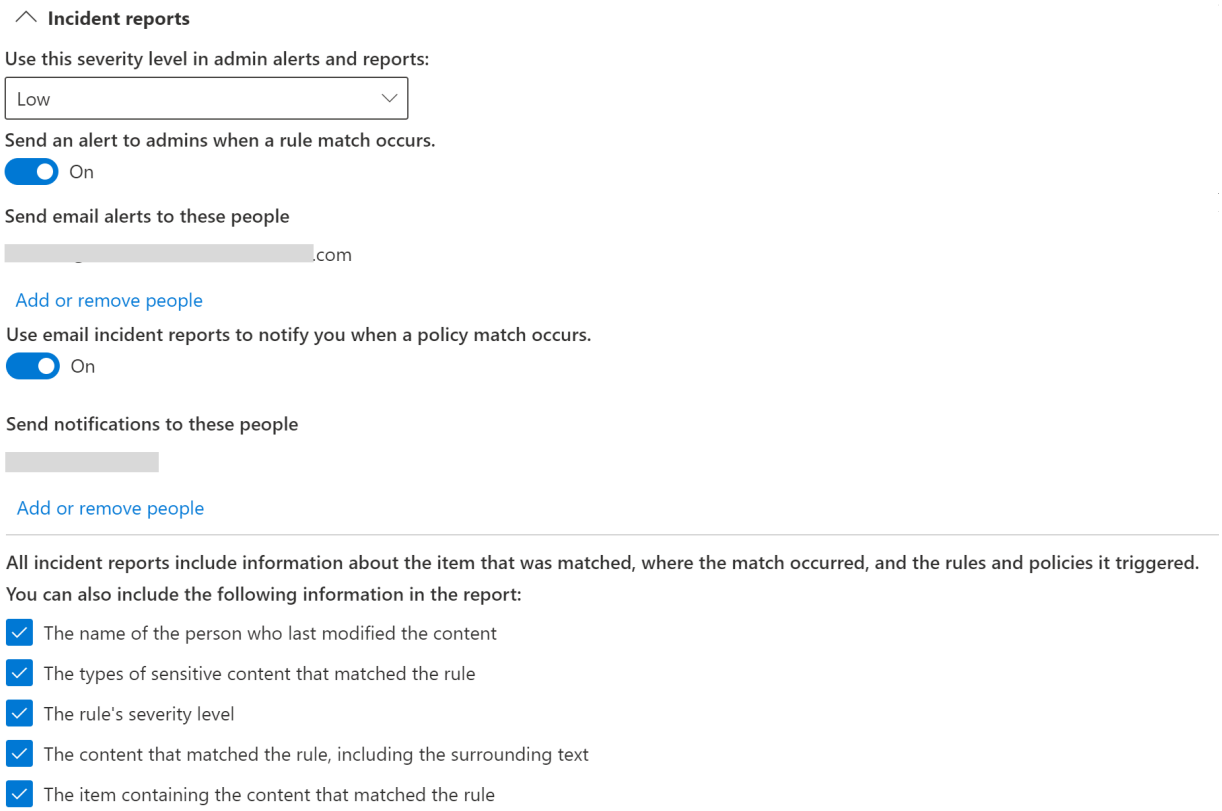

DLP scans email differently from items in SharePoint Online or OneDrive for Business. In SharePoint Online and OneDrive for Business, DLP scans existing items as well as new ones and generates an alert and incident report whenever a match is found. In Exchange Online, DLP only scans new email messages and generates a report if there is a policy match. DLP ***does not*** scan or match previously existing email items that are stored in a mailbox or archive.

Grouping and logical operators

Often your DLP policy has a straightforward requirement, such as to identify all content that contains a U.S. Social Security Number. However, in other scenarios, your DLP policy might need to identify more loosely defined data.

For example, to identify content subject to the U.S. Health Insurance Act (HIPAA), you need to look for:

- Content that contains specific types of sensitive information, such as a U.S. Social Security Number or Drug Enforcement Agency (DEA) Number.

  AND

- Content that's more difficult to identify, such as communications about a patient's care or descriptions of medical services provided. Identifying this content requires matching keywords from very large keyword lists, such as the International Classification of Diseases (ICD-9-CM or ICD-10-CM).

You can easily identify such loosely defined data by using grouping and logical operators (AND, OR). When you create a DLP policy, you can:

- Group sensitive information types.
- Choose the logical operator between the sensitive information types within a group and between the groups themselves.

Choosing the operator within a group

Within a group, you can choose whether any or all of the conditions in that group must be satisfied for the content to match the rule.

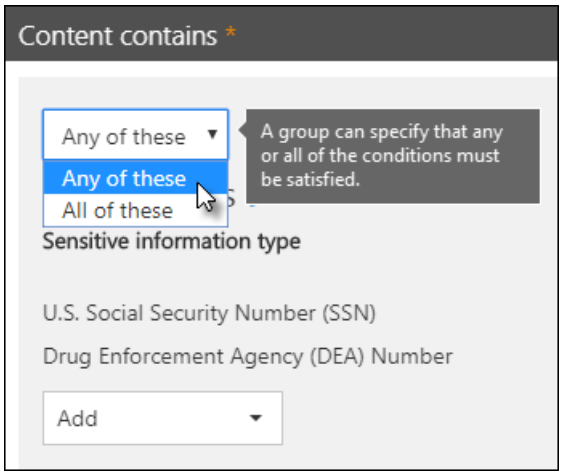

Adding a group

You can quickly add a group, which can have its own conditions and operator within that group.

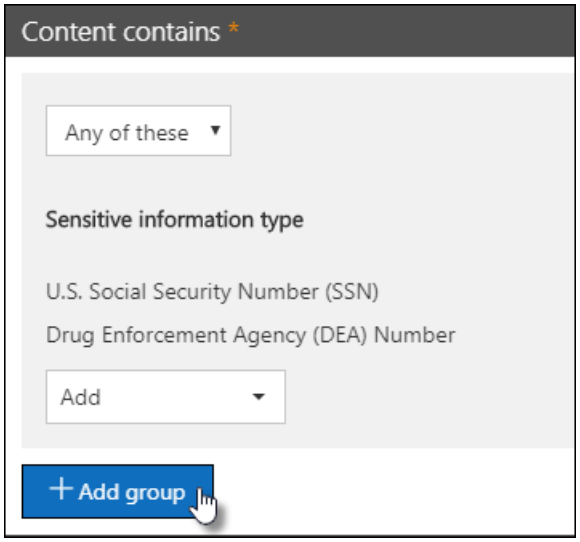

Choosing the operator between groups

Between groups, you can choose whether the conditions in just one group or all of the groups must be satisfied for the content to match the rule.

For example, the built-in **U.S. HIPAA** policy has a rule that uses an **AND** operator between the groups so that it identifies content that contains:

- from the group **PII Identifiers** (at least one SSN number **OR** DEA number)

  AND

- from the group **Medical Terms** (at least one ICD-9-CM keyword **OR** ICD-10-CM keyword)

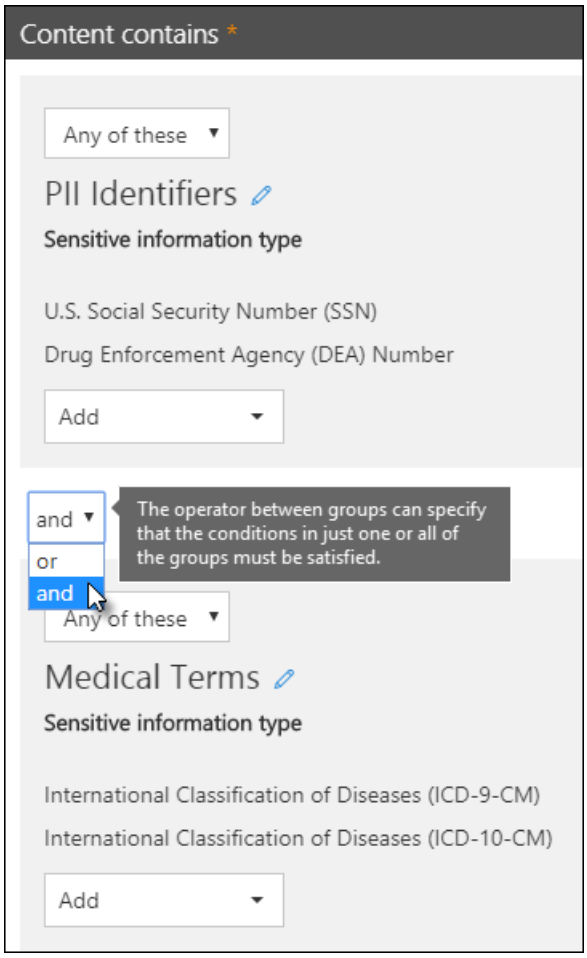

The priority by which rules are processed

When you create rules in a policy, each rule is assigned a priority in the order in which it's created — meaning, the rule created first has first priority, the rule created second has second priority, and so on.

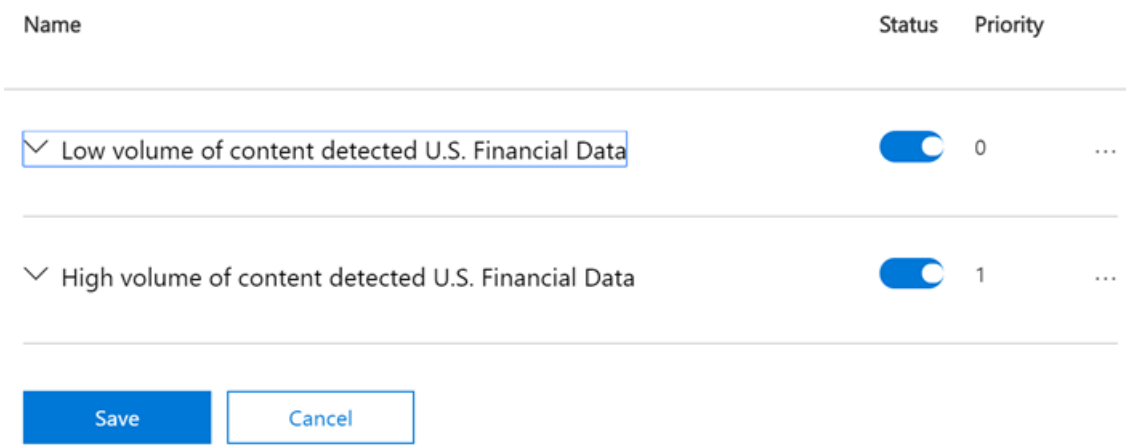

After you have set up more than one DLP policy, you can change the priority of one or more policies. To do that, select a policy, choose **Edit policy**, and use the **Priority** list to specify its priority.

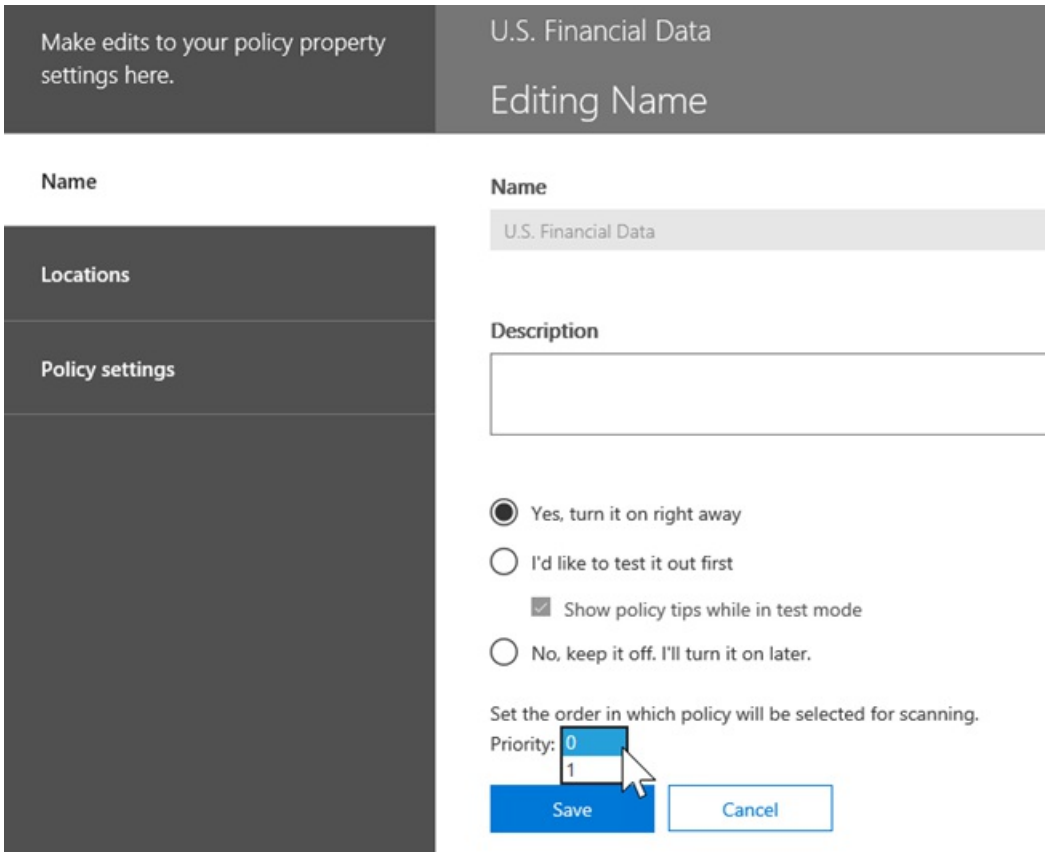

When content is evaluated against rules, the rules are processed in priority order. If content matches multiple rules, the rules are processed in priority order and the most restrictive action is enforced. For example, if content matches all of the following rules, Rule 3 is enforced because it's the highest priority, most restrictive rule:

- Rule 1: only notifies users
- Rule 2: notifies users, restricts access, and allows user overrides
- Rule 3: notifies users, restricts access, and does not allow user overrides
- Rule 4: only notifies users
- Rule 5: restricts access
- Rule 6: notifies users, restricts access, and does not allow user overrides

In this example, note that matches for all of the rules are recorded in the audit logs and shown in the DLP reports, even though only the most restrictive rule is enforced.

Regarding policy tips, note that:

- Only the policy tip from the highest priority, most restrictive rule will be shown. For example, a policy tip from a rule that blocks access to content will be shown over a policy tip from a rule that simply sends a notification. This prevents people from seeing a cascade of policy tips.
- If the policy tips in the most restrictive rule allow people to override the rule, then overriding this rule also overrides any other rules that the content matched.

Tuning rules to make them easier or harder to match

After people create and turn on their DLP policies, they sometimes run into these issues:

- Too much content that **is not** sensitive information matches the rules — in other words, too many false positives.

- Too little content that **is** sensitive information matches the rules. In other words, the protective actions aren't being enforced on the sensitive information.

To address these issues, you can tune your rules by adjusting the instance count and match accuracy to make it harder or easier for content to match the rules. Each sensitive information type used in a rule has both an instance count and match accuracy.

Instance count

Instance count means simply how many occurrences of a specific type of sensitive information must be present for content to match the rule. For example, content matches the rule shown below if between 1 and 9 unique U.S. or U.K. passport numbers are identified.

Note that the instance count includes only **unique** matches for sensitive information types and keywords. For example, if an email contains 10 occurrences of the same credit card number, those 10 occurrences count as a single instance of a credit card number.

To use instance count to tune rules, the guidance is straightforward:

- To make the rule easier to match, decrease the **min** count and/or increase the **max** count. You can also set **max** to **any** by deleting the numerical value.
- To make the rule harder to match, increase the **min** count.

Typically, you use less restrictive actions, such as sending user notifications, in a rule with a lower instance count (for example, 1-9). And you use more restrictive actions, such as restricting access to content without allowing user overrides, in a rule with a higher instance count (for example, 10-any).

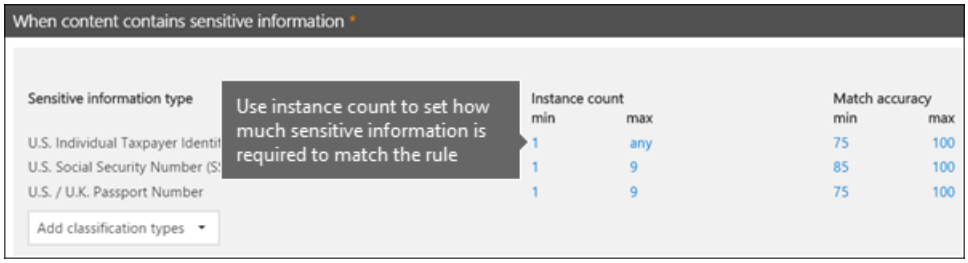

Match accuracy

As described above, a sensitive information type is defined and detected by using a combination of different types of evidence. Commonly, a sensitive information type is defined by multiple such combinations, called patterns. A pattern that requires less evidence has a lower match accuracy (or confidence level), while a pattern that requires more evidence has a higher match accuracy (or confidence level). To learn more about the actual patterns and confidence levels used by every sensitive information type, see Sensitive information type entity definitions.

For example, the sensitive information type named Credit Card Number is defined by two patterns:

- A pattern with 65% confidence that requires:
  - A number in the format of a credit card number.
  - A number that passes the checksum.
- A pattern with 85% confidence that requires:
  - A number in the format of a credit card number.
  - A number that passes the checksum.
  - A keyword or an expiration date in the right format.

You can use these confidence levels (or match accuracy) in your rules. Typically, you use less restrictive actions,

such as sending user notifications, in a rule with lower match accuracy. And you use more restrictive actions, such as restricting access to content without allowing user overrides, in a rule with higher match accuracy.

It's important to understand that when a specific type of sensitive information, such as a credit card number, is identified in content, only a single confidence level is returned:

- If all of the matches are for a single pattern, the confidence level for that pattern is returned.
- If there are matches for more than one pattern (that is, there are matches with two different confidence levels), a confidence level higher than any of the single patterns alone is returned. This is the tricky part. For example, for a credit card, if both the 65% and 85% patterns are matched, the confidence level returned for that sensitive information type is greater than 90% because more evidence means more confidence.

So if you want to create two mutually exclusive rules for credit cards, one for the 65% match accuracy and one for the 85% match accuracy, the ranges for match accuracy would look like this. The first rule picks up only matches of the 65% pattern. The second rule picks up matches with **at least one** 85% match and **can potentially have** other lower-confidence matches.

For these reasons, the guidance for creating rules with different match accuracies is:

- The lowest confidence level typically uses the same value for **min** and **max** (not a range).
- The highest confidence level is typically a range from just above the lower confidence level to 100.
- Any in-between confidence levels typically range from just above the lower confidence level to just below the higher confidence level.

Using a retention label as a condition in a DLP policy

When you use a previously created and published retention label as a condition in a DLP policy, there are some things to be aware of:

- The retention label must be created and published before you attempt to use it as a condition in a DLP policy.
- Published retention labels can take from one to seven days to sync. For more information, see When retention labels become available to apply for retention labels published in a retention policy, and How long it takes for retention labels to take effect for retention labels that are auto-published.
- Using a retention label in a policy **is only supported for items in SharePoint and OneDrive***.

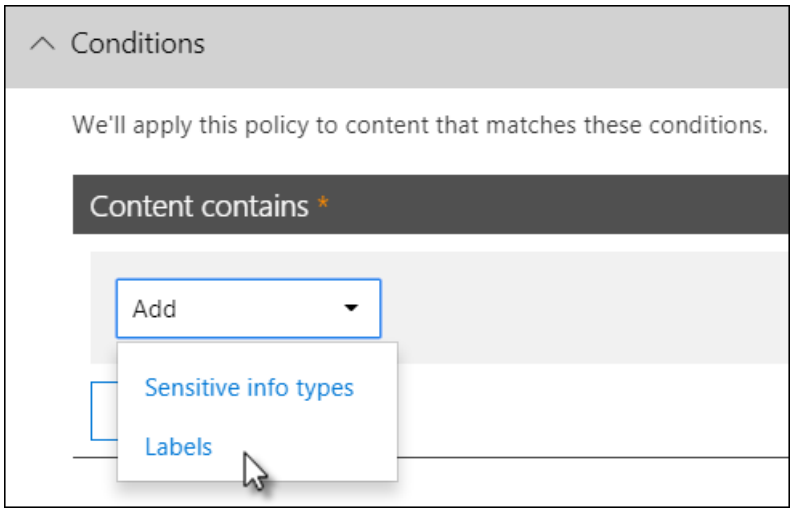

You might want to use a retention label in a DLP policy if you have items that are under retention and disposition, and you also want to apply other controls to them, for example:

- You published a retention label named **tax year 2018**, which when applied to tax documents from 2018 that are stored in SharePoint retains them for 10 years then disposes of them. You also don't want those items being shared outside your organization, which you can do with a DLP policy.

> **IMPORTANT**
> You'll get this error if you specify a retention label as a condition in a DLP policy and you also include Exchange and/or Teams as a location: **"Protecting labeled content in email and teams messages isn't supported. Either remove the label below or turn off Exchange and Teams as a location."** This is because Exchange transport does not evaluate the label metadata during message submission and delivery.

Using a sensitivity label as a condition in a DLP policy

Sensitivity label as a condition in DLP policies is currently in preview. Learn more.

How this feature relates to other features

Several features can be applied to content containing sensitive information:

- A retention label and a retention policy can both enforce **retention** actions on this content.
- A DLP policy can enforce **protection** actions on this content. And before enforcing these actions, a DLP policy can require other conditions to be met in addition to the content containing a label.

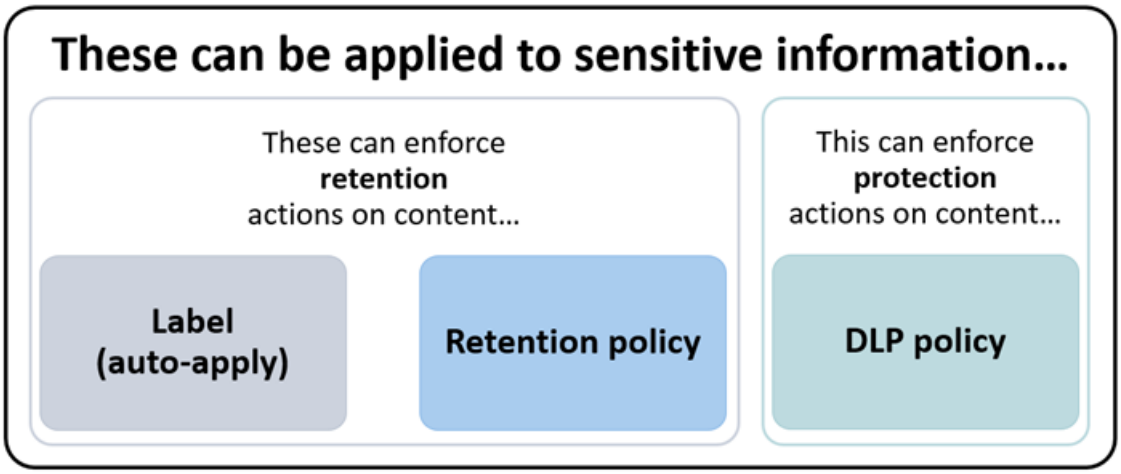

Note that a DLP policy has a richer detection capability than a label or retention policy applied to sensitive information. A DLP policy can enforce protective actions on content containing sensitive information, and if the sensitive information is removed from the content, those protective actions are undone the next time the

content's scanned. But if a retention policy or label is applied to content containing sensitive information, that's a one-time action that won't be undone even if the sensitive information is removed.

By using a label as a condition in a DLP policy, you can enforce both retention and protection actions on content with that label. You can think of content containing a label exactly like content containing sensitive information - both a label and a sensitive information type are properties used to classify content, so that you can enforce actions on that content.

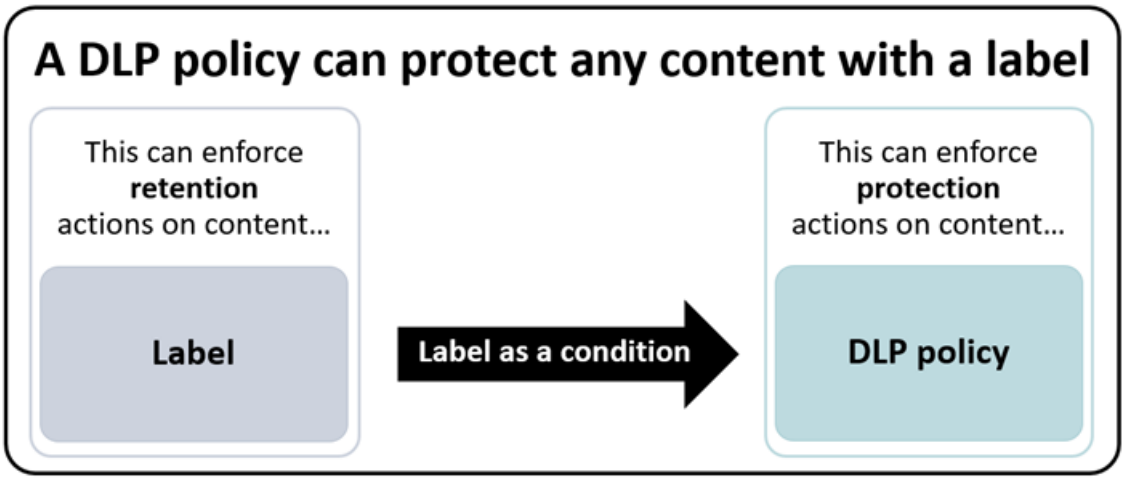

Simple settings vs. advanced settings

When you create a DLP policy, you'll choose between simple or advanced settings:

- **Simple settings** make it easy to create the most common type of DLP policy without using the rule editor to create or modify rules.
- **Advanced settings** use the rule editor to give you complete control over every setting for your DLP policy.

Don't worry, under the covers, simple settings and advanced settings work exactly the same, by enforcing rules comprised of conditions and actions -- only with simple settings, you don't see the rule editor. It's a quick way to create a DLP policy.

Simple settings

By far, the most common DLP scenario is creating a policy to help protect content containing sensitive information from being shared with people outside your organization, and taking an automatic remediating action such as restricting who can access the content, sending end-user or admin notifications, and auditing the event for later investigation. People use DLP to help prevent the inadvertent disclosure of sensitive information.

To simplify achieving this goal, when you create a DLP policy, you can choose **Use simple settings**. These settings provide everything you need to implement the most common DLP policy, without having to go into the rule editor.

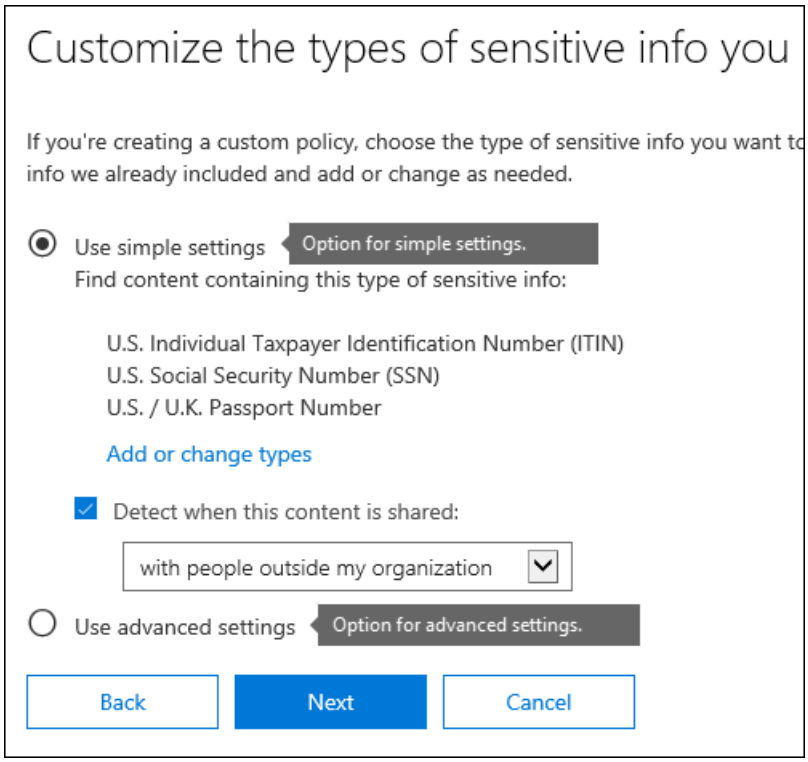

Advanced settings

If you need to create more customized DLP policies, you can choose **Use advanced settings**.

The advanced settings present you with the rule editor, where you have full control over every possible option, including the instance count and match accuracy (confidence level) for each rule.

To jump to a section quickly, click an item in the top navigation of the rule editor to go to that section below.

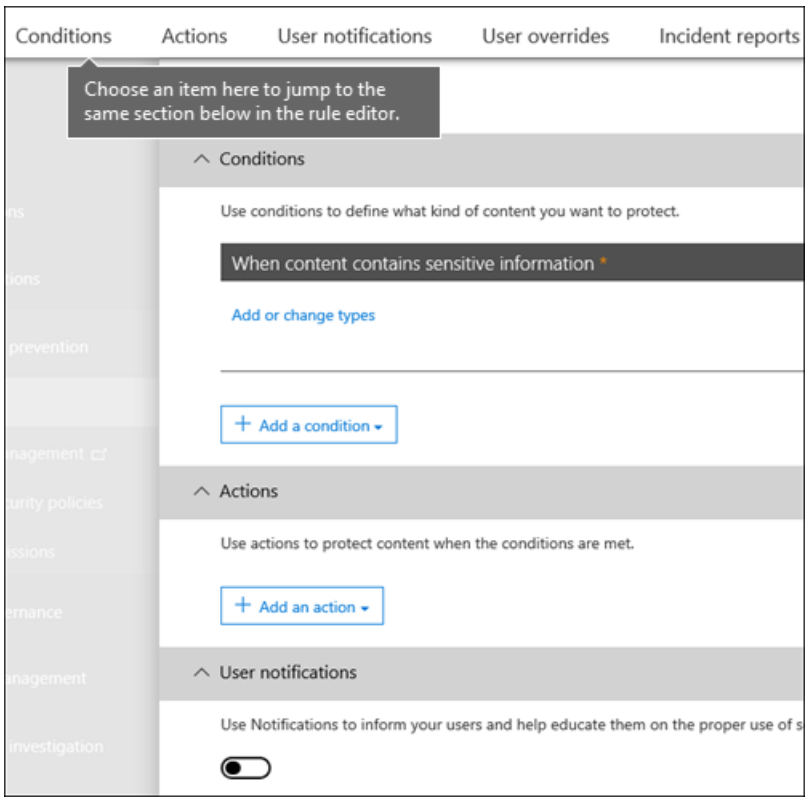

DLP policy templates

The first step in creating a DLP policy is choosing what information to protect. By starting with a DLP template, you save the work of building a new set of rules from scratch, and figuring out which types of information should be included by default. You can then add to or modify these requirements to fine tune the rule to meet

your organization's specific requirements.

A preconfigured DLP policy template can help you detect specific types of sensitive information, such as HIPAA data, PCI-DSS data, Gramm-Leach-Bliley Act data, or even locale-specific personally identifiable information (P.I.). To make it easy for you to find and protect common types of sensitive information, the policy templates included in Microsoft 365 already contain the most common sensitive information types necessary for you to get started.

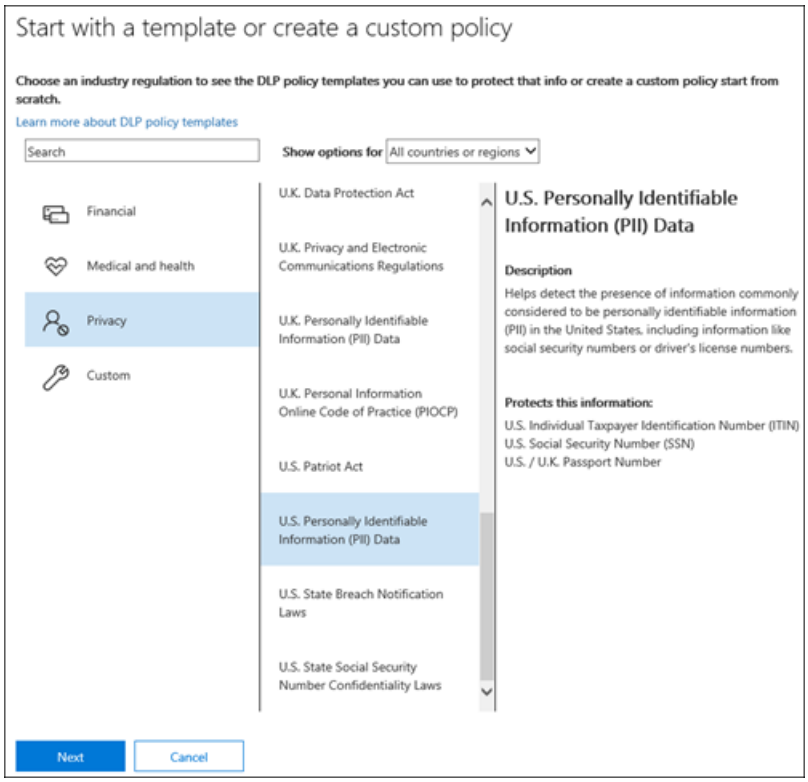

Your organization may also have its own specific requirements, in which case you can create a DLP policy from scratch by choosing the **Custom policy** option. A custom policy is empty and contains no premade rules.

Roll out DLP policies gradually with test mode

When you create your DLP policies, you should consider rolling them out gradually to assess their impact and test their effectiveness before fully enforcing them. For example, you don't want a new DLP policy to unintentionally block access to thousands of documents that people require access to in order to get their work done.

If you're creating DLP policies with a large potential impact, we recommend following this sequence:

1. **Start in test mode without Policy Tips** and then use the DLP reports and any incident reports to assess the impact. You can use DLP reports to view the number, location, type, and severity of policy matches. Based on the results, you can fine tune the rules as needed. In test mode, DLP policies will not impact the productivity of people working in your organization.
2. **Move to Test mode with notifications and Policy Tips** so that you can begin to teach users about your compliance policies and prepare them for the rules that are going to be applied. At this stage, you can also ask users to report false positives so that you can further refine the rules.
3. **Start full enforcement on the policies** so that the actions in the rules are applied and the content's protected. Continue to monitor the DLP reports and any incident reports or notifications to make sure that the results are what you intend.

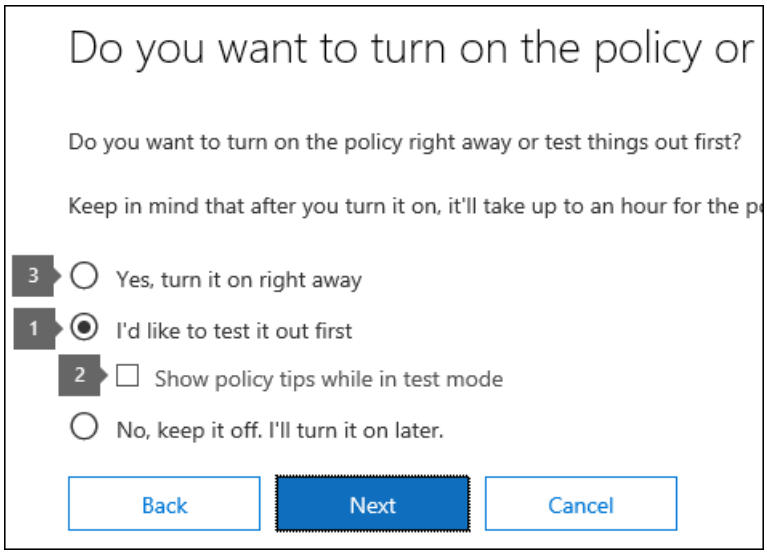

You can turn off a DLP policy at any time, which affects all rules in the policy. However, each rule can also be turned off individually by toggling its status in the rule editor.

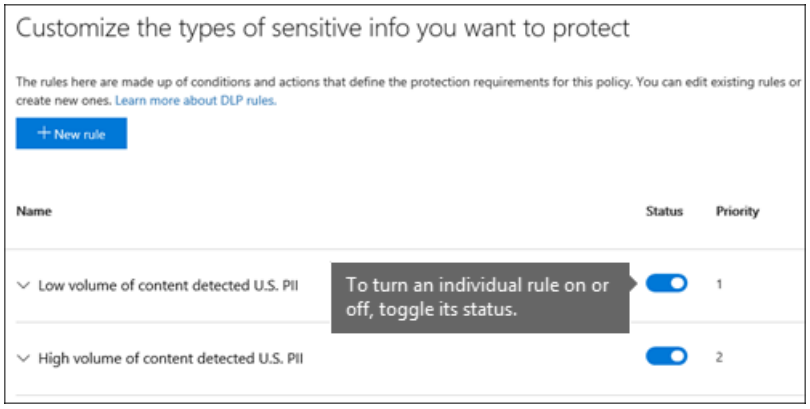

You can also change the priority of multiple rules in a policy. To do that, open a policy for editing. In a row for a rule, choose the ellipses (...), and then choose an option, such as **Move down** or **Bring to last**.

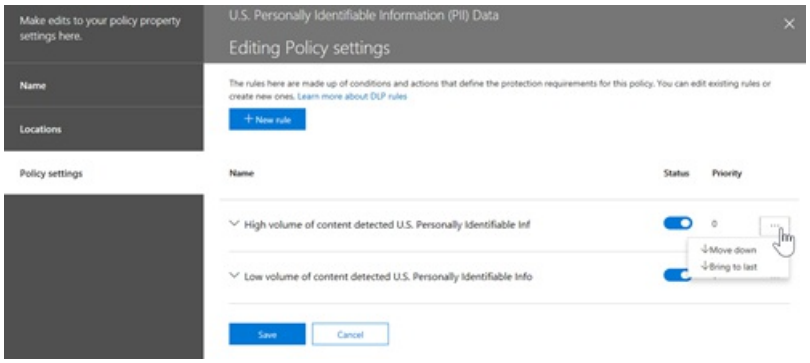

DLP reports

After you create and turn on your DLP policies, you'll want to verify that they're working as you intended and helping you stay compliant. With DLP reports, you can quickly view the number of DLP policy and rule matches over time, and the number of false positives and overrides. For each report, you can filter those matches by location, time frame, and even narrow it down to a specific policy, rule, or action.

With the DLP reports, you can get business insights and:

- Focus on specific time periods and understand the reasons for spikes and trends.
- Discover business processes that violate your organization's compliance policies.
- Understand any business impact of the DLP policies.

In addition, you can use the DLP reports to fine tune your DLP policies as you run them.

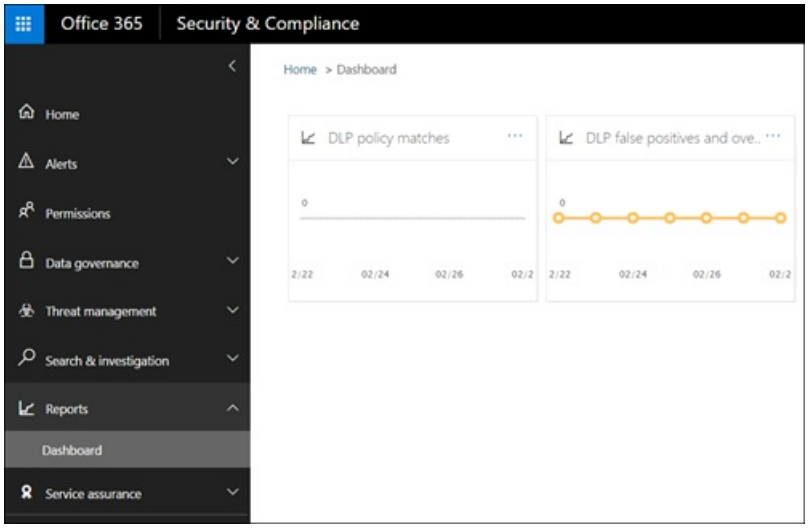

How DLP policies work

DLP detects sensitive information by using deep content analysis (not just a simple text scan). This deep content analysis uses keyword matches, dictionary matches, the evaluation of regular expressions, internal functions, and other methods to detect content that matches your DLP policies. Potentially only a small percentage of your data is considered sensitive. A DLP policy can identify, monitor, and automatically protect just that data, without impeding or affecting people who work with the rest of your content.

Policies are synced

After you create a DLP policy in the Security & Compliance Center, it's stored in a central policy store, and then synced to the various content sources, including:

- Exchange Online, and from there to Outlook on the web and Outlook.
- OneDrive for Business sites.
- SharePoint Online sites.
- Office desktop programs (Excel, PowerPoint, and Word).
- Microsoft Teams channels and chat messages.

After the policy's synced to the right locations, it starts to evaluate content and enforce actions.

Policy evaluation in OneDrive for Business and SharePoint Online sites

Across all of your SharePoint Online sites and OneDrive for Business sites, documents are constantly changing — they're continually being created, edited, shared, and so on. This means documents can conflict or become compliant with a DLP policy at any time. For example, a person can upload a document that contains no sensitive information to their team site, but later, a different person can edit the same document and add sensitive information to it.

For this reason, DLP policies check documents for policy matches frequently in the background. You can think of this as asynchronous policy evaluation.

How it works

As people add or change documents in their sites, the search engine scans the content, so that you can search for it later. While this is happening, the content's also scanned for sensitive information and to check if it's shared. Any sensitive information that's found is stored securely in the search index, so that only the compliance team can access it, but not typical users. Each DLP policy that you've turned on runs in the background (asynchronously), checking search frequently for any content that matches a policy, and applying actions to

protect it from inadvertent leaks.

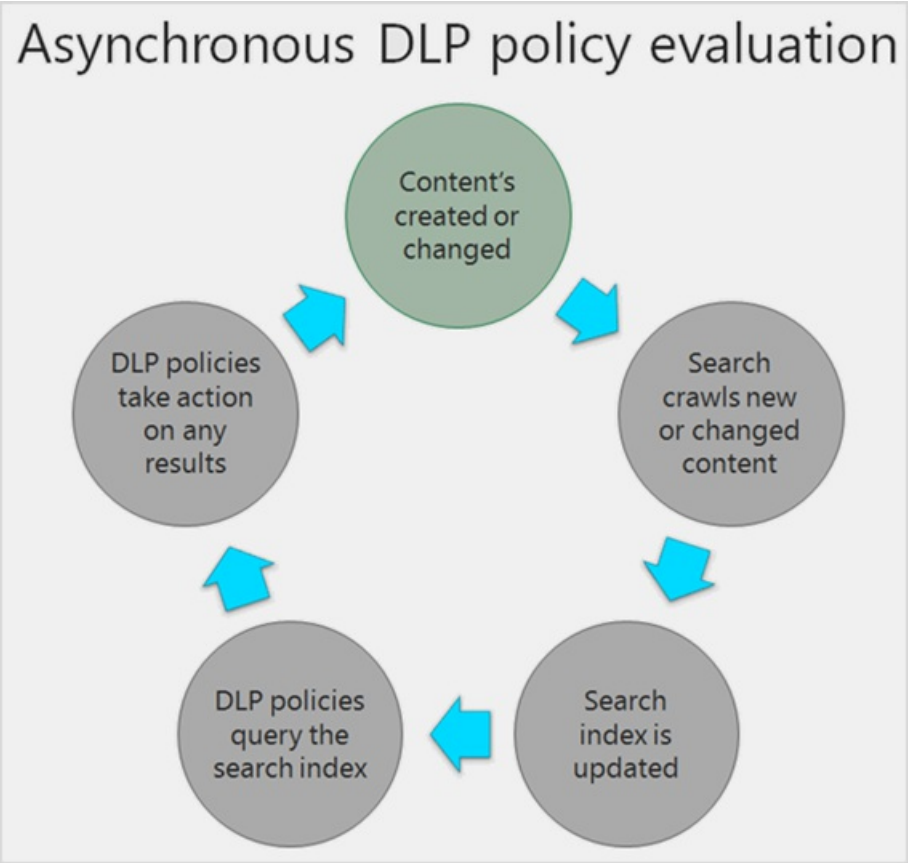

Finally, documents can conflict with a DLP policy, but they can also become compliant with a DLP policy. For example, if a person adds credit card numbers to a document, it might cause a DLP policy to block access to the document automatically. But if the person later removes the sensitive information, the action (in this case, blocking) is automatically undone the next time the document is evaluated against the policy.

DLP evaluates any content that can be indexed. For more information on what file types are crawled by default, see Default crawled file name extensions and parsed file types in SharePoint Server.

> **NOTE**
> External sharing of new files in SharePoint can be blocked by default until at least one DLP policy scans the new item. See, Mark new files as sensitive by default for detailed information.

Policy evaluation in Exchange Online, Outlook, and Outlook on the web

When you create a DLP policy that includes Exchange Online as a location, the policy's synced from the Office 365 Security & Compliance Center to Exchange Online, and then from Exchange Online to Outlook on the web and Outlook.

When a message is being composed in Outlook, the user can see policy tips as the content being created is evaluated against DLP policies. And after a message is sent, it's evaluated against DLP policies as a normal part of mail flow, along with Exchange mail flow rules (also known as transport rules) and DLP policies created in the Exchange admin center. DLP policies scan both the message and any attachments.

Policy evaluation in the Office desktop programs

Excel, PowerPoint, and Word include the same capability to identify sensitive information and apply DLP policies as SharePoint Online and OneDrive for Business. These Office programs sync their DLP policies directly from the central policy store, and then continuously evaluate the content against the DLP policies when people work with documents opened from a site that's included in a DLP policy.

DLP policy evaluation in Office is designed not to affect the performance of the programs or the productivity of people working on content. If they're working on a large document, or the user's computer is busy, it might take a few seconds for a policy tip to appear.

Policy evaluation in Microsoft Teams

When you create a DLP policy that includes Microsoft Teams as a location, the policy's synced from the Office 365 Security & Compliance Center to user accounts and Microsoft Teams channels and chat messages. Depending on how DLP policies are configured, when someone attempts to share sensitive information in a Microsoft Teams chat or channel message, the message can be blocked or revoked. And, documents that contain sensitive information and that are shared with guests (external users) won't open for those users. To learn more, see Data loss prevention and Microsoft Teams.

Permissions

Members of your compliance team who will create DLP policies need permissions to the Security & Compliance Center. By default, your tenant admin will have access to this location and can give compliance officers and other people access to the Security & Compliance Center, without giving them all of the permissions of a tenant admin. To do this, we recommend that you:

1. Create a group in Microsoft 365 and add compliance officers to it.
2. Create a role group on the **Permissions** page of the Security & Compliance Center.
3. While creating the role group, use the **Choose Roles** section to add the following role to the Role Group: **DLP Compliance Management**.
4. Use the **Choose Members** section to add the Microsoft 365 group you created before to the role group.

You can also create a role group with view-only privileges to the DLP policies and DLP reports by granting the **View-Only DLP Compliance Management** role.

For more information, see Give users access to the Office 365 Compliance Center.

These permissions are required only to create and apply a DLP policy. Policy enforcement does not require access to the content.

Find the DLP cmdlets

To use most of the cmdlets for the Security & Compliance Center, you need to:

1. Connect to the Office 365 Security & Compliance Center using remote PowerShell.
2. Use any of these policy-and-compliance-dlp cmdlets.

However, DLP reports need pull data from across Microsoft 365, including Exchange Online. For this reason, **the cmdlets for the DLP reports are available in Exchange Online Powershell -- not in Security & Compliance Center Powershell**. Therefore, to use the cmdlets for the DLP reports, you need to:

1. Connect to Exchange Online using remote PowerShell.
2. Use any of these cmdlets for the DLP reports:
   - Get-DlpDetectionsReport
   - Get-DlpDetailReport

More information

- Create a DLP policy from a template

- Send notifications and show policy tips for DLP policies
- Create a DLP policy to protect documents with FCI or other properties
- What the DLP policy templates include
- Sensitive information type entity definitions
- What the DLP functions look for
- Create a custom sensitive information type

Learn about Microsoft 365 Endpoint data loss prevention

2/18/2021 • 4 minutes to read • Edit Online

You can use Microsoft 365 data loss prevention (DLP) to monitor the actions that are being taken on items you've determined to be sensitive and to help prevent the unintentional sharing of those items. For more information on DLP, see Overview of data loss prevention.

Endpoint data loss prevention (Endpoint DLP) extends the activity monitoring and protection capabilities of DLP to sensitive items that are on Windows 10 devices. Once devices are onboarded into the Microsoft 365 compliance solutions, the information about what users are doing with sensitive items is made visible in activity explorer and you can enforce protective actions on those items via DLP policies.

Endpoint activities you can monitor and take action on

Microsoft Endpoint DLP enables you to audit and manage the following types of activities users take on sensitive items on devices running Windows 10.

ACTIVITY	DESCRIPTION	AUDITABLE/RESTICTABLE
upload to cloud service, or access by unallowed browsers	Detects when a user attempts to upload an item to a restricted service domain or access an item through a browser. If they are using a browser that is listed in DLP as an being an unallowed browser, the upload activity will be blocked and the user is redirected to use Edge Chromium. Edge Chromium will then either allow or block the upload or access based on the DLP policy configuration	auditable and restrictable
copy to other app	Detects when a user attempts to copy information from a protected item and then paste it into another app, process or item. Copying and pasting information within the same app, process, or item is not detected by this activity.	auditable and restrictable
copy to USB removable media	Detects when a user attempts to copy an item or information to removable media or USB device.	auditable and restrictable
copy to a network share	Detects when a user attempts to copy an item to a network share or mapped network drive	auditable and restrictable
print a document	Detects when a user attempts to print a protected item to a local or network printer.	auditable and restrictable

ACTIVITY	DESCRIPTION	AUDITABLE/RESTICTABLE
create an item	Detects when a user creates an item	auditable
rename an item	Detects when a user renames an item	auditable

Monitored files

Endpoint DLP supports monitoring of these file types:

- Word files
- PowerPoint files
- Excel files
- PDF files
- .csv files
- .tsv files
- .txt files
- .rtf files
- .c files
- .class files
- .cpp files
- .cs files
- .h files
- .java files

By default, endpoint DLP audits the activities for these file types, even if there isn't a policy match. If you only want monitoring data from policy matches, you can turn off the **Always audit file activity for devices** in the endpoint DLP global settings. No matter what, activities on any Word, PowerPoint, Excel, PDF, and .csv file are always audited.

Endpoint DLP monitors activity-based on MIME type, so activities will be captured even if the file extension is changed.

What's different in Endpoint DLP

There are a few extra concepts that you need to be aware of before you dig into Endpoint DLP.

Enabling Device management

Device management is the functionality that enables the collection of telemetry from devices and brings it into Microsoft 365 compliance solutions like Endpoint DLP and Insider Risk management. You'll need to onboard all devices you want to use as locations in DLP policies.

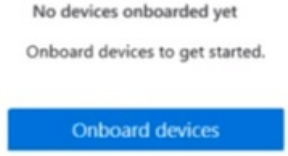

Onboarding and offboarding are handled via scripts you download from the Device management center. The center has custom scripts for each of these deployment methods:

- local script (up to 10 machines)
- Group policy

- System Center Configuration Manager (version 1610 or later)
- Mobile Device Management/Microsoft Intune
- VDI onboarding scripts for non-persistent machines

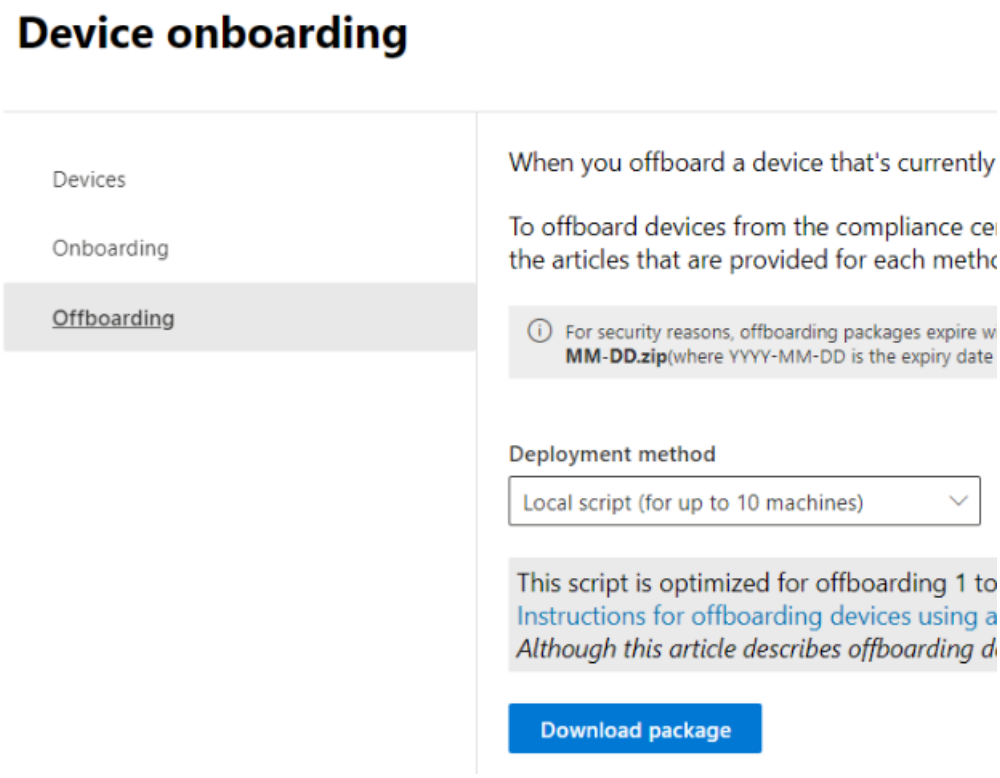

Use the procedures in Getting started with Microsoft 365 Endpoint DLP to onboard devices.

If you have onboarded devices through Microsoft Defender for Endpoint, those devices will automatically show up in the list of devices.

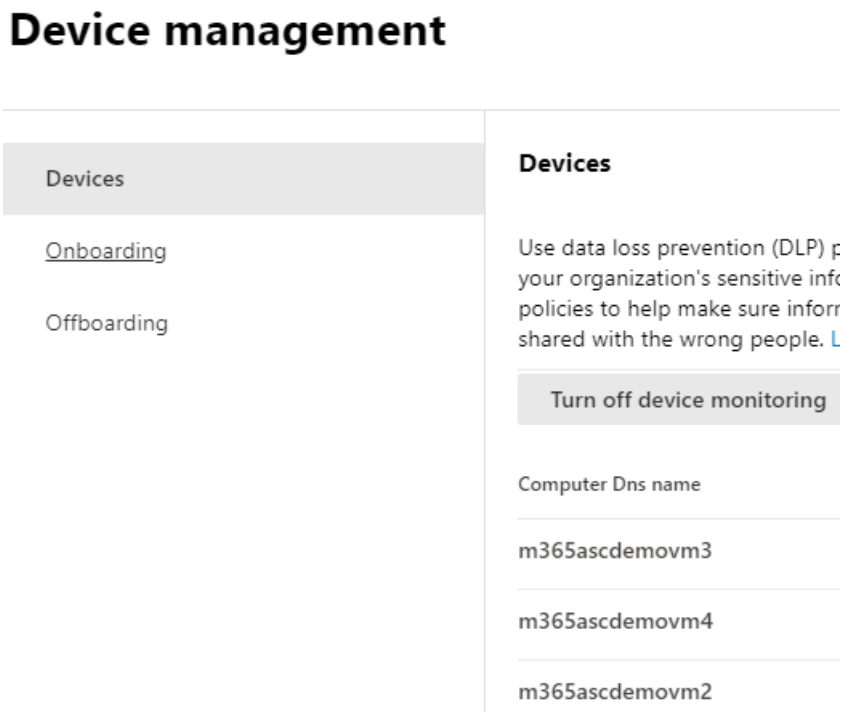

Viewing Endpoint DLP data

You can view alerts related to DLP policies enforced on endpoint devices by going to the DLP Alerts Management Dashboard.

Alert: DLP policy match for document 'CC.Data.docx' on a device

Details Events

Alert information

Alert ID
845cad2a-210e-b185-d400-08d8595cb61a

Alert status
Investigating

Alert severity
■■■ High

Time detected
Sep 15, 2020 3:22 PM

Number of events
1

DLP policy matched
Block CC Data

Locations
Endpoint

You can also view details of the associated event with rich metadata in the same dashboard

Event: Sensitive info in 'CC.Data.docx' - File copied to clipboard

Details | Sensitive info types

Event details

ID
0e653e5a-c175-4887-81bf-0b5cf9eba8d7

Location
Endpoint

Time of activity
Sep 15, 2020 3:19 PM

Impacted entities

User

Hostname

IP address

File
CC.Data.docx

File path
C:\DLP.test.files\CC.Data.docx

Policy details

DLP policy matched
Block CC Data

Rule matched
Rule to stop sharing credit card

Sensitive info types detected
Credit Card Number (1, 85%)

Violating action
File copied to clipboard

Once a device is onboarded, information about audited activities flows into Activity explorer even before you configure and deploy any DLP policies that have devices as a location.

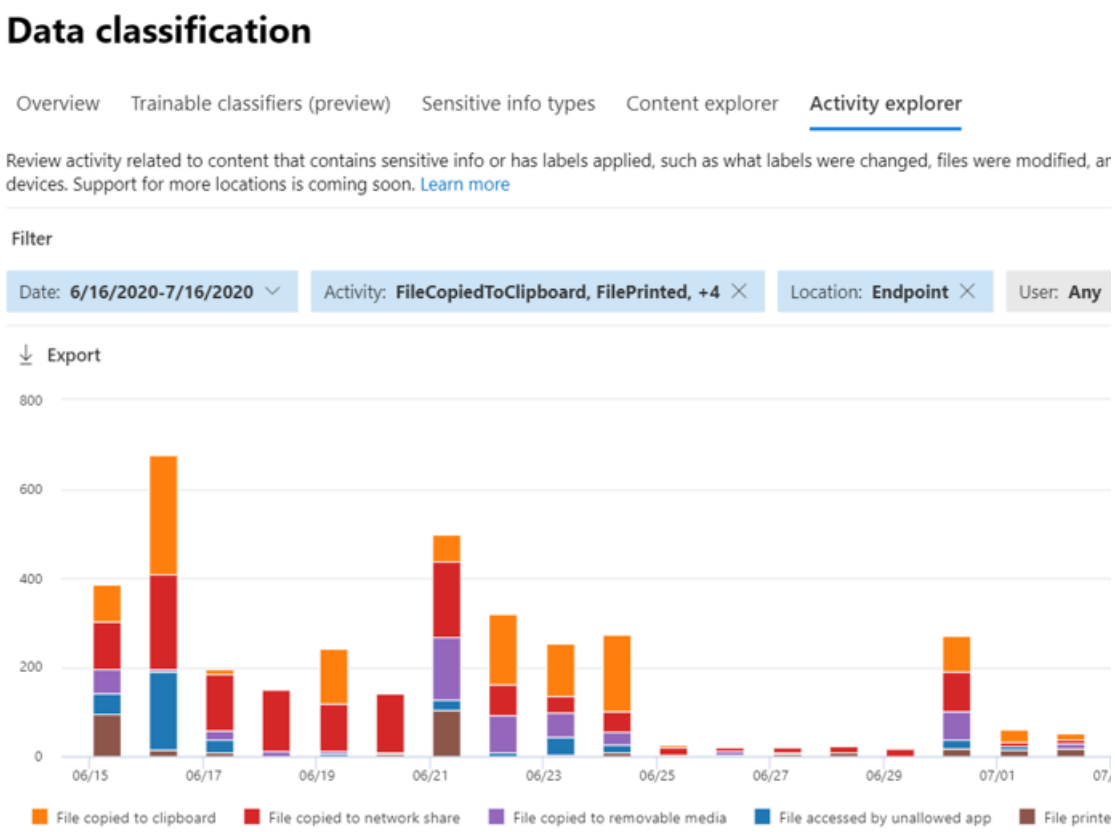

Endpoint DLP collects extensive information on audited activity.

For example, if a file is copied to removable USB media, you'd see these attributes in the activity details:

- activity type
- client IP
- target file path
- happened timestamp
- file name
- user
- file extension
- file size
- sensitive information type (if applicable)
- sha1 value
- sha256 value
- previous file name
- location
- parent
- filepath
- source location type
- platform
- device name
- destination location type
- application that performed the copy
- Microsoft Defender for Endpoint device ID (if applicable)
- removable media device manufacturer
- removable media device model
- removable media device serial number

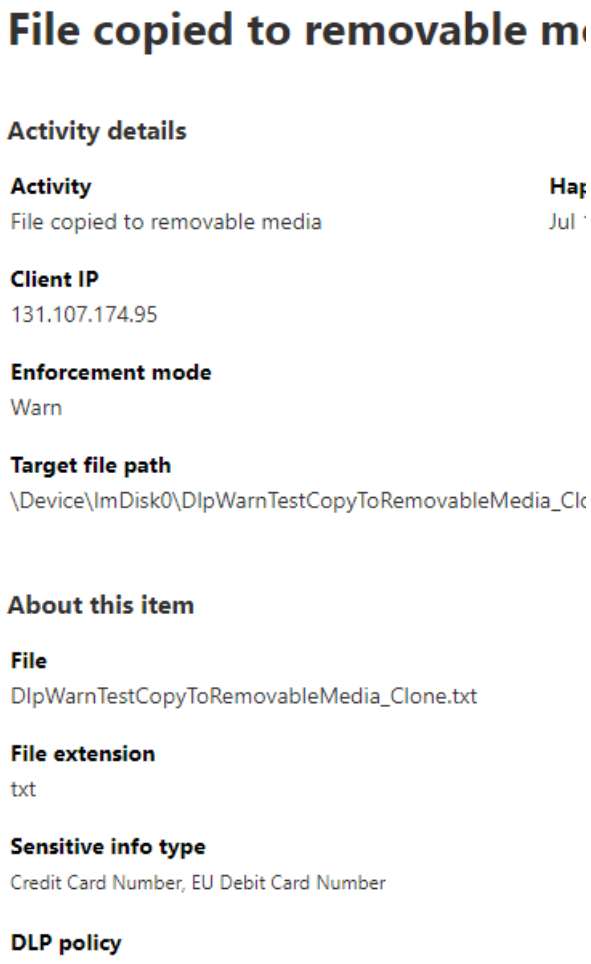

Next steps

Now that you've learned about Endpoint DLP, your next steps are:

1. Getting started with Microsoft Endpoint data loss prevention
2. Using Microsoft Endpoint data loss prevention

See also

- Getting started with Microsoft Endpoint data loss prevention
- Using Microsoft Endpoint data loss prevention
- Overview of data loss prevention
- Create, test, and tune a DLP policy
- Get started with Activity explorer
- Microsoft Defender for Endpoint
- Insider Risk management

Get started with the default DLP policy

11/2/2020 • 3 minutes to read • Edit Online

Before you even create your first data loss prevention (DLP) policy, DLP is helping to protect your sensitive information with a default policy. This default policy and its recommendation (shown below) help keep your sensitive content secure by notifying you when email or documents containing a credit card number were shared with someone outside your organization. You'll see this recommendation on the **Home** page of the Security & Compliance Center.

You can use this widget to quickly view when and how much sensitive information was shared, and then refine the default DLP policy in just a click or two. You can also edit the default DLP policy at any time because it's fully customizable. Note that if you don't see the recommendation at first, try clicking **+More** at the bottom of the **Recommended for you** section.

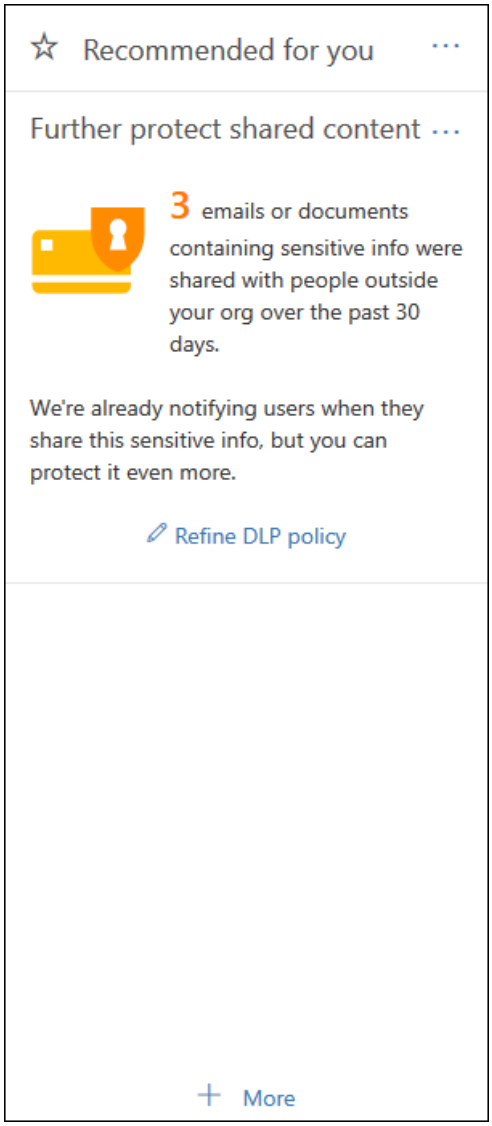

View the report and refine the default DLP policy

When the widget shows you that users have shared sensitive information with people outside your organization, choose **Refine DLP policy** at the bottom.

The detailed report shows you when and how much content containing credit card numbers was shared in the past 30 days. Note that rule matches can take up to 48 hours to show up in the widget.

To help protect the sensitive information, the default DLP policy:

- Detects when content in Exchange, SharePoint, and OneDrive that contains at least one credit card number is shared with people outside your organization.
- Shows a policy tip and sends an email notification to users when they attempt to share this sensitive information with people outside your organization. For more information on these options, see Send email notifications and show policy tips for DLP policies.
- Generates detailed activity reports so that you can track things like who shared the content with people outside your organization and when they did it. You can use the DLP reports and audit log data (where **Activity** = **DLP**) to see this information.

To quickly refine the default DLP policy, you can choose to have it:

- Send you an incident report email when users share this sensitive information with people outside your organization.
- Add other users to the email incident report.
- Block access to the content containing the sensitive information, but allow the user to override and share or send if they need to.

For more information on incident reports or restricting access, see Overview of data loss prevention policies.

If you want to change these options later, you can edit the default DLP policy at any time - see the next section.

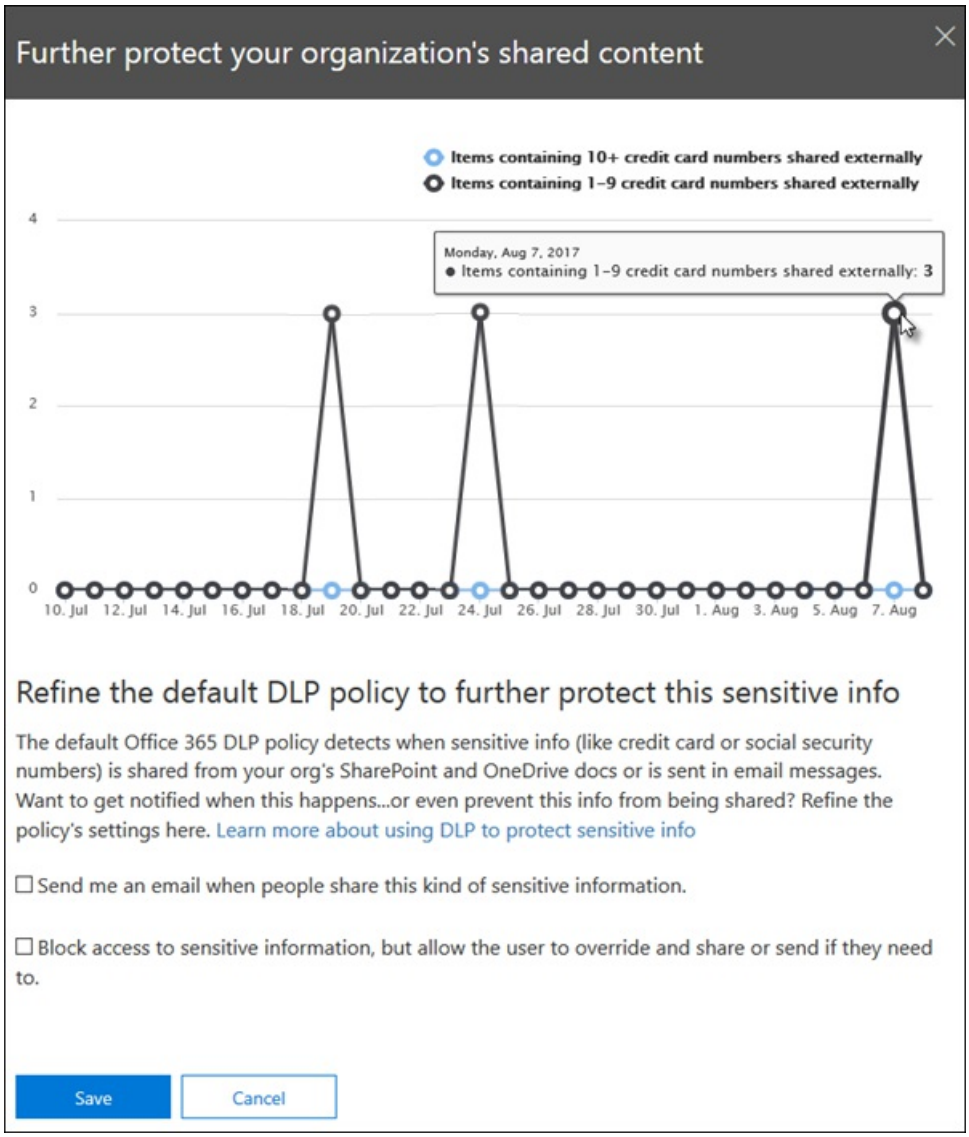

Edit the default DLP policy

This policy is named **Default DLP policy** and appears under **Data loss prevention** on the **Policy** page of the Security & Compliance Center.

This policy is fully customizable, the same as any DLP policy that you create yourself from scratch. You can also turn off or delete the policy, so that your users no longer receive policy tips or email notifications.

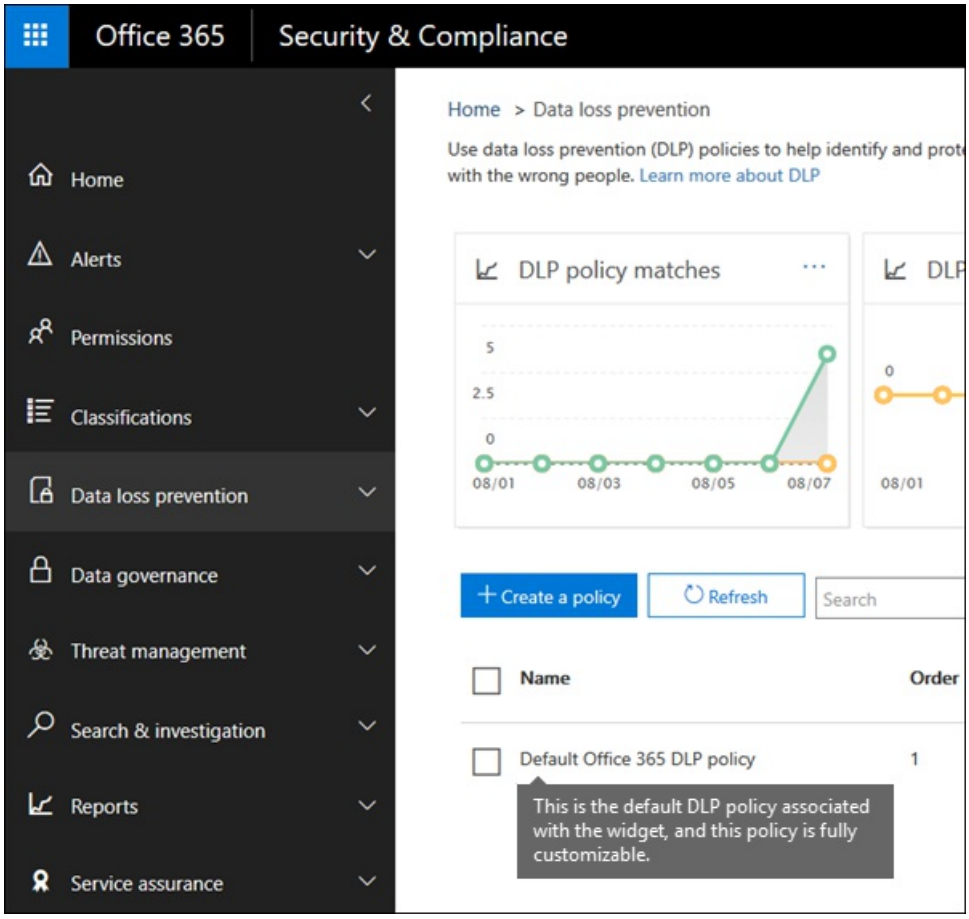

When the widget does and does not appear

The widget named **Further protect shared content** appears in the **Recommended for you** section of the **Home** page of the Security & Compliance Center.

This widget appears only when:

- There are no data loss prevention policies in the Security & Compliance Center or Exchange admin center. This widget is intended to help you get started with DLP, so it doesn't appear if you already have DLP policies.
- Content containing least one credit card has been shared with someone outside your organization in the past 30 days.

Note that rule matches can take up to 48 hours to be available to the widget, so after sensitive information shared externally is detected, it may take up to two days for the recommendation to appear.

Finally, after you use the widget to refine the default DLP policy, the widget disappears from the **Home** page.

Create a DLP policy from a template

5/20/2020 • 10 minutes to read • Edit Online

The easiest, most common way to get started with DLP policies is to use one of the templates included in Office 365. You can use one of these templates as is, or customize the rules to meet your organization's specific compliance requirements.

Microsoft 365 includes over 40 ready-to-use templates that can help you meet a wide range of common regulatory and business policy needs. For example, there are DLP policy templates for:

- Gramm-Leach-Bliley Act (GLBA)
- Payment Card Industry Data Security Standard (PCI-DSS)
- United States Personally Identifiable Information (U.S. PII)
- United States Health Insurance Act (HIPAA)

You can fine tune a template by modifying any of the existing rules or adding new ones. For example, you can add new types of sensitive information to a rule, modify the counts in a rule to make it harder or easier to trigger, allow people to override the actions in a rule by providing a business justification, or change who notifications and incident reports are sent to. A DLP policy template is a flexible starting point for many common compliance scenarios.

You can also choose the Custom template, which has no default rules, and configure your DLP policy from scratch, to meet the specific compliance requirements for your organization.

Example: Identify sensitive information across all OneDrive for Business sites and restrict access for people outside your organization

OneDrive for Business accounts make it easy for people across your organization to collaborate and share documents. But a common concern for compliance officers is that sensitive information stored in OneDrive for Business accounts may be inadvertently shared with people outside your organization. A DLP policy can help mitigate this risk.

In this example, you'll create a DLP policy that identifies U.S. PII data, which includes Individual Taxpayer Identification Numbers (ITIN), Social Security Numbers, and U.S. passport numbers. You'll get started by using a template, and then you'll modify the template to meet your organization's compliance requirements—specifically, you'll:

- Add a couple of types of sensitive information—U.S. bank account numbers and U.S. driver's license numbers—so that the DLP policy protects even more of your sensitive data.
- Make the policy more sensitive, so that a single occurrence of sensitive information is enough to restrict access for external users.
- Allow users to override the actions by providing a business justification or reporting a false positive. This way, your DLP policy won't prevent people in your organization from getting their work done, provided they have a valid business reason for sharing the sensitive information.

Create a DLP policy from a template

1. Go to https://protection.office.com.
2. Sign in using your work or school account. You're now in the Security & Compliance Center.

3. In the Security & Compliance Center > left navigation > **Data loss prevention** > **Policy** > **+ Create a policy**.

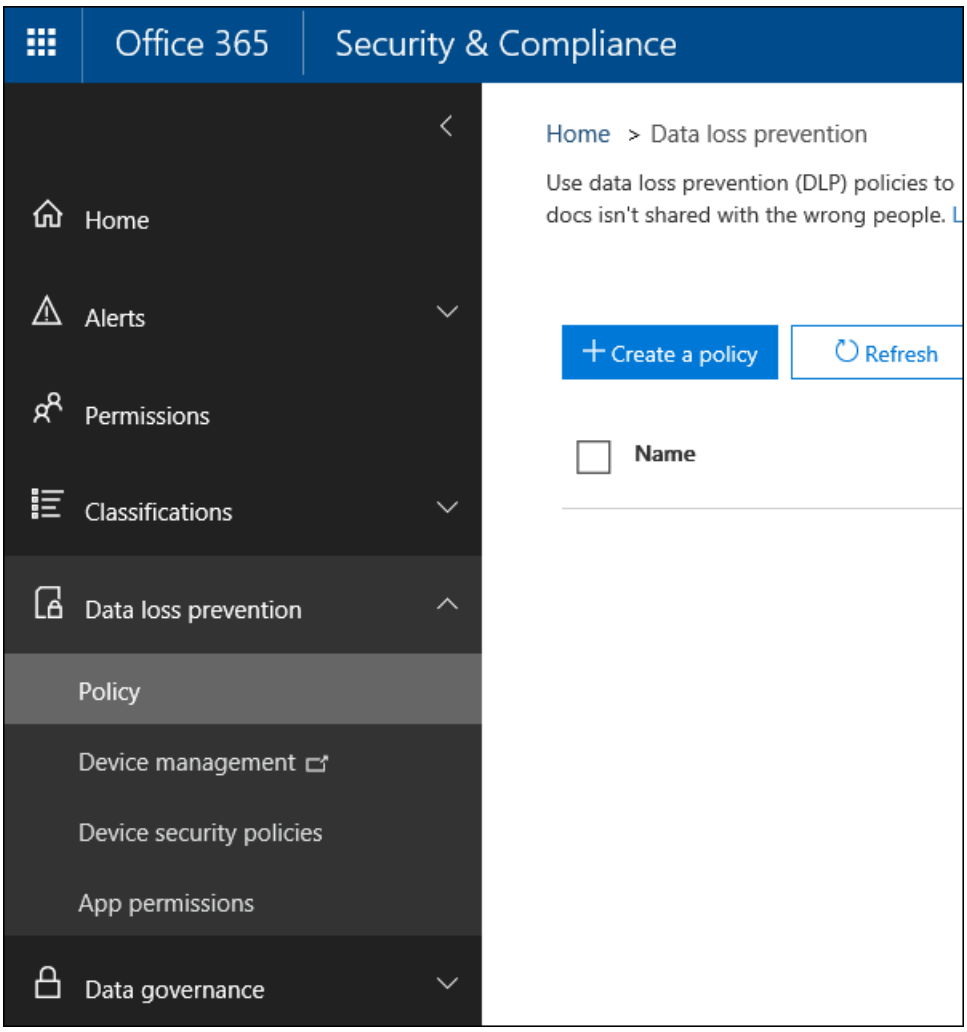

4. Choose the DLP policy template that protects the types of sensitive information that you need > **Next**.

   In this example, you'll select **Privacy** > **U.S. Personally Identifiable Information (PII) Data** because it already includes most of the types of sensitive information that you want to protect—you'll add a couple later.

   When you select a template, you can read the description on the right to learn what types of sensitive information the template protects.

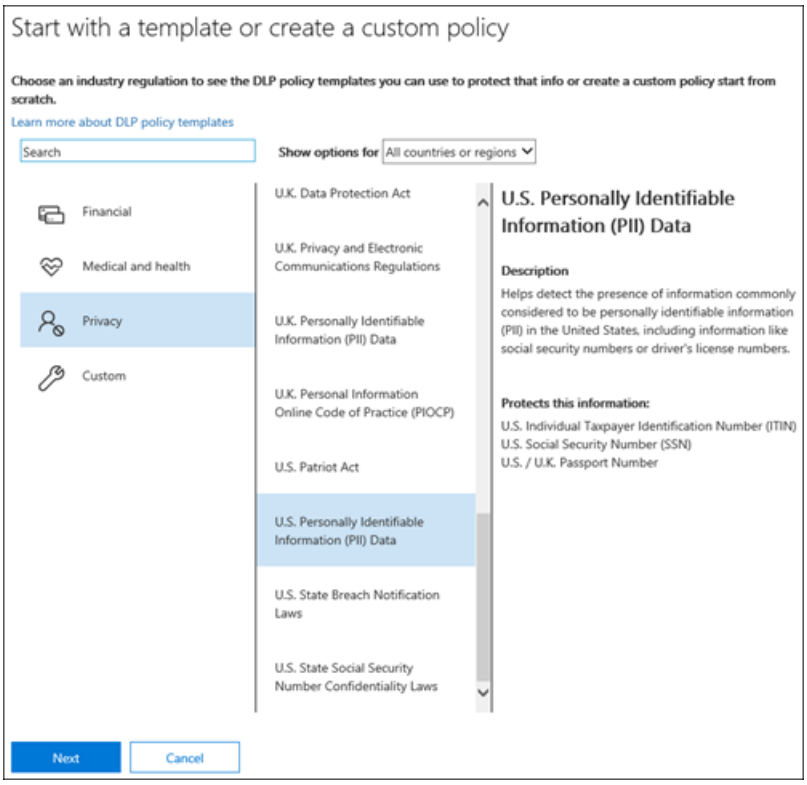

5. Name the policy > **Next**.

6. To choose the locations that you want the DLP policy to protect, do one of the following:

- Choose **All locations in Office 365** > **Next**.
- Choose **Let me choose specific locations** > **Next**. For this example, choose this.

  To include or exclude an entire location such as all Exchange email or all OneDrive accounts, switch the **Status** of that location on or off.

  To include only specific SharePoint sites or OneDrive for Business accounts, switch the **Status** to on, and then click the links under **Include** to choose specific sites or accounts. When you apply a policy to a site, the rules configured in that policy are automatically applied to all subsites of that site.

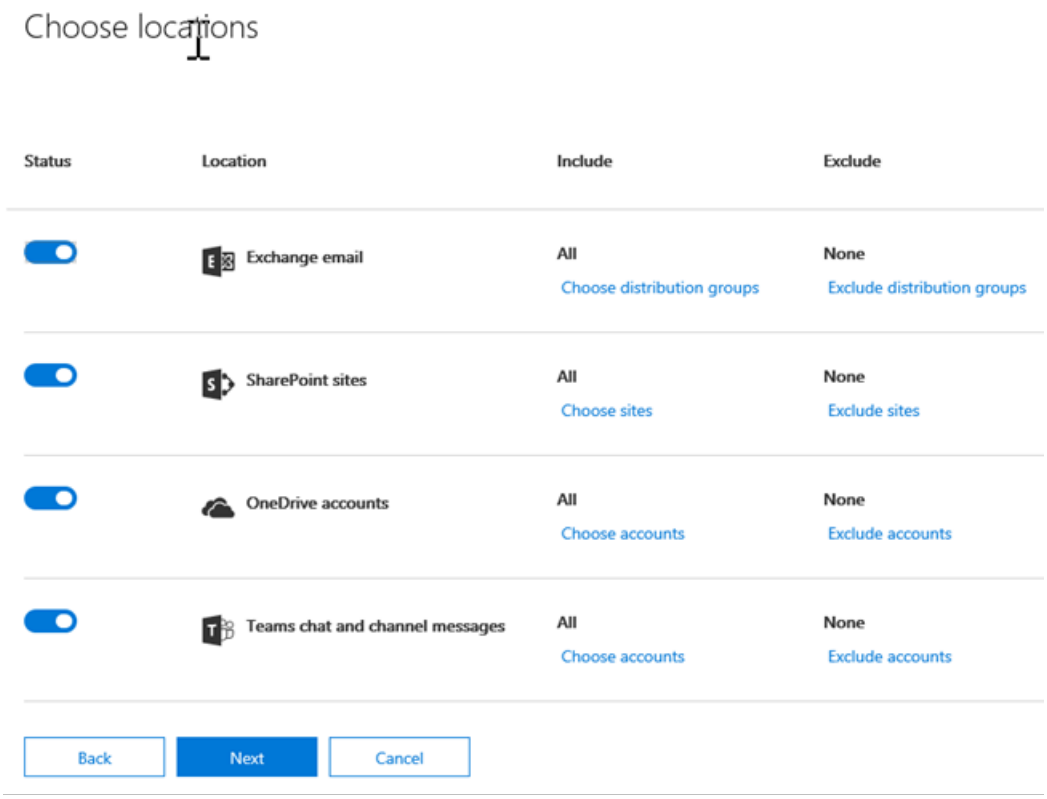

In this example, to protect sensitive information stored in all OneDrive for Business accounts, turn off the **Status** for both **Exchange email** and **SharePoint sites**, and leave the **Status** on for **OneDrive accounts**.

7. Choose **Use advanced settings** > **Next**.
8. A DLP policy template contains predefined rules with conditions and actions that detect and act upon specific types of sensitive information. You can edit, delete, or turn off any of the existing rules, or add new ones. When done, click **Next**.

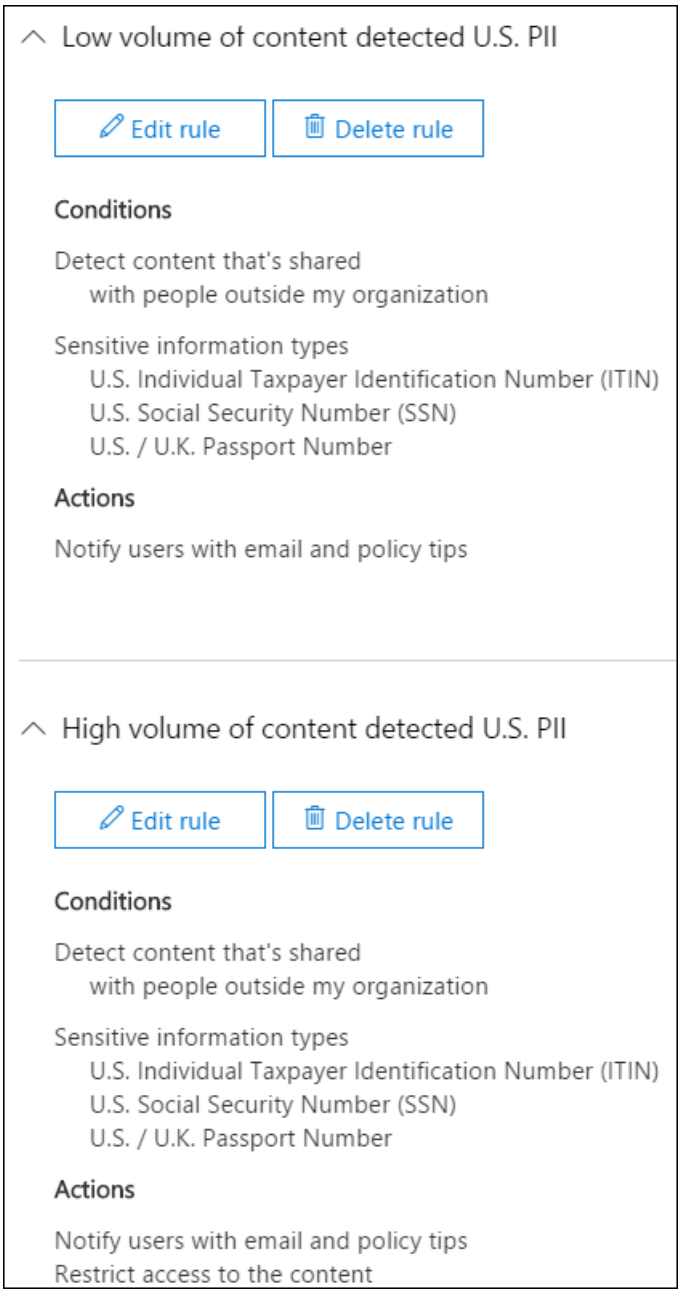

In this example, the U.S. PII Data template includes two predefined rules:

- **Low volume of content detected U.S. PII** This rule looks for files containing between 1 and 10 occurrences of each of three types of sensitive information (ITIN, SSN, and U.S. passport numbers), where the files are shared with people outside the organization. If found, the rule sends an email notification to the primary site collection administrator, document owner, and person who last modified the document.
- **High volume of content detected U.S. PII** This rule looks for files containing 10 or more occurrences of each of the same three sensitive information types, where the files are shared with people outside the organization. If found, this action also sends an email notification, plus it restricts access to the file. For content in a OneDrive for Business account, this means that permissions for the document are restricted for everyone except the primary site collection administrator, document owner, and person who last modified the document.

  To meet your organization's specific requirements, you may want to make the rules easier to trigger, so that a single occurrence of sensitive information is enough to block access for external users. After looking at these rules, you understand that you don't need low and high count rules—you need only a single rule that blocks access if any occurrence of sensitive information is found.

  So you expand the rule named **Low volume of content detected U.S. PII** > **Delete rule**.

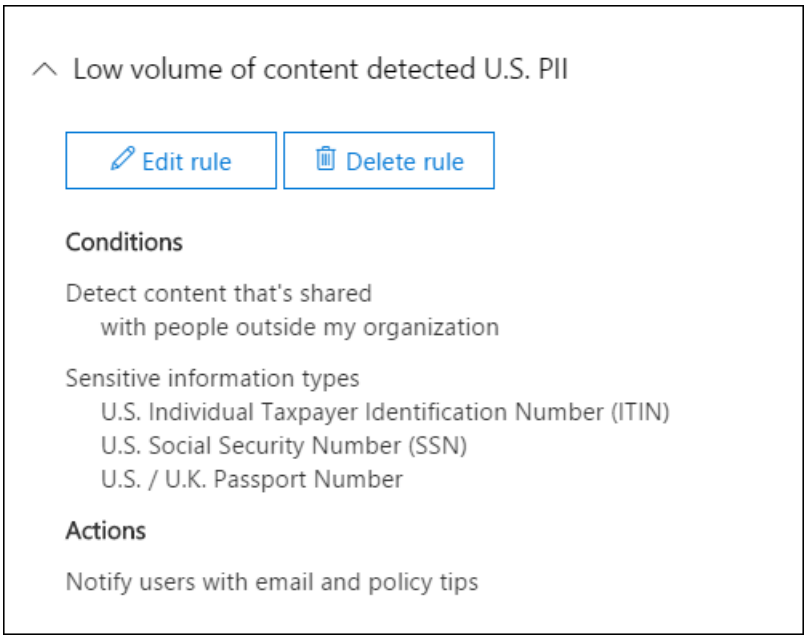

9. Now, in this example, you need to add two sensitive information types (U.S. bank account numbers and U.S. driver's license numbers), allow people to override a rule, and change the count to any occurrence. You can do all of this by editing one rule, so select **High volume of content detected U.S. PII** > **Edit rule**.

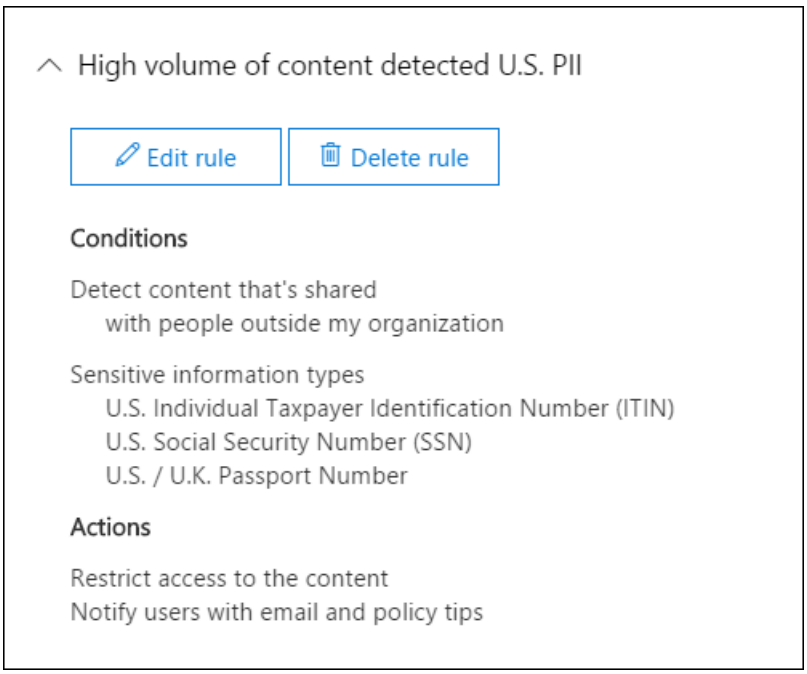

10. To add a sensitive information type, in the **Conditions** section > **Add or change types**. Then, under **Add or change types** > choose **Add** > select **U.S. Bank Account Number** and **U.S. Driver's License Number** > **Add** > **Done**.

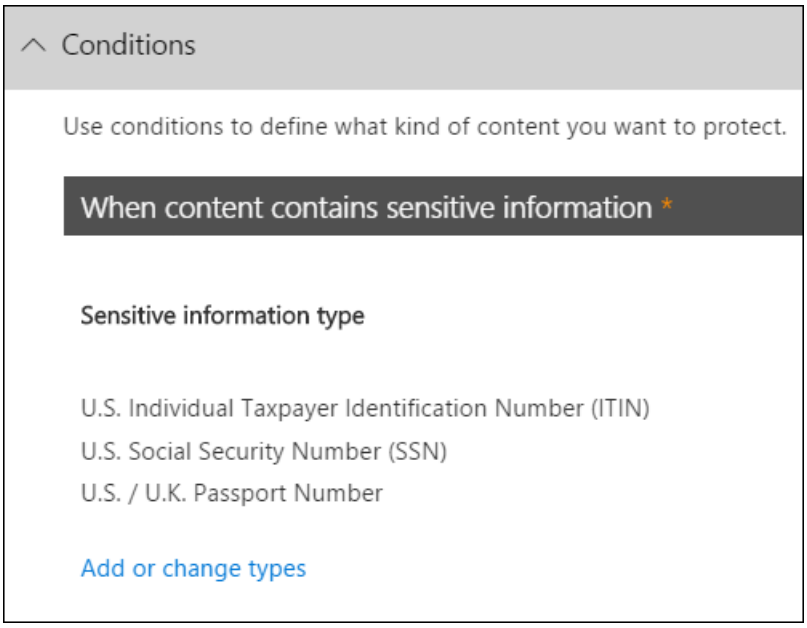

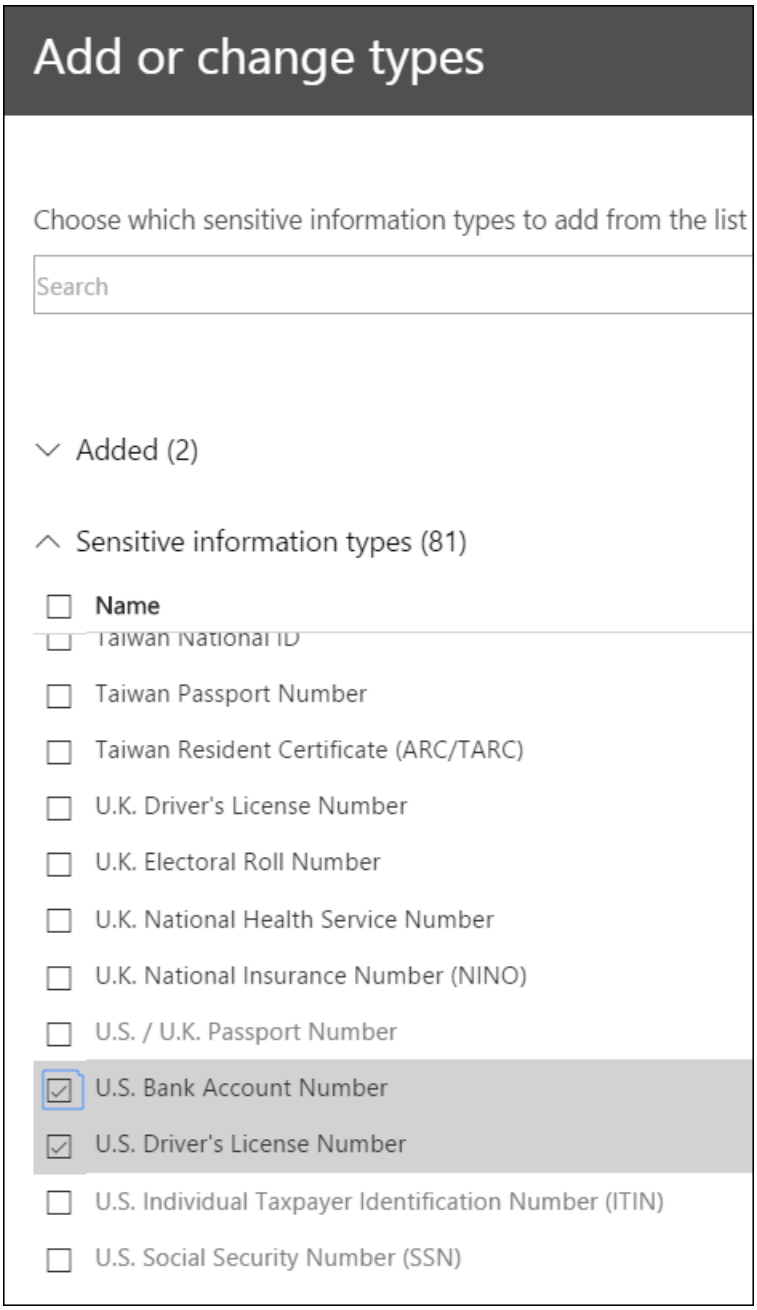

11. To change the count (the number of instances of sensitive information required to trigger the rule), under **Instance count** > choose the **min** value for each type > enter 1. The minimum count cannot be empty.

The maximum count can be empty; an empty **max** value convert to **any**.

When finished, the min count for all of the sensitive information types should be **1** and the max count should be **any**. In other words, any occurrence of this type of sensitive information will satisfy this condition.

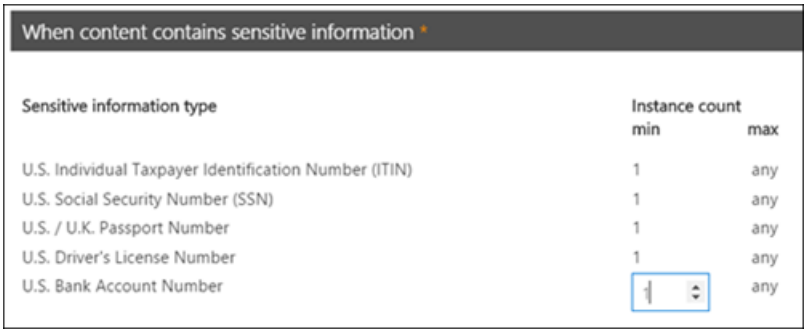

12. For the final customization, you don't want your DLP policies to block people from doing their work when they have a valid business justification or encounter a false positive, so you want the user notification to include options to override the blocking action.

    In the **User notifications** section, you can see that email notifications and policy tips are turned on by default for this rule in the template.

    In the **User overrides** section, you can see that overrides for a business justification are turned on, but overrides to report false positives are not. Choose **Override the rule automatically if they report it as a false positive**.

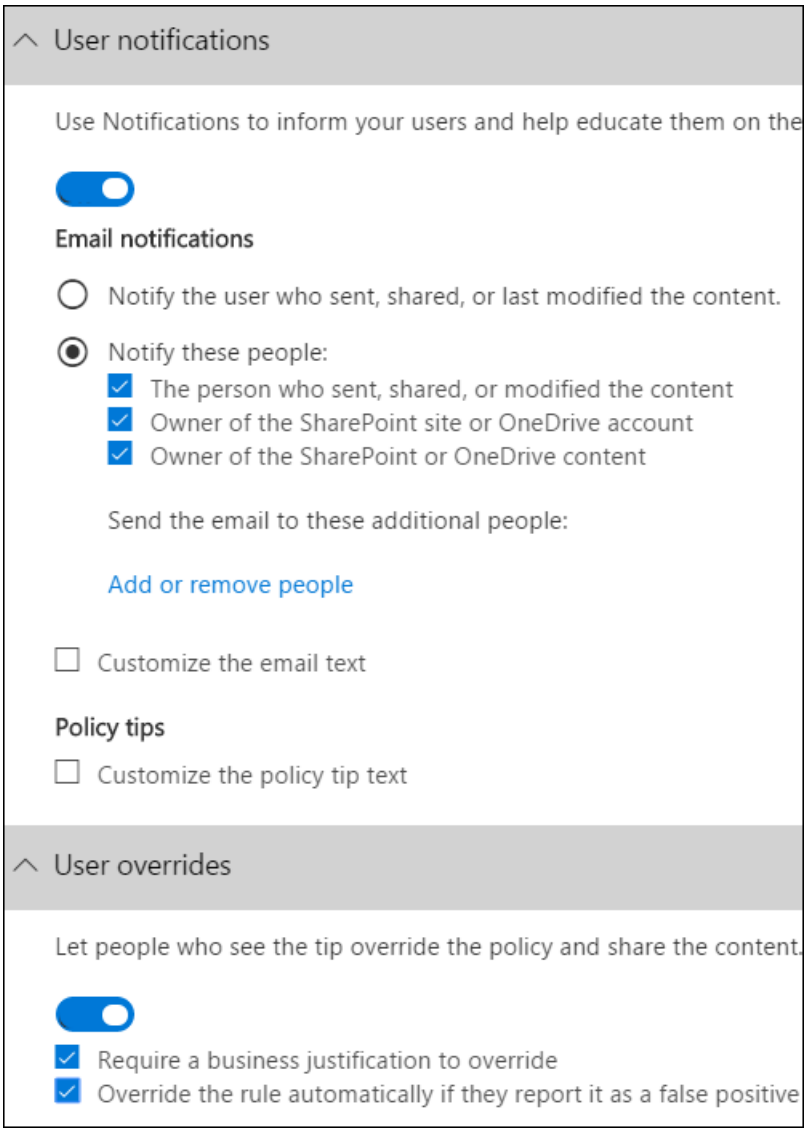

13. At the top of the rule editor, change the name of this rule from the default **High volume of content detected U.S. PII** to **Any content detected with U.S. PII** because it's now triggered by any occurrence of its sensitive information types.

14. At the bottom of the rule editor > **Save**.

15. Review the conditions and actions for this rule > **Next**.

    On the right, notice the **Status** switch for the rule. If you turn off an entire policy, all rules contained in the policy are also turned off. However, here you can turn off a specific rule without turning off the entire policy. This can be useful when you need to investigate a rule that is generating a large number of false positives.

16. On the next page, read and understand the following, and then choose whether to turn on the rule or test it out first > **Next**.

    Before you create your DLP policies, you should consider rolling them out gradually to assess their impact and test their effectiveness before you fully enforce them. For example, you don't want a new DLP policy to unintentionally block access to thousands of documents that people require to get their work done.

    If you're creating DLP policies with a large potential impact, we recommend following this sequence:

17. Start in test mode without Policy Tips and then use the DLP reports to assess the impact. You can use DLP reports to view the number, location, type, and severity of policy matches. Based on the results, you can fine tune the rules as needed. In test mode, DLP policies will not impact the productivity of people working in your organization.

18. Move to Test mode with notifications and Policy Tips so that you can begin to teach users about your compliance policies and prepare them for the rules that are going to be applied. At this stage, you can also ask users to report false positives so that you can further refine the rules.

19. Turn on the policies so that the rules are enforced and the content's protected. Continue to monitor the DLP reports and any incident reports or notifications to make sure that the results are what you intend.

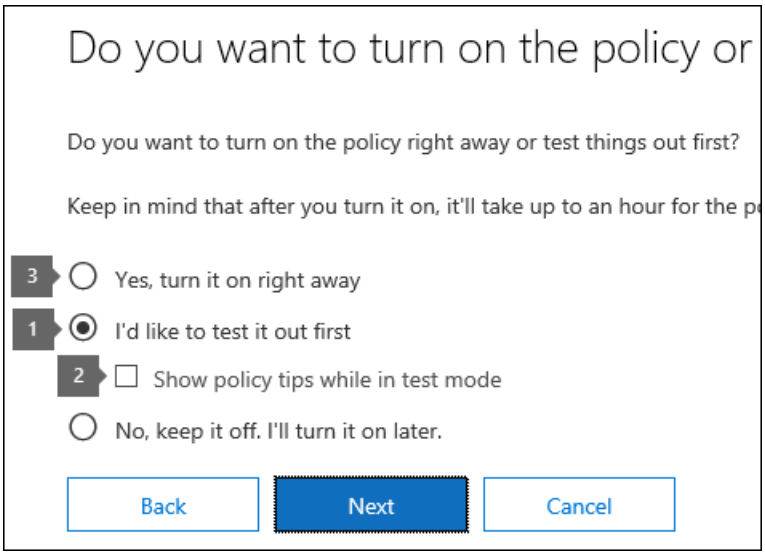

20. Review your settings for this policy > choose **Create**.

After you create and turn on a DLP policy, it's deployed to any content sources that it includes, such as SharePoint Online sites or OneDrive for Business accounts, where the policy begins automatically enforcing its rules on that content.

View the status of a DLP policy

At any time, you can view the status of your DLP policies on the **Policy** page in the **Data loss prevention** section of the Security & Compliance Center. Here you can find important information, such as whether a policy was successfully enabled or disabled, or whether the policy is in test mode.

Here are the different statuses and what they mean.

STATUS	EXPLANATION
Turning on...	The policy is being deployed to the content sources that it includes. The policy is not yet enforced on all sources.
Testing, with notifications	The policy is in test mode. The actions in a rule are not applied, but policy matches are collected and can be viewed by using the DLP reports. Notifications about policy matches are sent to the specified recipients.
Testing, without notifications	The policy is in test mode. The actions in a rule are not applied, but policy matches are collected and can be viewed by using the DLP reports. Notifications about policy matches are not sent to the specified recipients.
On	The policy is active and enforced. The policy was successfully deployed to all its content sources.
Turning off...	The policy is being removed from the content sources that it includes. The policy may still be active and enforced on some sources. Turning off a policy may take up to 45 minutes.
Off	The policy is not active and not enforced. The settings for the policy (sources, keywords, duration, etc) are saved.
Deleting...	The policy is in the process of being deleted. The policy is not active and not enforced. It normally takes an hour for a policy to delet

Turn off a DLP policy

You can edit or turn off a DLP policy at any time. Turning off a policy disables all of the rules in the policy.

To edit or turn off a DLP policy, on the **Policy** page > select the policy > **Edit policy**.

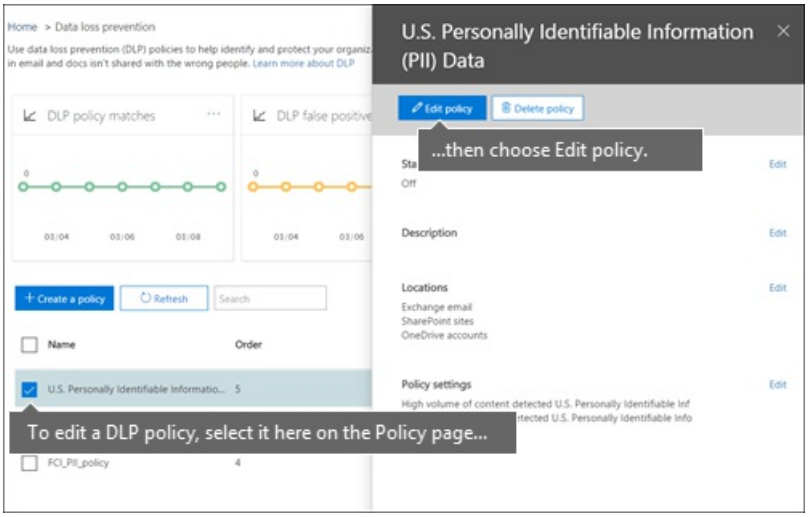

In addition, you can turn off each rule individually by editing the policy and then toggling off the **Status** of that rule, as described above.

More information

- Overview of data loss prevention policies
- Send notifications and show policy tips for DLP policies
- Create a DLP policy to protect documents with FCI or other properties
- What the DLP policy templates include
- Sensitive information type entity definitions

Create, test, and tune a DLP policy

2/18/2021 • 12 minutes to read • Edit Online

Data loss prevention (DLP) helps you prevent the unintentional or accidental sharing of sensitive information.

DLP examines email messages and files for sensitive information, like a credit card number. Using DLP you can detect sensitive information, and take action such as:

- Log the event for auditing purposes
- Display a warning to the end user who is sending the email or sharing the file
- Actively block the email or file sharing from taking place

Permissions

Members of your compliance team who will create DLP policies need permissions to the Compliance Center. By default, your tenant admin will have access can give compliance officers and other people access. Follow these steps:

1. Create a group in Microsoft 365 and add compliance officers to it.
2. Create a role group on the **Permissions** page of the Security & Compliance Center.
3. While creating the role group, use the **Choose Roles** section to add the following role to the role group: **DLP Compliance Management**.
4. Use the **Choose Members** section to add the Microsoft 365 group you created before to the role group.

Use the **View-Only DLP Compliance Management** role to create role group with view-only privileges to the DLP policies and DLP reports.

For more information, see Give users access to the Office 365 Compliance Center.

These permissions are required to create and apply a DLP policy not to enforce policies.

How sensitive information is detected by DLP

DLP finds sensitive information by regular expression (RegEx) pattern matching, in combination with other indicators such as the proximity of certain keywords to the matching patterns. For example, a VISA credit card number has 16 digits. But, those digits can be written in different ways, such as 1111-1111-1111-1111, 1111 1111 1111 1111, or 1111111111111111.

Any 16-digit string is not necessarily a credit card number, it could be a ticket number from a help desk system, or a serial number of a piece of hardware. To tell the difference between a credit card number and a harmless 16-digit string, a calculation is performed (checksum) to confirm that the numbers match a known pattern from the various credit card brands.

If DLP finds keywords such as "VISA" or "AMEX", near date values that might be the credit card expiry date, DLP also uses that data to help it decide whether the string is a credit card number or not.

In other words, DLP is smart enough to recognize the difference between these two strings of text in an email:

- "Can you order me a new laptop. Use my VISA number 1111-1111-1111-1111, expiry 11/22, and send me the estimated delivery date when you have it."
- "My laptop serial number is 2222-2222-2222-2222 and it was purchased on 11/2010. By the way, is my

travel visa approved yet?"

See Sensitive information type entity definitions that explains how each information type is detected.

Where to start with data loss prevention

When the risks of data leakage aren't entirely obvious, it's difficult to work out where exactly you should start with implementing DLP. Fortunately, DLP policies can be run in "test mode", allowing you to gauge their effectiveness and accuracy before you turn them on.

DLP policies for Exchange Online can be managed through the Exchange admin center. But you can configure DLP policies for all workloads through the Security & Compliance Center, so that's what I'll use for demonstrations in this article. In the Security & Compliance Center, you'll find the DLP policies under **Data loss prevention** > **Policy**. Choose **Create a policy** to start.

Microsoft 365 provides a range of DLP policy templates you can use to create policies. Let's say that you're an Australian business. You can filter the templates on Australia, and choose Financial, Medical and Health, and Privacy.

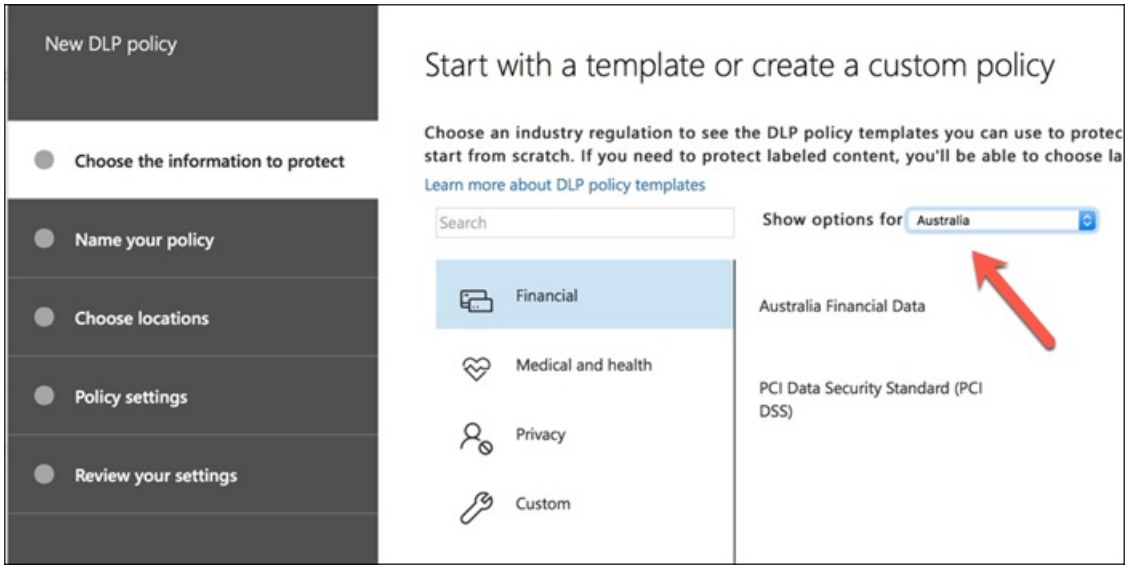

For this demonstration I'll choose Australian Personally Identifiable Information (PII) Data, which includes the information types of Australian Tax File Number (TFN) and Driver's License Number.

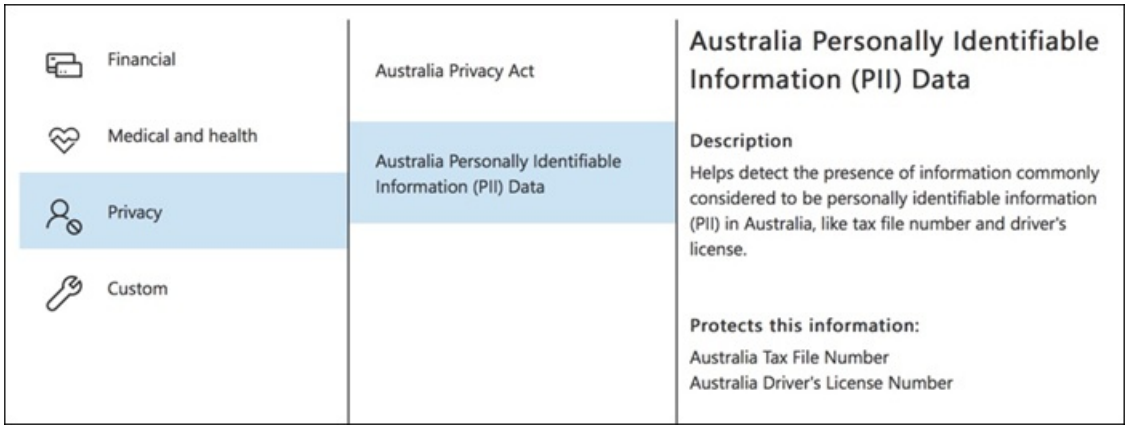

Give your new DLP policy a name. The default name will match the DLP policy template, but you should choose a more descriptive name of your own, because multiple policies can be created from the same template.

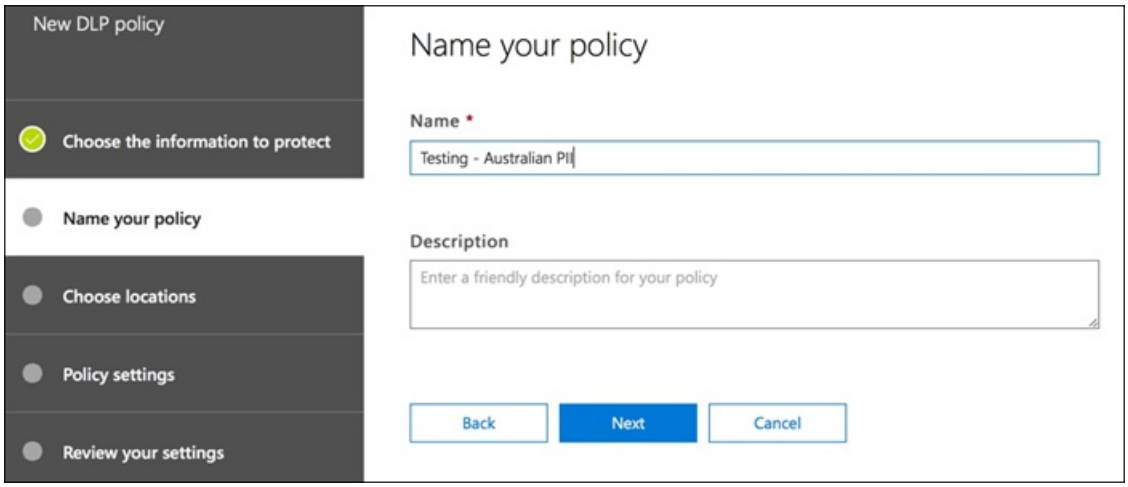

Choose the locations that the policy will apply to. DLP policies can apply to Exchange Online, SharePoint Online, and OneDrive for Business. I am going to leave this policy configured to apply to all locations.

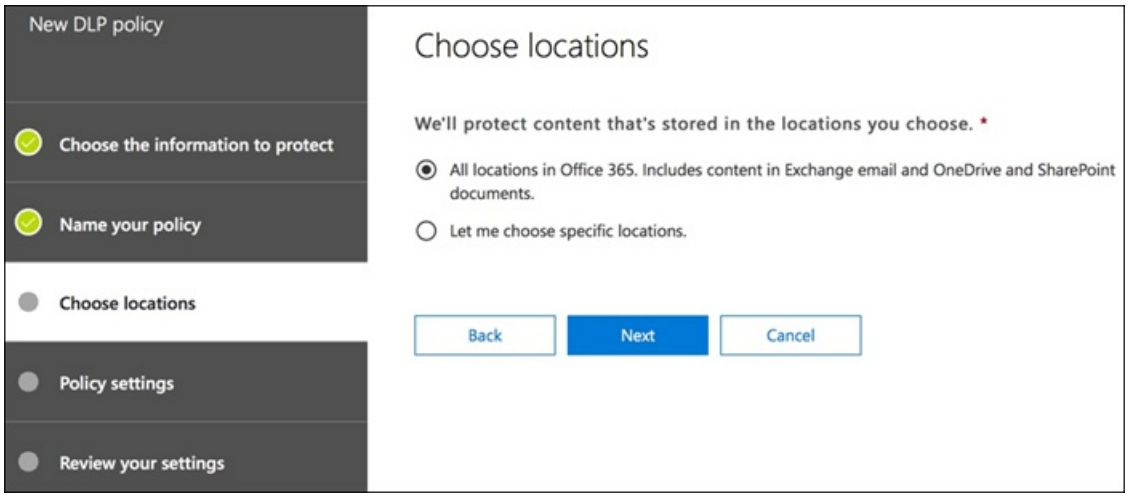

At the first **Policy Settings** step, just accept the defaults for now. You can customize DLP policies, but the defaults are a fine place to start.

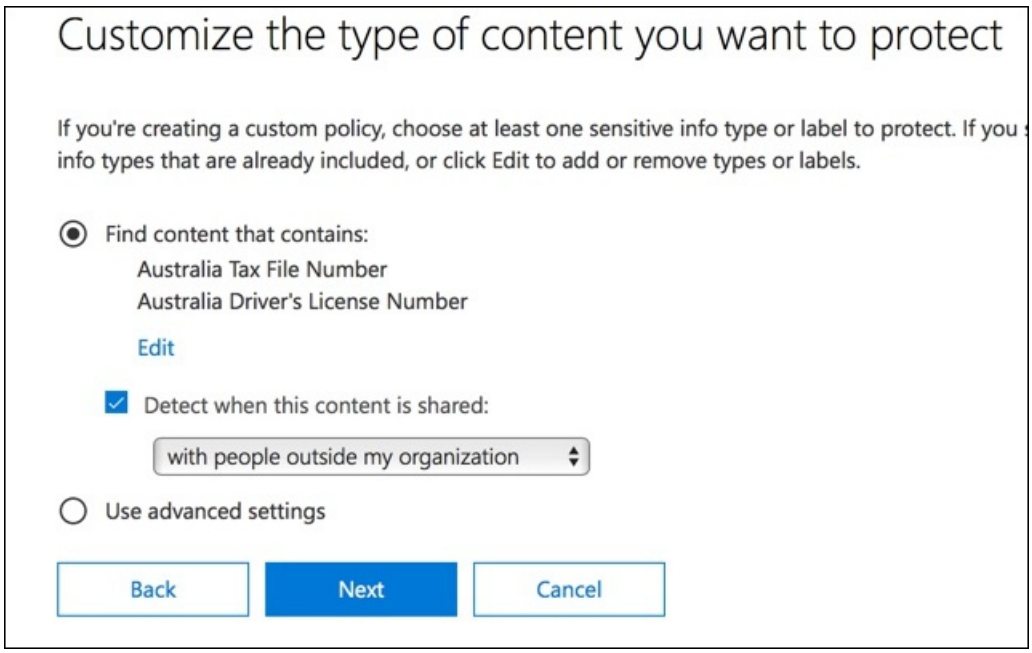

After clicking Next,** you'll be presented with an additional **Policy Settings** page with more customization options. For a policy that you are just testing, here's where you can start to make some adjustments.

- I've turned off policy tips for now, which is a reasonable step to take if you're just testing things out and don't want to display anything to users yet. Policy tips display warnings to users that they're about to violate a DLP

policy. For example, an Outlook user will see a warning that the file they've attached contains credit card numbers and will cause their email to be rejected. The goal of policy tips is to stop the non-compliant behaviour before it happens.

- I've also decreased the number of instances from 10 to 1, so that this policy will detect any sharing of Australian PII data, not just bulk sharing of the data.
- I've also added another recipient to the incident report email.

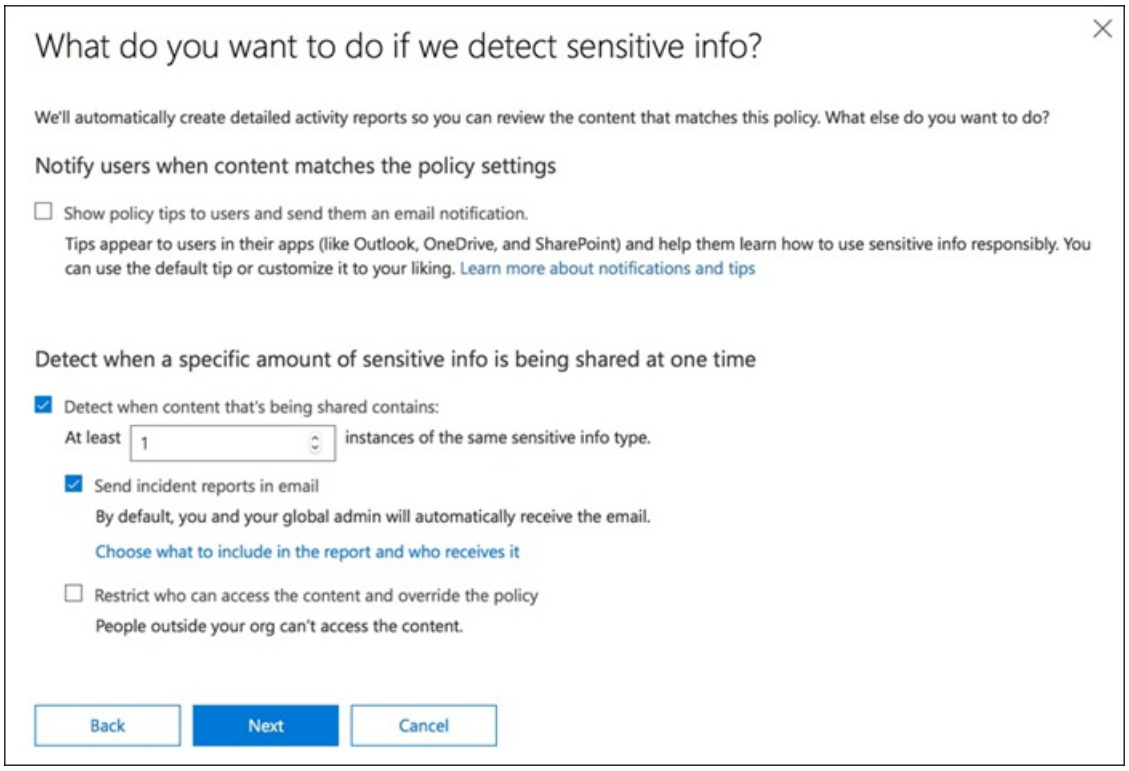

Finally, I've configured this policy to run in test mode initially. Notice there's also an option here to disable policy tips while in test mode. This gives you the flexibility to have policy tips enabled in the policy, but then decide whether to show or suppress them during your testing.

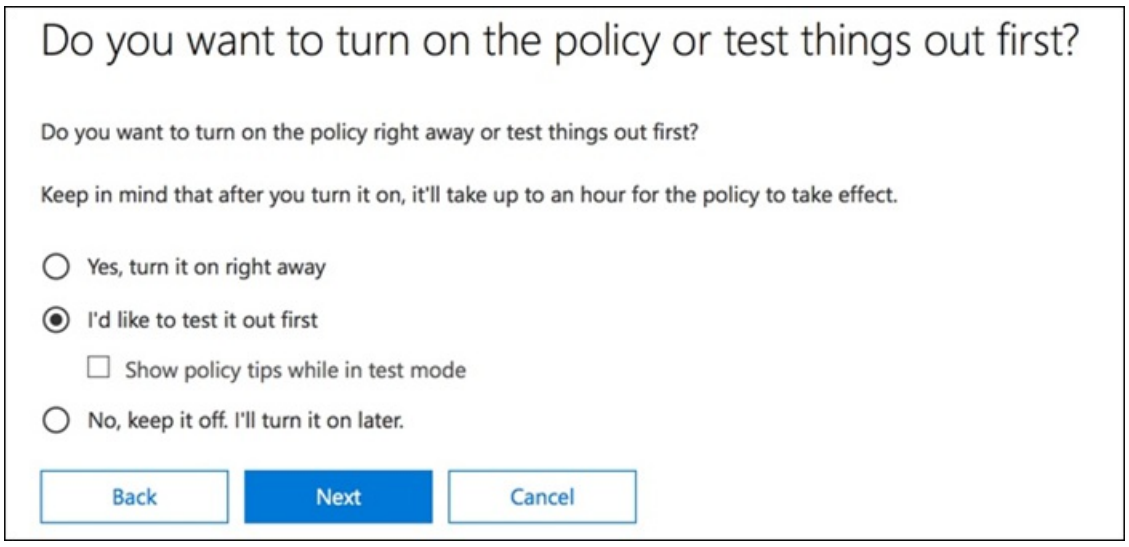

On the final review screen click **Create** to finish creating the policy.

Test a DLP policy

Your new DLP policy will begin to take effect within about 1 hour. You can sit and wait for it to be triggered by normal user activity, or you can try to trigger it yourself. Earlier I linked to Sensitive information type entity definitions, which provides you with information about how to trigger DLP matches.

As an example, the DLP policy I created for this article will detect Australian tax file numbers (TFN). According to the documentation, the match is based on the following criteria.

Australia Tax File Number

Format: 8-9 digits

Pattern: 8-9 digits typically presented with spaces as follows:

- Three digits
- An optional space
- Three digits
- An optional space
- 2-3 digits where the last digit is a check digit

Checksum: Yes

Definition:

A DLP policy is 95% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_australian_tax_file_number` finds content that matches the pattern.
- A keyword from `Keyword_Australia_Tax_File_Number` is found.
- No keyword from `Keyword_number_exclusions` is found.
- The checksum passes.

A DLP policy is 85% confident that it's detected this type of sensitive information if, within a proximity of 300 characters:

- The function `Func_australian_tax_file_number` finds content that matches the pattern.
- No keyword from `Keyword_Australia_Tax_File_Number` or `Keyword_number_exclusions` is found.
- The checksum passes.

```
<!-- Australia Tax File Number -->
<Entity id="e29bc95f-ff70-4a37-aa01-04d17360a4c5" patternsProximity="300" recommendedConfidence="85">
    <Pattern confidenceLevel="95">
        <IdMatch idRef="Func_australian_tax_file_number" />
        <Any minMatches="1">
          <Match idRef="Keyword_Australia_Tax_File_Number" />
        </Any>
        <Any minMatches="0" maxMatches="0">
          <Match idRef="Keyword_number_exclusions" />
        </Any>
  </Pattern>
```

To demonstrate TFN detection in a rather blunt manner, an email with the words "Tax file number" and a 9 digit string in close proximity will sail through without any issues. The reason it does not trigger the DLP policy is that the 9-digit string must pass the checksum that indicates it is a valid TFN and not just a harmless string of numbers.

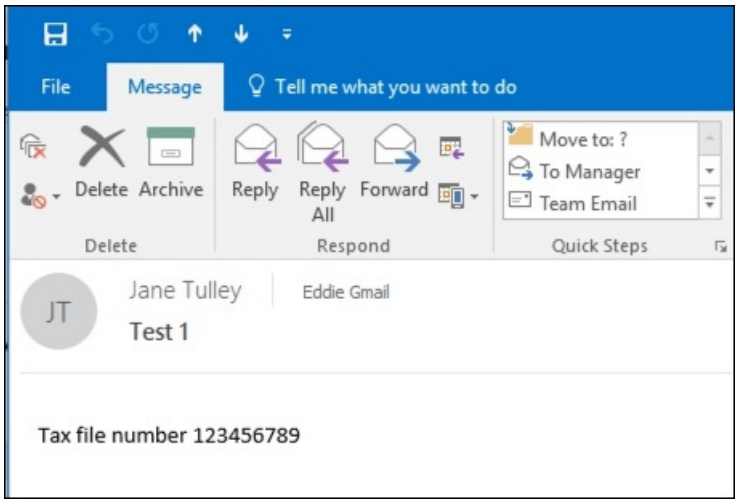

In comparison, an email with the words "Tax file number" and a valid TFN that passes the checksum will trigger the policy. For the record here, the TFN I'm using was taken from a website that generates valid, but not genuine, TFNs. Such sites are very useful because one of the most common mistakes when testing a DLP policy is using a fake number that's not valid and won't pass the checksum (and therefore won't trigger the policy).

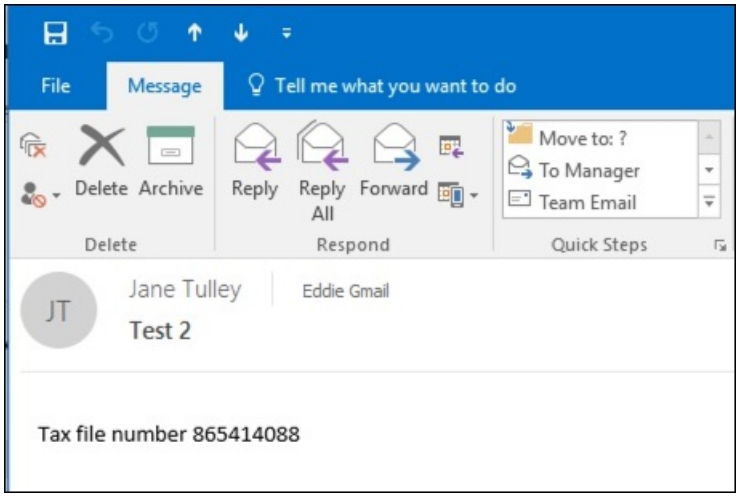

The incident report email includes the type of sensitive information that was detected, how many instances were detected, and the confidence level of the detection.

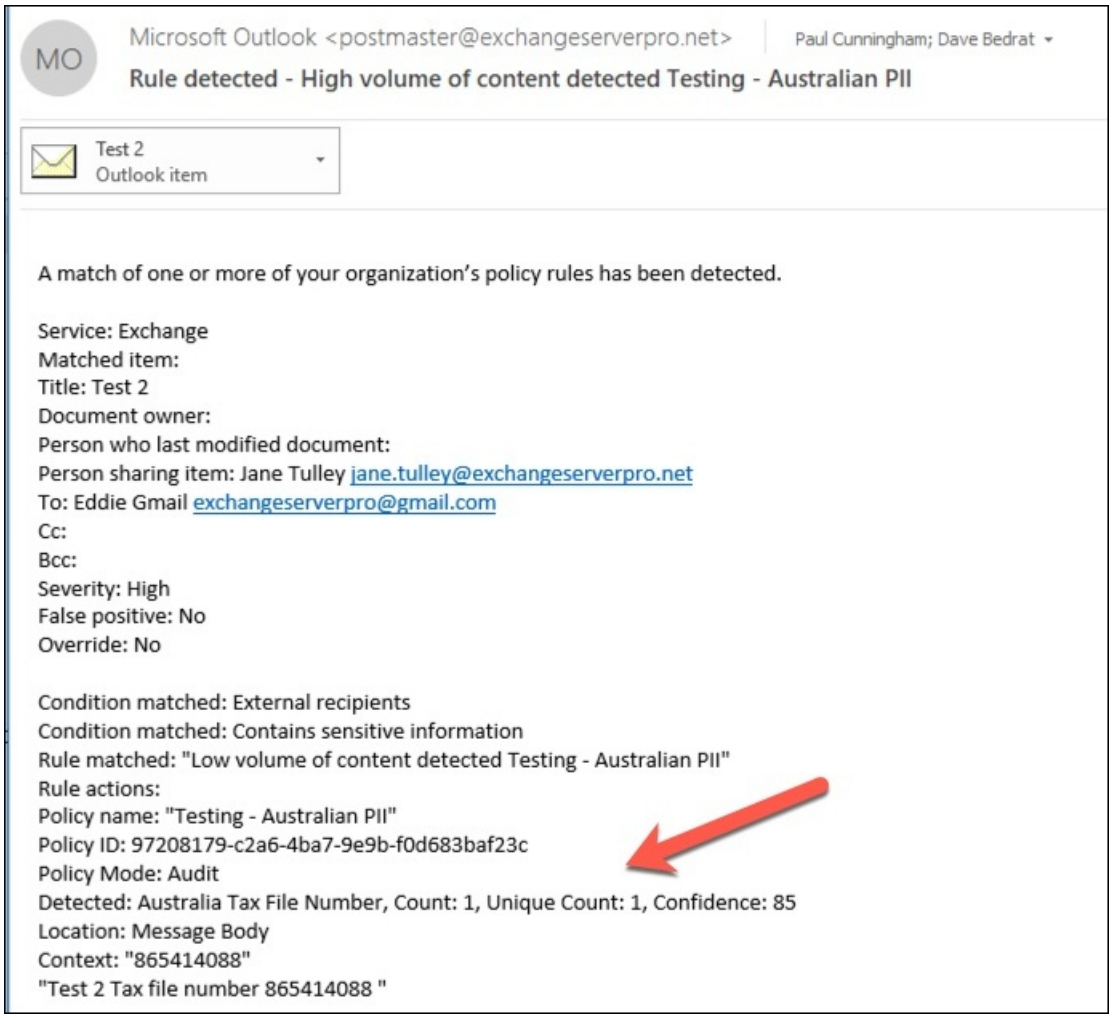

If you leave your DLP policy in test mode and analyze the incident report emails, you can start to get a feel for the accuracy of the DLP policy and how effective it will be when it is enforced. In addition to the incident reports, you can use the DLP reports to see an aggregated view of policy matches across your tenant.

Tune a DLP policy

As you analyze your policy hits you might want to make some adjustments to how the policies behave. As a simple example, you might determine that one TFN in email is not a problem (I think it still is, but let's go with it for the sake of demonstration), but two or more instances is a problem. Multiple instances could be a risky scenario such as an employee emailing a CSV export from the HR database to an external party, for example an external accounting service. Definitely something you would prefer to detect and block.

In the Security & Compliance Center you can edit an existing policy to adjust the behaviour.

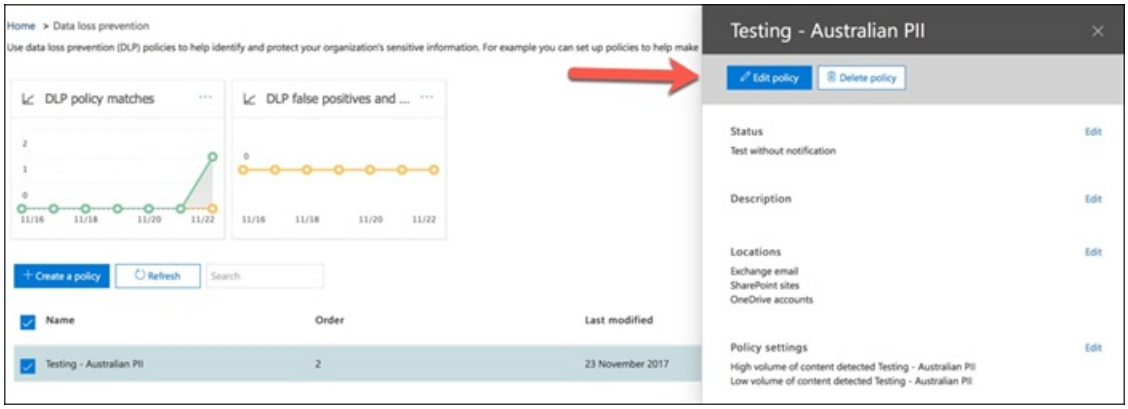

You can adjust the location settings so that the policy is applied only to specific workloads, or to specific sites and accounts.

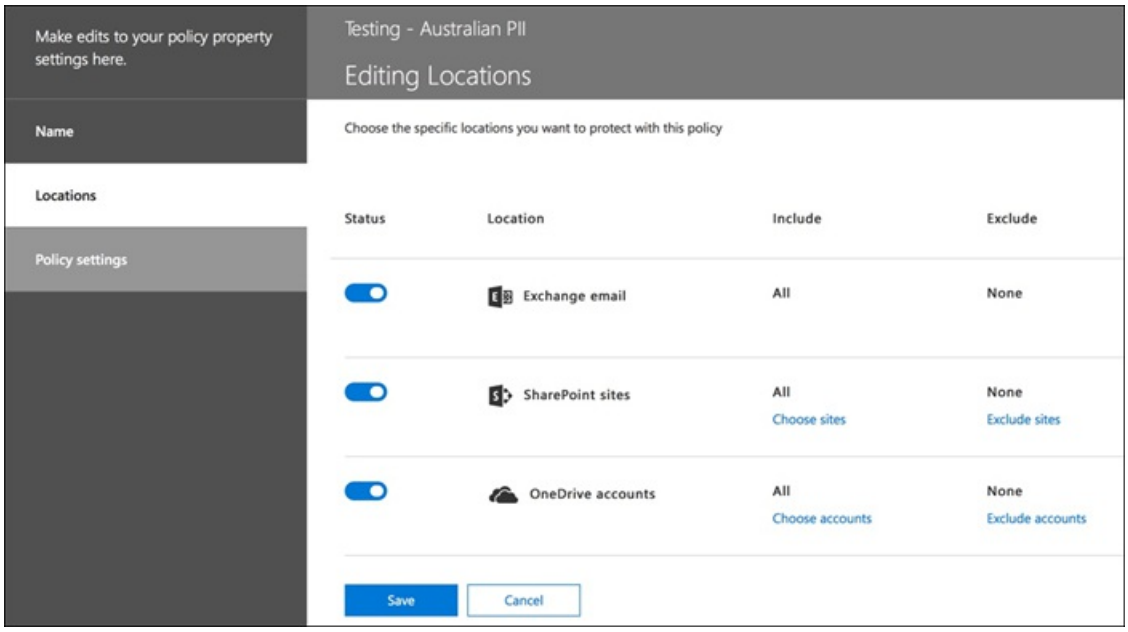

You can also adjust the policy settings and edit the rules to better suit your needs.

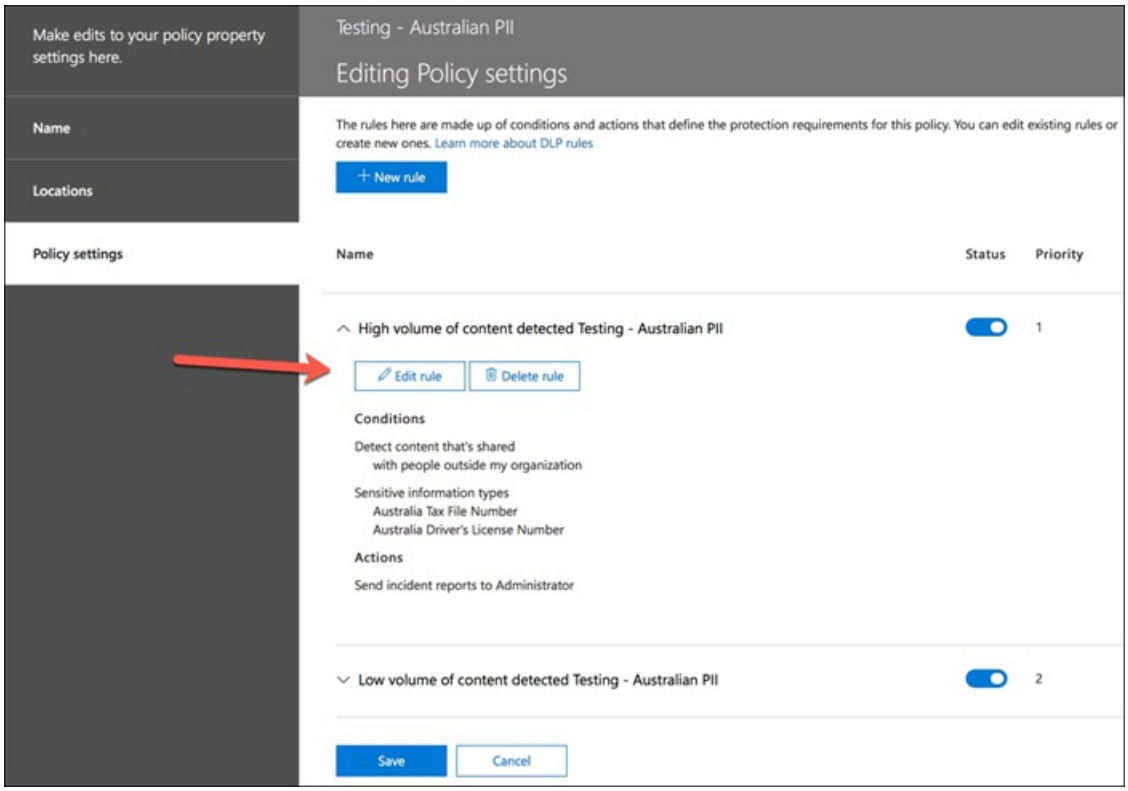

When editing a rule within a DLP policy you can change:

- The conditions, including the type and number of instances of sensitive data that will trigger the rule.
- The actions that are taken, such as restricting access to the content.
- User notifications, which are policy tips that are displayed to the user in their email client or web browser.
- User overrides, which determines whether users can choose to proceed with their email or file sharing anyway.
- Incident reports, to notify administrators.

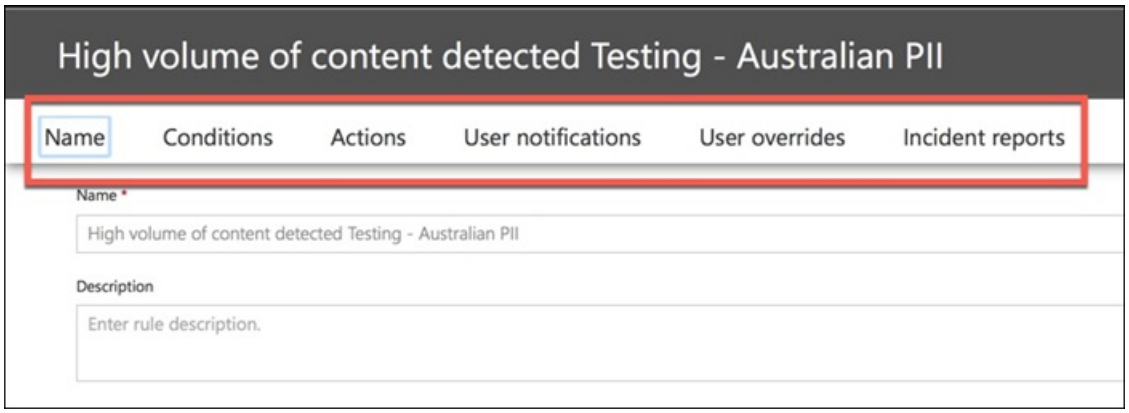

For this demonstration I've added user notifications to the policy (be careful of doing this without adequate user awareness training), and allowed users to override the policy with a business justification or by flagging it as a false positive. Note that you can also customize the email and policy tip text if you want to include any additional information about your organization's policies, or prompt users to contact support if they have questions.

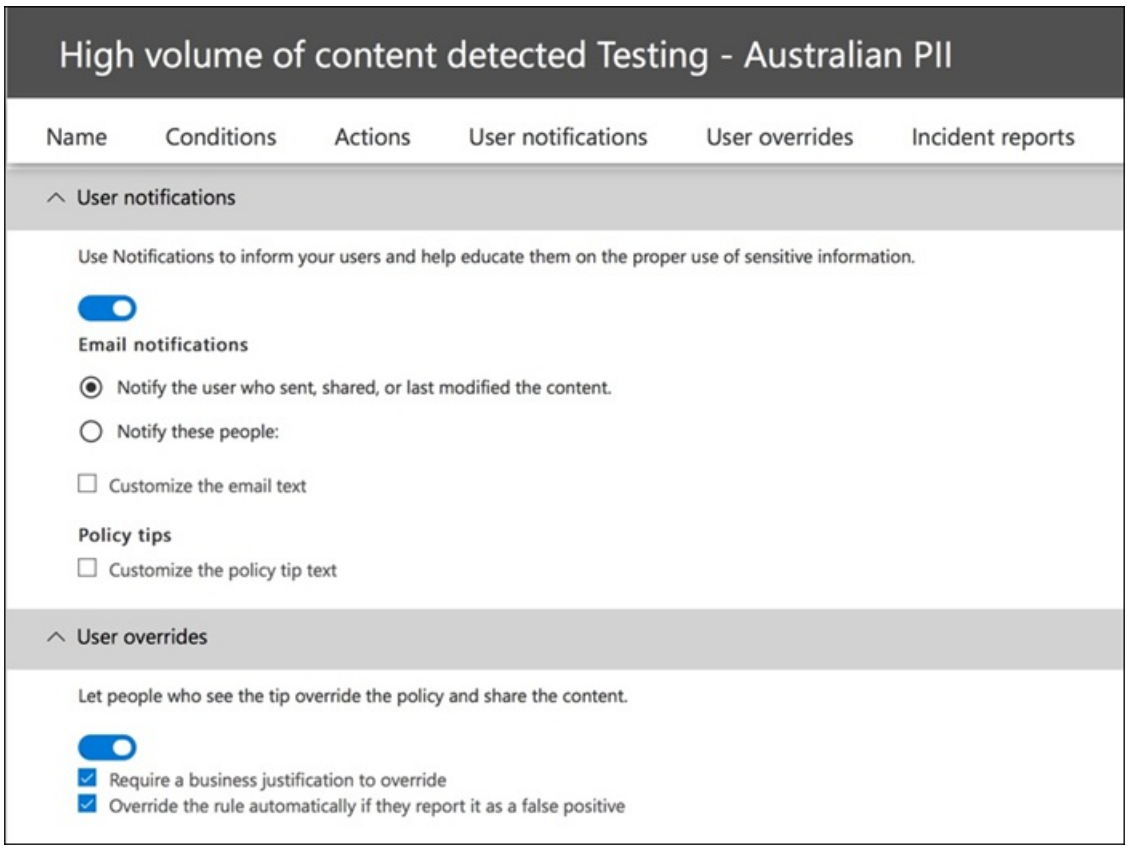

The policy contains two rules for handling of high volume and low volume, so be sure to edit both with the actions that you want. This is an opportunity to treat cases differently depending on their characteristics. For example, you might allow overrides for low volume violations, but not allow overrides for high volume violations.

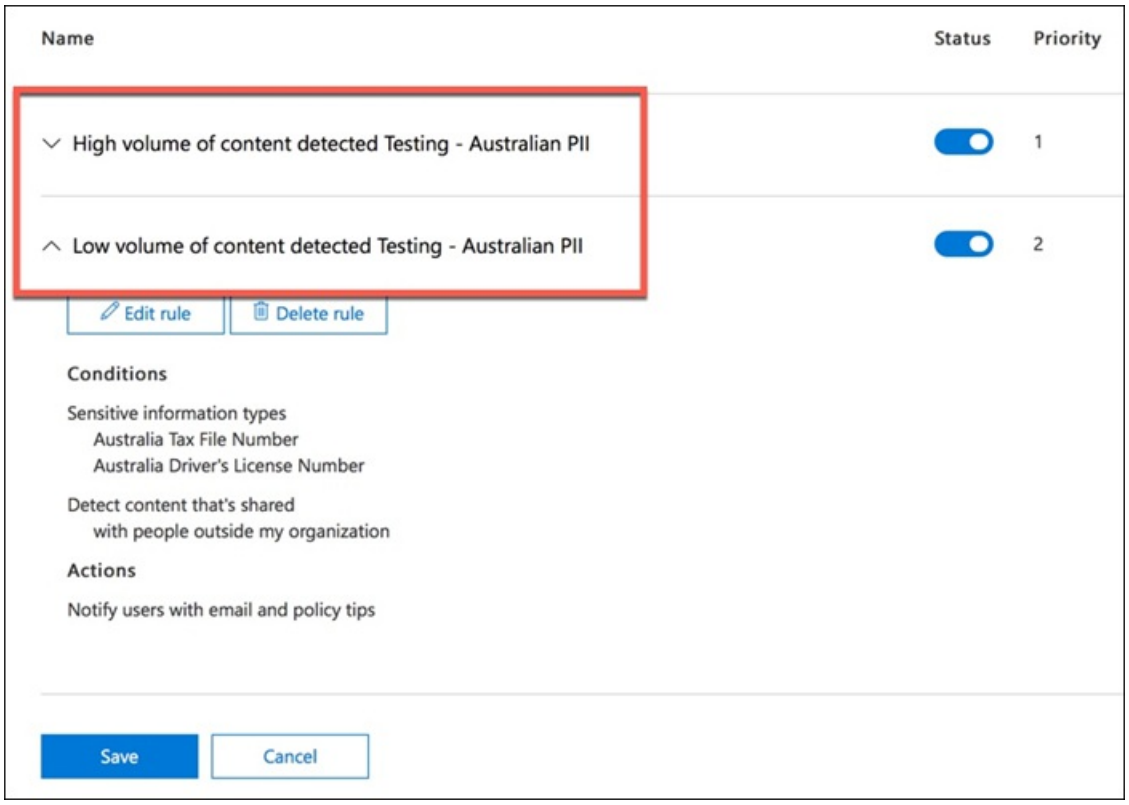

Also, if you want to actually block or restrict access to content that is in violation of policy, you need to configure an action on the rule to do so.

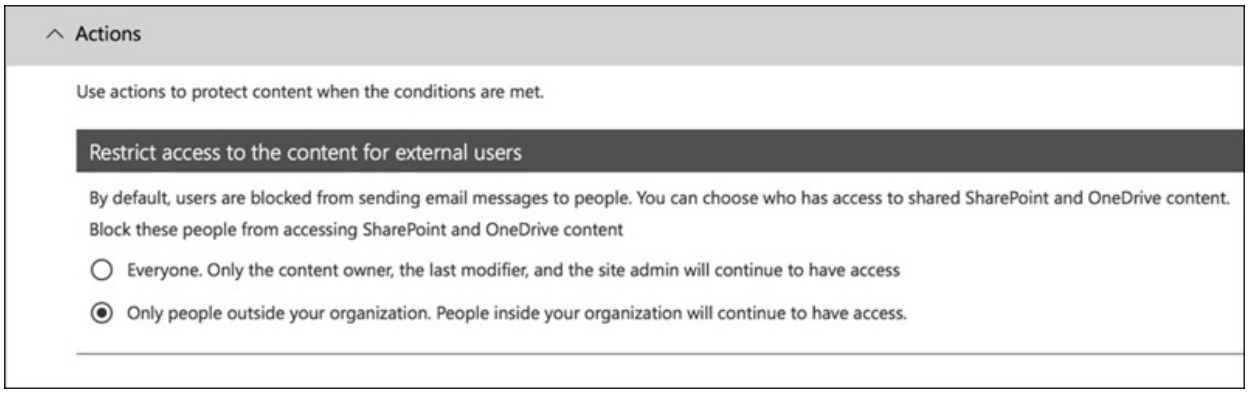

After saving those changes to the policy settings, I also need to return to the main settings page for the policy and enable the option to show policy tips to users while the policy is in test mode. This is an effective way to introduce DLP policies to your end users, and do user awareness training, without risking too many false positives that impact their productivity.

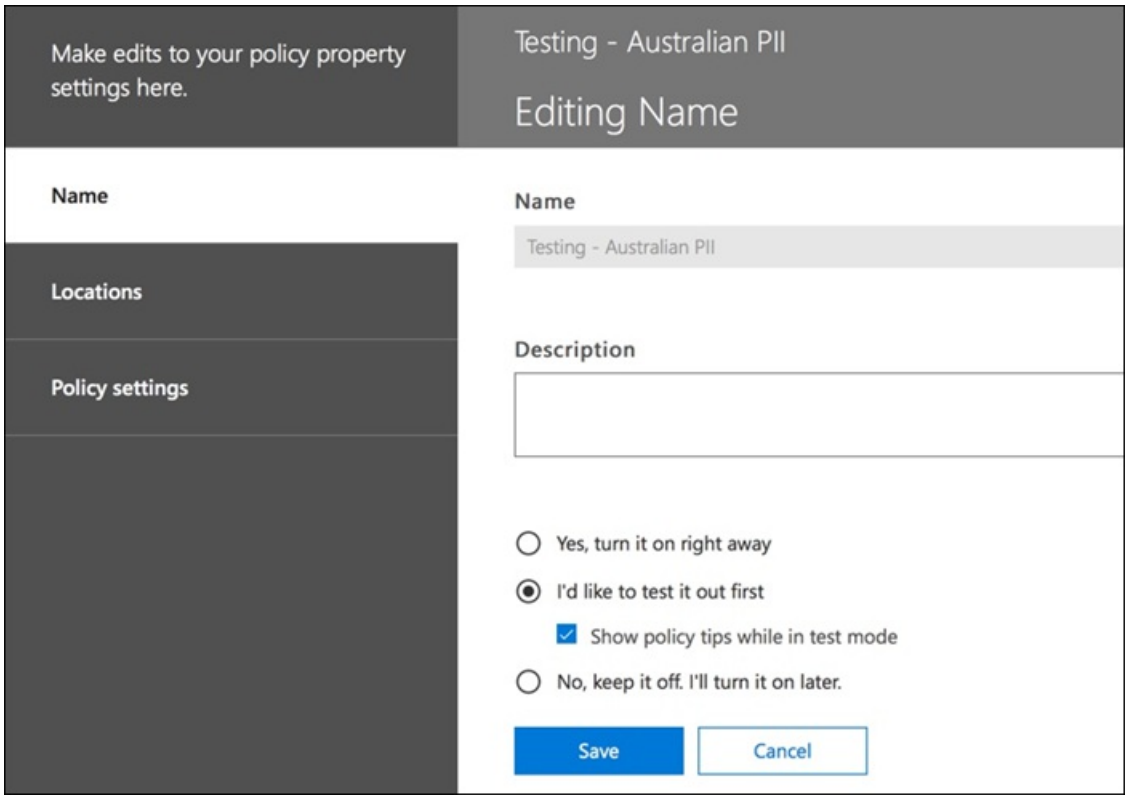

On the server side (or cloud side if you prefer), the change may not take effect immediately, due to various processing intervals. If you're making a DLP policy change that will display new policy tips to a user, the user may not see the changes take effect immediately in their Outlook client, which checks for policy changes every 24 hours. If you want to speed things up for testing, you can use this registry fix to clear the last download time stamp from the PolicyNudges key. Outlook will download the latest policy information the next time you restart it and begin composing an email message.

If you have policy tips enabled, the user will begin to see the tips in Outlook, and can report false positives to you when they occur.

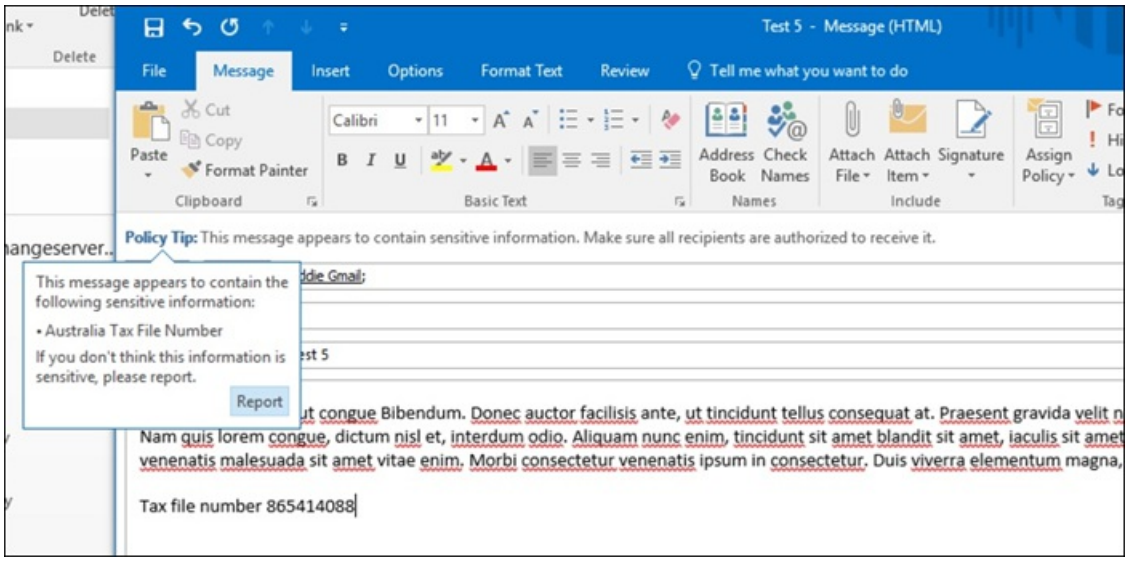

Investigate false positives

DLP policy templates are not perfect straight out of the box. It's likely that you'll find some false positives occurring in your environment, which is why it's so important to ease your way into a DLP deployment, taking the time to adequately test and tune your policies.

Here's an example of a false positive. This email is quite harmless. The user is providing their mobile phone number to someone, and including their email signature.

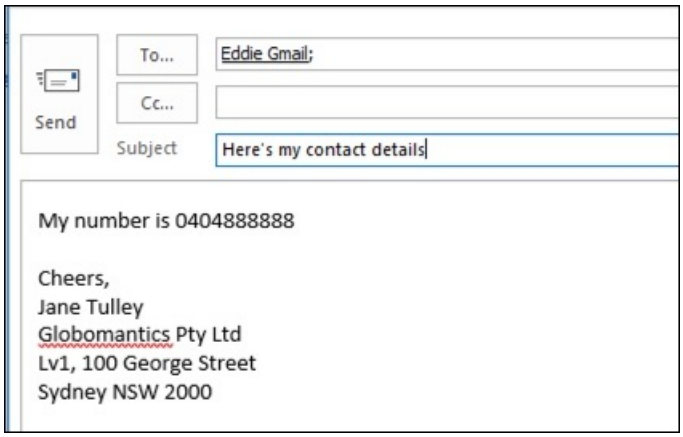

But the user sees a policy tip warning them that the email contains sensitive information, specifically, an Australian driver's license number.

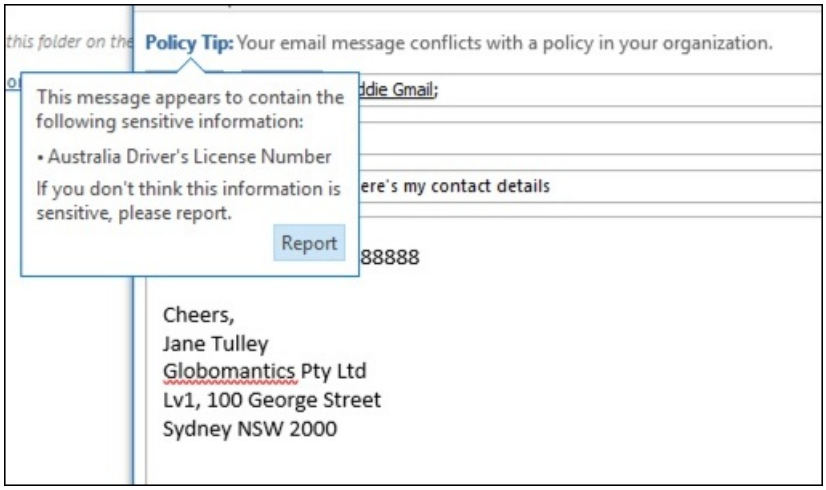

The user can report the false positive, and the administrator can look into why it has occurred. In the incident report email, the email is flagged as a false positive.

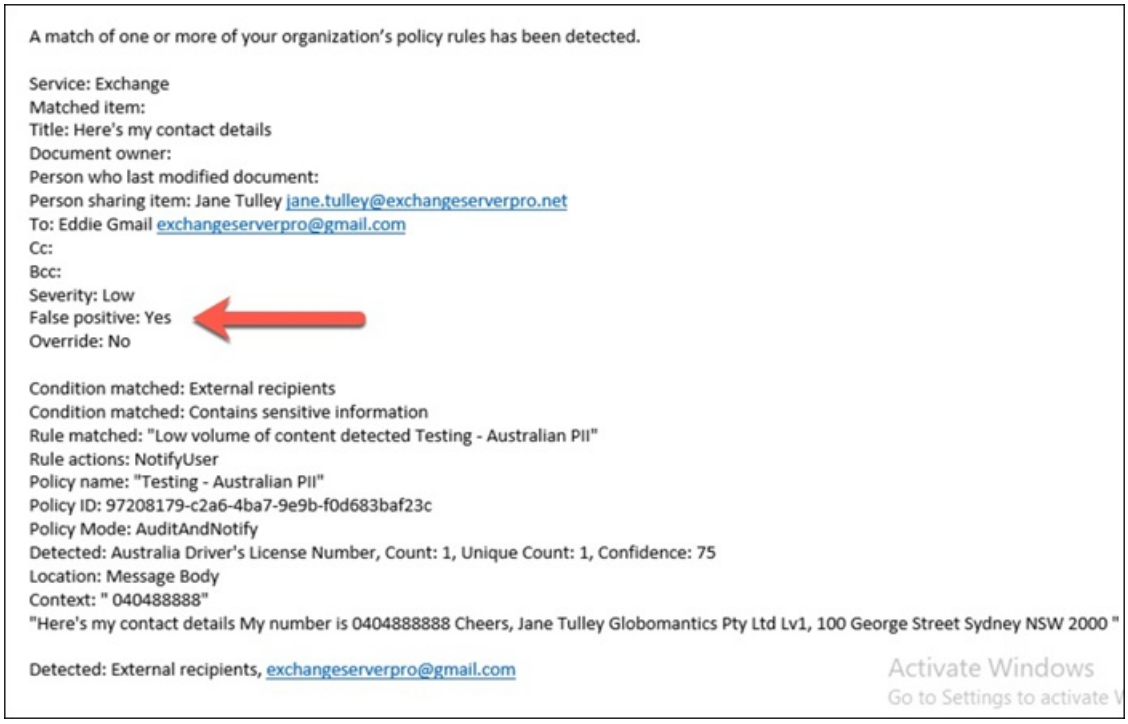

This driver's license case is a good example to dig into. The reason this false positive has occurred is that the "Australian Driver's License" type will be triggered by any 9-digit string (even one that is part of a 10-digit string), within 300 characters proximity to the keywords "sydney nsw" (not case sensitive). So it's triggered by the phone number and email signature, only because the user happens to be in Sydney.

One option is to remove the Australian driver's license information type from the policy. It's in there because it's part of the DLP policy template, but we're not forced to use it. If you're only interested in Tax File Numbers and not driver's licenses, you can just remove it. For example, you can remove it from the low volume rule in the policy, but leave it in the high volume rule so that lists of multiple drivers licenses are still detected.

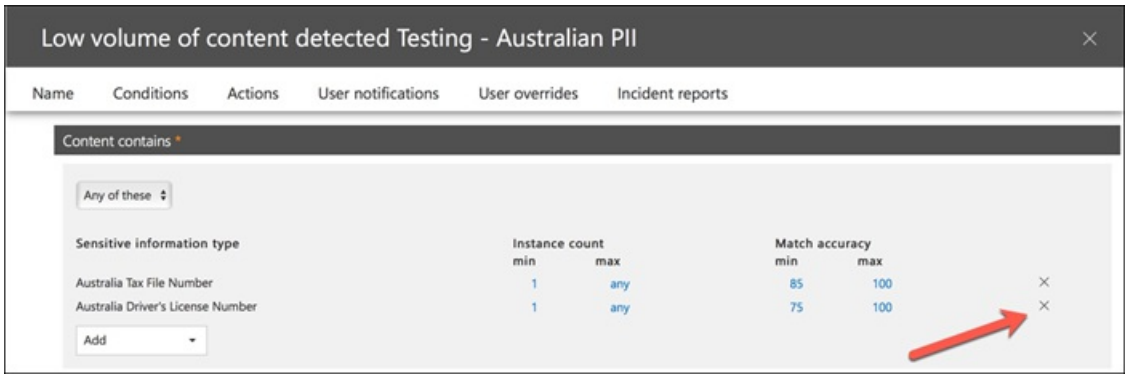

Another option is to simply increase the instance count, so that a low volume of driver's licenses is only detected when there are multiple instances.

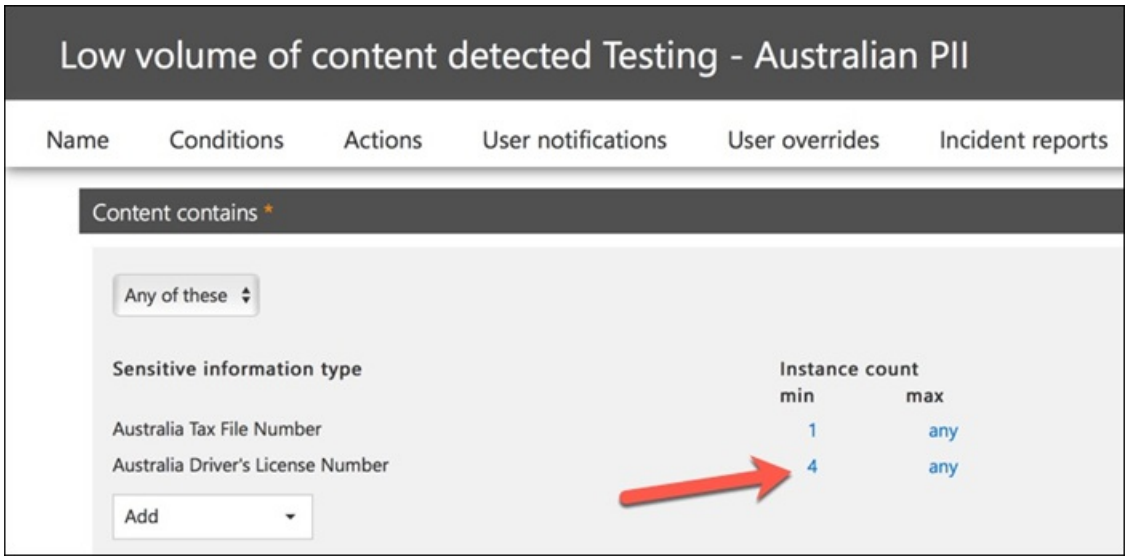

In addition to changing the instance count, you can also adjust the match accuracy (or confidence level). If your sensitive information type has multiple patterns, you can adjust the match accuracy in your rule, so that your rule matches only specific patterns. For example, to help reduce false positives, you can set the match accuracy of your rule so that it matches only the pattern with the highest confidence level. Understanding how confidence level is calculated is a bit tricky (and beyond the scope of this post), but here's a good explanation of how to use confidence level to tune your rules.

Finally, if you want to get even a bit more advanced, you can customize any sensitive information type -- for example, you can remove "Sydney NSW" from the list of keywords for Australia driver's license number, to eliminate the false positive triggered above. To learn how to do this by using XML and PowerShell, see customizing a built-in sensitive information type.

Turn on a DLP policy

When you're happy that your DLP policy is accurately and effectively detecting sensitive information types, and that your end users are ready to deal with the policies being in place, then you can enable the policy.

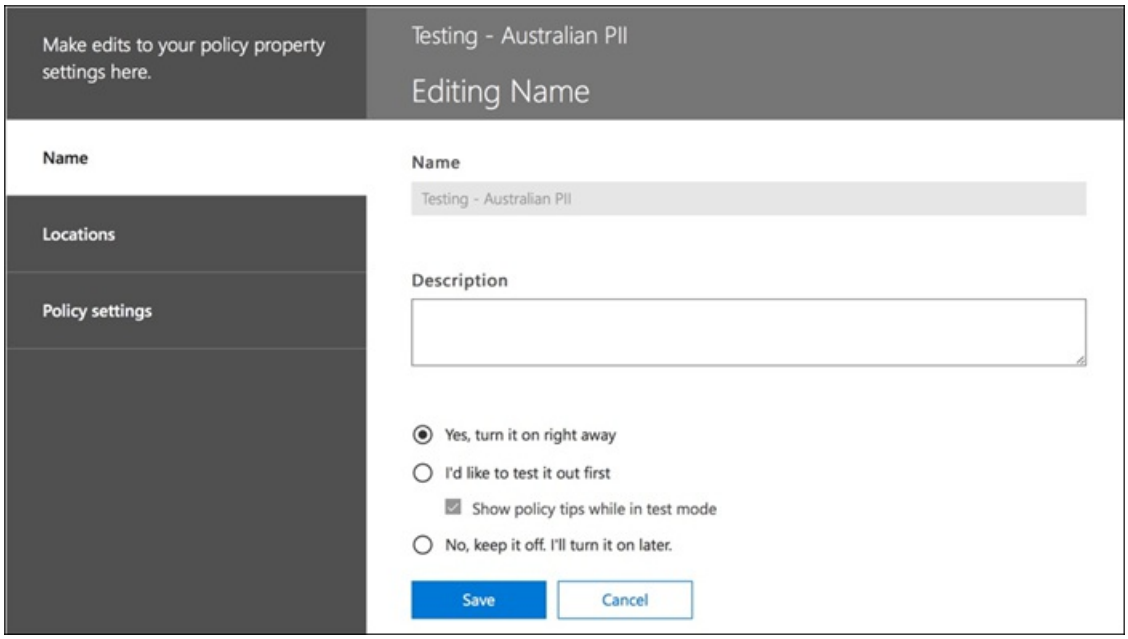

If you're waiting to see when the policy will take effect, Connect to Security & Compliance Center PowerShell and run the Get-DlpCompliancePolicy cmdlet to see the DistributionStatus.

```
PS C:\> Connect-IPPSSession


PS C:\> Get-DlpCompliancePolicy "Testing - Australian PII" | Select DistributionStatus

DistributionStatus
------------------
Pending


PS C:\> Get-DlpCompliancePolicy "Testing - Australian PII" | Select DistributionStatus

DistributionStatus
------------------
Success
```

After turning on the DLP policy, you should run some final tests of your own to make sure that the expected policy actions are occurring. If you're trying to test things like credit card data, there are websites online with information on how to generate sample credit card or other personal information that will pass checksums and trigger your policies.

Policies that allow user overrides will present that option to the user as part of the policy tip.

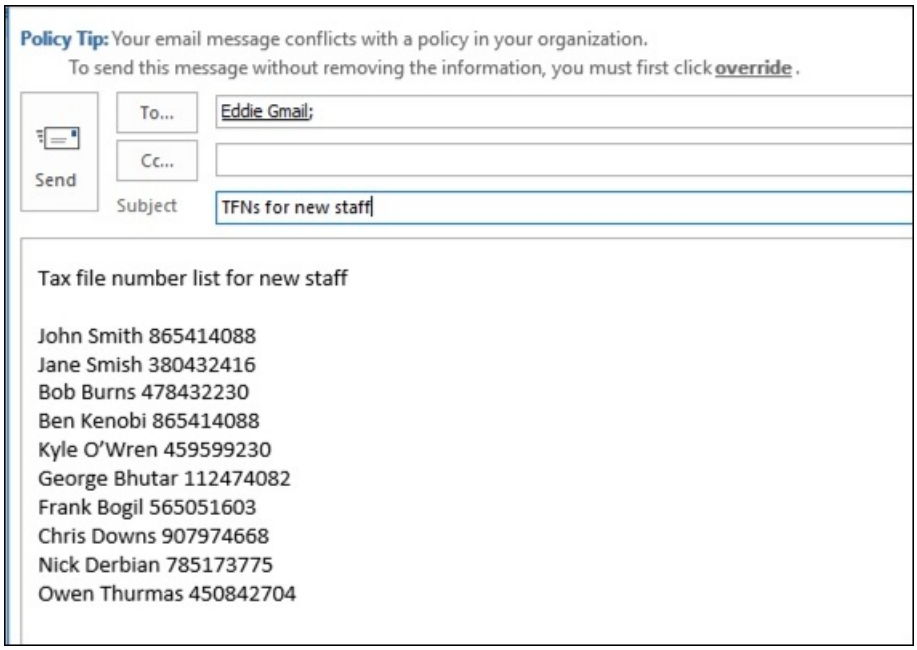

Policies that restrict content will present the warning to the user as part of the policy tip, and prevent them from sending the email.

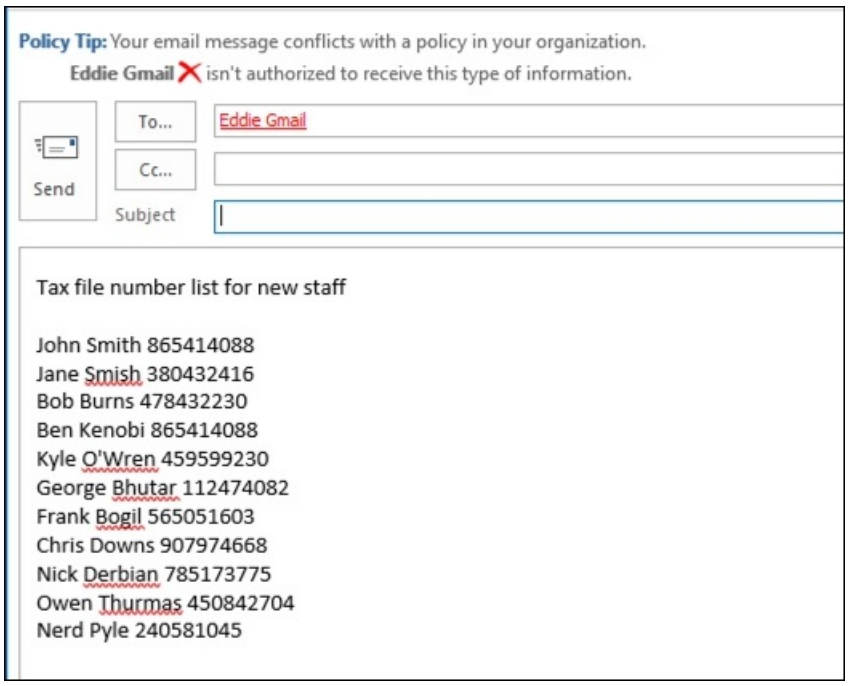

Summary

Data loss prevention policies are useful for organizations of all types. Testing some DLP policies is a low risk exercise due to the control you have over things like policy tips, end user overrides, and incident reports. You can quietly test some DLP policies to see what type of violations are already occurring in your organization, and then craft policies with low false positive rates, educate your users on what is allowed and not allowed, and then roll out your DLP policies to the organization.

Get started with Endpoint data loss prevention

2/18/2021 • 7 minutes to read • Edit Online

Microsoft Endpoint data loss prevention (Endpoint DLP) is part of the Microsoft 365 data loss prevention (DLP) suite of features you can use to discover and protect sensitive items across Microsoft 365 services. For more information about all of Microsoft's DLP offerings, see Overview of data loss prevention. To learn more about Endpoint DLP, see Learn about Endpoint data loss prevention

Microsoft Endpoint DLP allows you to monitor Windows 10 devices and detect when sensitive items are used and shared. This gives you the visibility and control you need to ensure that they are used and protected properly, and to help prevent risky behavior that might compromise them.

Before you begin

SKU/subscriptions licensing

Before you get started with Endpoint DLP, you should confirm your Microsoft 365 subscription and any add-ons. To access and use Endpoint DLP functionality, you must have one of these subscriptions or add-ons.

- Microsoft 365 E5
- Microsoft 365 A5 (EDU)
- Microsoft 365 E5 compliance
- Microsoft 365 A5 compliance
- Microsoft 365 E5 information protection and governance
- Microsoft 365 A5 information protection and governance

Permissions

To enable device management, the account you use must be a member of any one of these roles:

- Global admin
- Security admin
- Compliance admin

If you want to use a custom account to view the device management settings, it must be in one of these roles:

- Global admin
- Compliance admin
- Compliance data admin
- Global reader

If you want to use a custom account to access the onboarding/offboarding page, it must be in one of these roles:

- Global admin
- Compliance admin

If you want to use a custom account to turn on/off device monitoring, it must be in one of these roles:

- Global admin
- Compliance admin

Data from Endpoint DLP can be viewed in Activity explorer. There are four roles that grant permission to activity explorer, the account you use for accessing the data must be a member of any one of them.

- Global admin
- Compliance admin
- Security admin
- Compliance data admin
- Global reader
- Security reader
- Reports reader

Prepare your endpoints

Make sure that the Windows 10 devices that you plan on deploying Endpoint DLP to meet these requirements.

1. Must be running Windows 10 x64 build 1809 or later.
2. Antimalware Client Version is 4.18.2009.7 or newer. Check your current version by opening Windows Security app, select the Settings icon, and then select About. The version number is listed under Antimalware Client Version. Update to the latest Antimalware Client Version by installing Windows Update KB4052623.

   > **NOTE**
   > None of Windows Security components need to be active, you can run Endpoint DLP independent of Windows Security status, but the Real-time protection and Behavior monitor) must be enabled.

3. The following Windows Updates are installed.

   > **NOTE**
   > These updates are not a pre-requisite to onboard a device to Endpoint DLP, but contain fixes for important issues thus must be installed before using the product.

   - For Windows 10 1809 - KB4559003, KB4577069, KB4580390
   - For Windows 10 1903 or 1909 - KB4559004, KB4577062, KB4580386
   - For Windows 10 2004 - KB4568831, KB4577063
   - For devices running Office 2016 (and not any other Office version) - KB4577063
4. All devices must be Azure Active Directory (Azure AD) joined, or Hybrid Azure AD joined.
5. Install Microsoft Chromium Edge browser on the endpoint device to enforce policy actions for the upload to cloud activity. See, Download the new Microsoft Edge based on Chromium.
6. If you are on Monthly Enterprise Channel of Microsoft 365 Apps versions 2004-2008, there is a known issue with Endpoint DLP classifying Office content and you need to update to version 2009 or later. See Update history for Microsoft 365 Apps (listed by date) for current versions. To learn more about this issue, see the Office Suite section of Release notes for Current Channel releases in 2020.
7. If you have endpoints that use a device proxy to connect to the internet, follow the procedures in Configure device proxy and internet connection settings for Endpoint DLP.

Onboarding devices into device management

You must enable device monitoring and onboard your endpoints before you can monitor and protect sensitive items on a device. Both of these actions are done in the Microsoft 365 Compliance portal.

When you want to onboard devices that haven't been onboarded yet, you'll download the appropriate script and deploy it to those devices. Follow the Onboarding devices procedure.

If you already have devices onboarded into Microsoft Defender for Endpoint, they will already appear in the managed devices list. Follow the With devices onboarded into Microsoft Defender for Endpoint procedure.

Onboarding devices

In this deployment scenario, you'll onboard devices that have not been onboarded yet, and you just want to monitor and protect sensitive items from unintentional sharing on Windows 10 devices.

1. Open the Microsoft compliance center.
2. Open the Compliance Center settings page and choose **Onboard devices**.

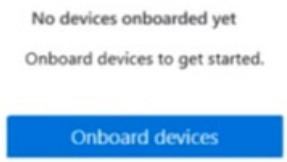

> **NOTE**
> While it usually takes about 60 seconds for device onboarding to be enabled, please allow up to 30 minutes before engaging with Microsoft support.

3. Choose **Device management** to open the **Devices** list. The list will be empty until you onboard devices.
4. Choose **Onboarding** to begin the onboarding process.
5. Choose the way you want to deploy to these additional devices from the **Deployment method** list and then **download package**.

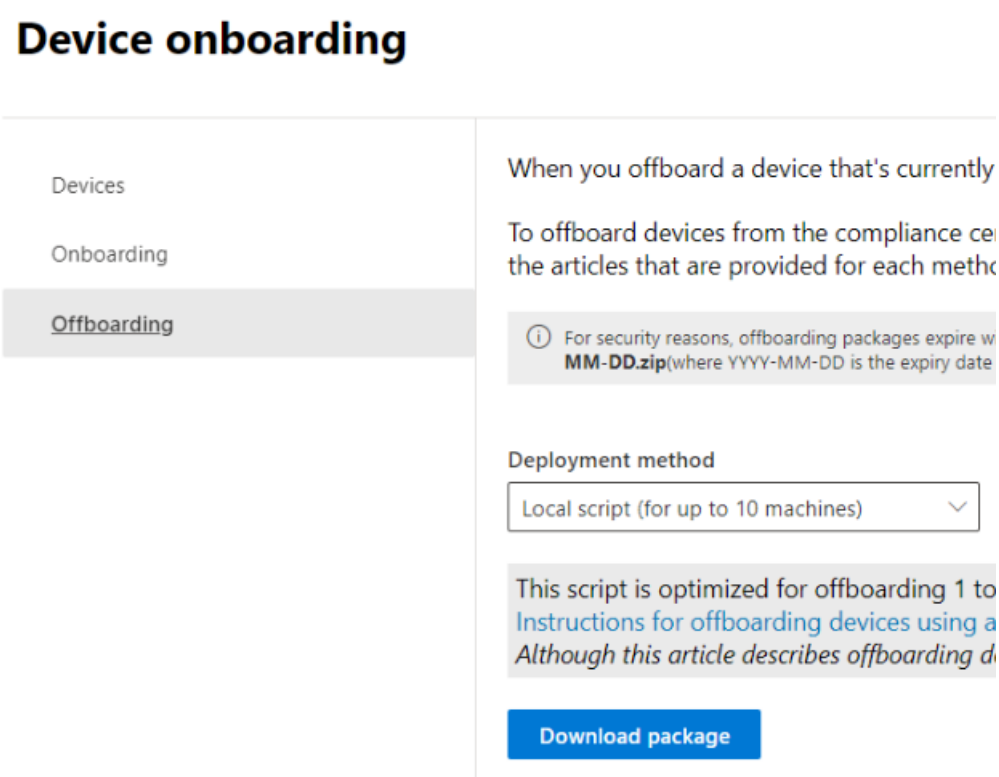

6. Follow the appropriate procedures in Onboarding tools and methods for Windows 10 machines. This link takes you to a landing page where you can access Microsoft Defender for Endpoint procedures that match the deployment package you selected in step 5:
   - Onboard Windows 10 machines using Group Policy
   - Onboard Windows machines using Microsoft Endpoint Configuration Manager

- Onboard Windows 10 machines using Mobile Device Management tools
- Onboard Windows 10 machines using a local script
- Onboard non-persistent virtual desktop infrastructure (VDI) machines.

Once done and endpoint is onboarded, it should be visible in the devices list and also start reporting audit activity logs to Activity explorer.

> **NOTE**
> This experience is under license enforcement. Without the required license, data will not be visible or accessible.

With devices onboarded into Microsoft Defender for Endpoint

In this scenario, Microsoft Defender for Endpoint is already deployed and there are endpoints reporting in. All these endpoints will appear in the managed devices list. You can continue to onboard new devices into Endpoint DLP to expand coverage by using the Onboarding devices procedure.

1. Open the Microsoft compliance center.
2. Open the Compliance Center settings page and choose **Enable device monitoring**.
3. Choose **Device management** to open the **Devices** list. You should see the list of devices that are already reporting in to Microsoft Defender for Endpoint.

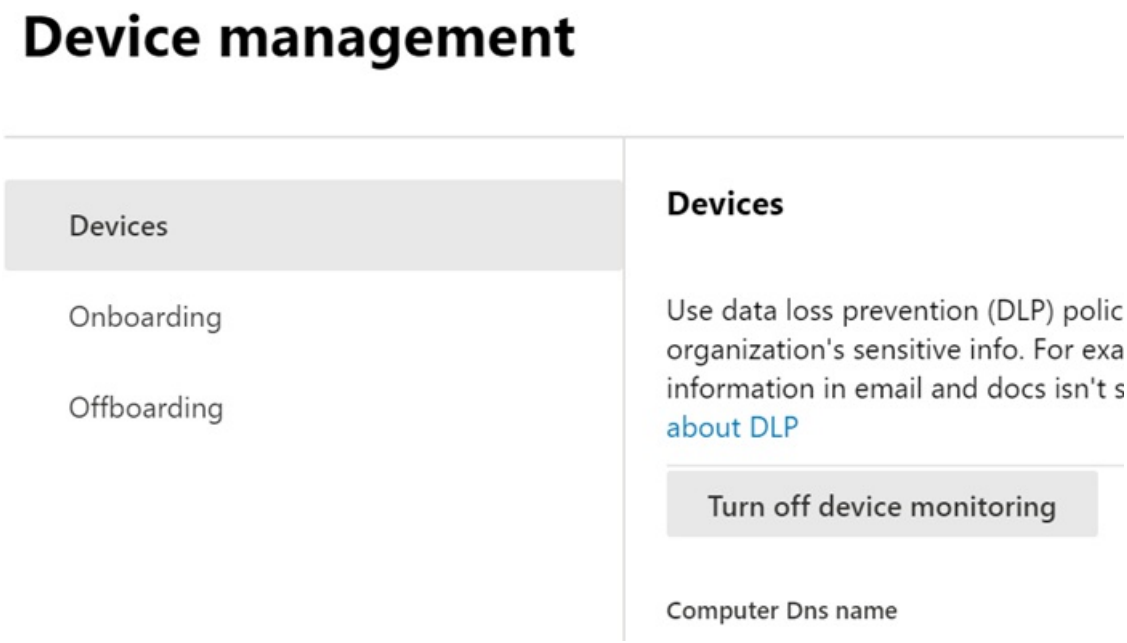

4. Choose **Onboarding** if you need to onboard additional devices.
5. Choose the way you want to deploy to these additional devices from the **Deployment method** list and then **Download package**.
6. Follow the appropriate procedures in Onboarding tools and methods for Windows 10 machines. This link takes you to a landing page where you can access Microsoft Defender for Endpoint procedures that match the deployment package you selected in step 5:

   - Onboard Windows 10 machines using Group Policy
   - Onboard Windows machines using Microsoft Endpoint Configuration Manager
   - Onboard Windows 10 machines using Mobile Device Management tools
   - Onboard Windows 10 machines using a local script
   - Onboard non-persistent virtual desktop infrastructure (VDI) machines.

Once done and endpoint is onboarded, it should be visible under the **Devices** table and also start reporting audit logs to the **Activity Explorer**.

> **NOTE**
>
> This experience is under license enforcement. Without the required license, data will not be visible or accessible.

Viewing Endpoint DLP alerts in DLP Alerts Management dashboard

1. Open the Data loss prevention page in the Microsoft 365 Compliance center and choose Alerts.
2. Refer to the procedures in How to configure and view alerts for your DLP policies to view alerts for your Endpoint DLP policies.

Viewing Endpoint DLP data in activity explorer

1. Open the Data classification page for your domain in the Microsoft 365 Compliance center and choose Activity explorer.
2. Refer to the procedures in Get started with Activity explorer to access and filter all the data for your Endpoint devices.

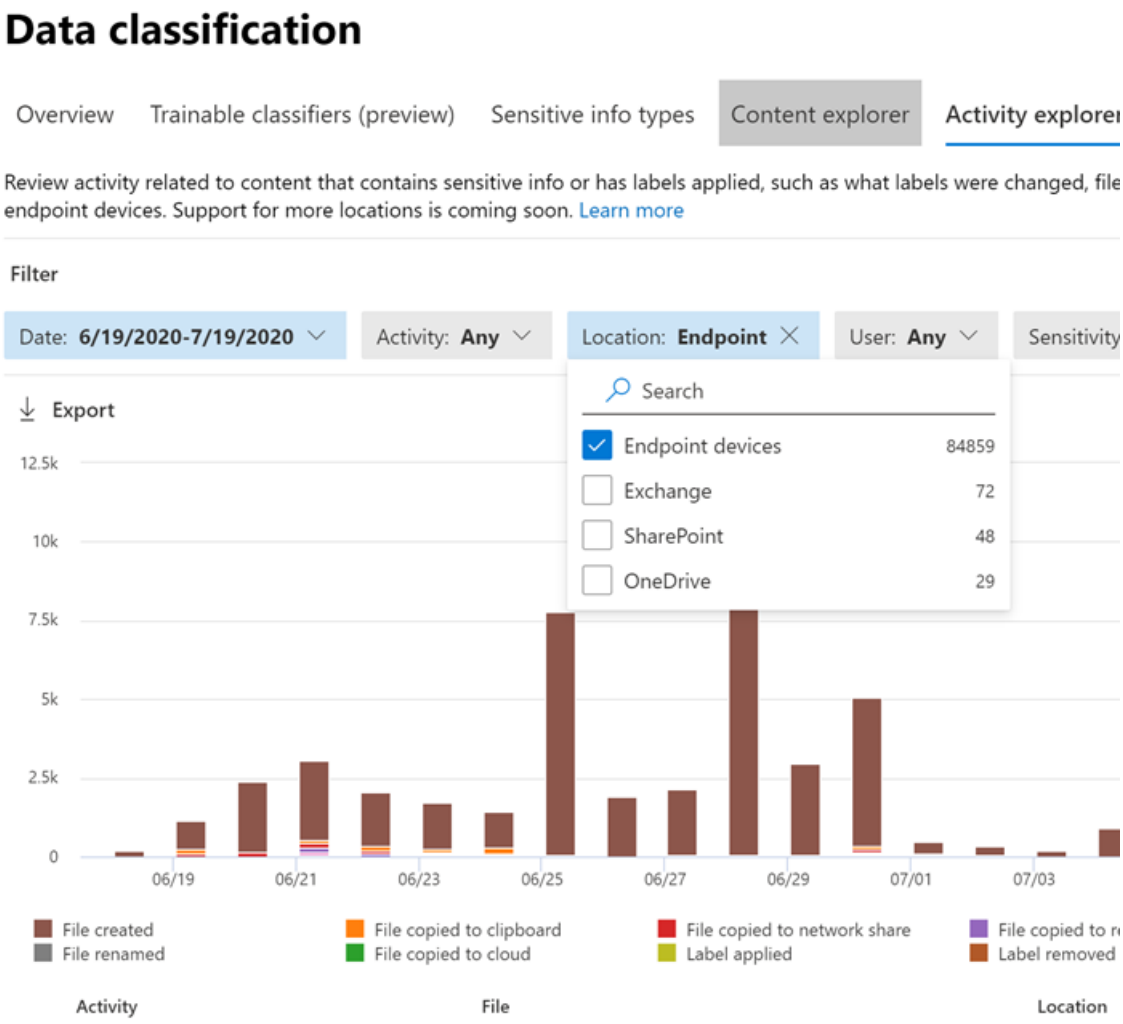

Next steps

Now that you have onboarded devices and can view the activity data in Activity explorer, you are ready to move on to your next step where you create DLP policies that protect your sensitive items.

- Using Endpoint data loss prevention

See also

- Learn about Endpoint data loss prevention
- Using Endpoint data loss prevention
- Overview of data loss prevention
- Create, test, and tune a DLP policy
- Get started with Activity explorer
- Microsoft Defender for Endpoint
- Onboarding tools and methods for Windows 10 machines
- Microsoft 365 subscription
- Azure AD joined devices
- Download the new Microsoft Edge based on Chromium

Configure device proxy and internet connection settings for Endpoint DLP

2/18/2021 • 5 minutes to read • Edit Online

Microsoft Endpoint DLP uses Microsoft Windows HTTP (WinHTTP) to report data and communicate with the Microsoft endpoint cloud service. The embedded Endpoint DLP runs in system context using the LocalSystem account.

> **TIP**
>
> For organizations that use forward proxies as a gateway to the Internet, you can use network protection to investigate behind a proxy. For more information, see Investigate connection events that occur behind forward proxies.

The WinHTTP configuration setting is independent of the Windows Internet (WinINet) Internet browsing proxy settings and can only discover a proxy server by using the following auto discovery methods:

- Transparent proxy
- Web Proxy Auto-discovery Protocol (WPAD)

> **NOTE**
>
> If you're using Transparent proxy or WPAD in your network topology, you don't need special configuration settings. For more information on Defender for Endpoint URL exclusions in the proxy, see Enable access to Endpoint DLP cloud service URLs in the proxy server.

- Manual static proxy configuration:
  - Registry based configuration
  - WinHTTP configured using netsh command – Suitable only for desktops in a stable topology (for example: a desktop in a corporate network behind the same proxy)

Configure the proxy server manually using a registry-based static proxy

For endpoint devices that aren't permitted to connect to the Internet, you need to configure a registry-based static proxy. You need to configure this to allow only Microsoft Endpoint DLP to report diagnostic data and communicate with Microsoft endpoint cloud service.

The static proxy is configurable through Group Policy (GP). The group policy can be found under:

1. Open **Administrative Templates > Windows Components > Data Collection and Preview Builds > Configure Authenticated Proxy usage for the Connected User Experience and Telemetry Service**
2. Set it to **Enabled** and select **Disable Authenticated Proxy usage**:

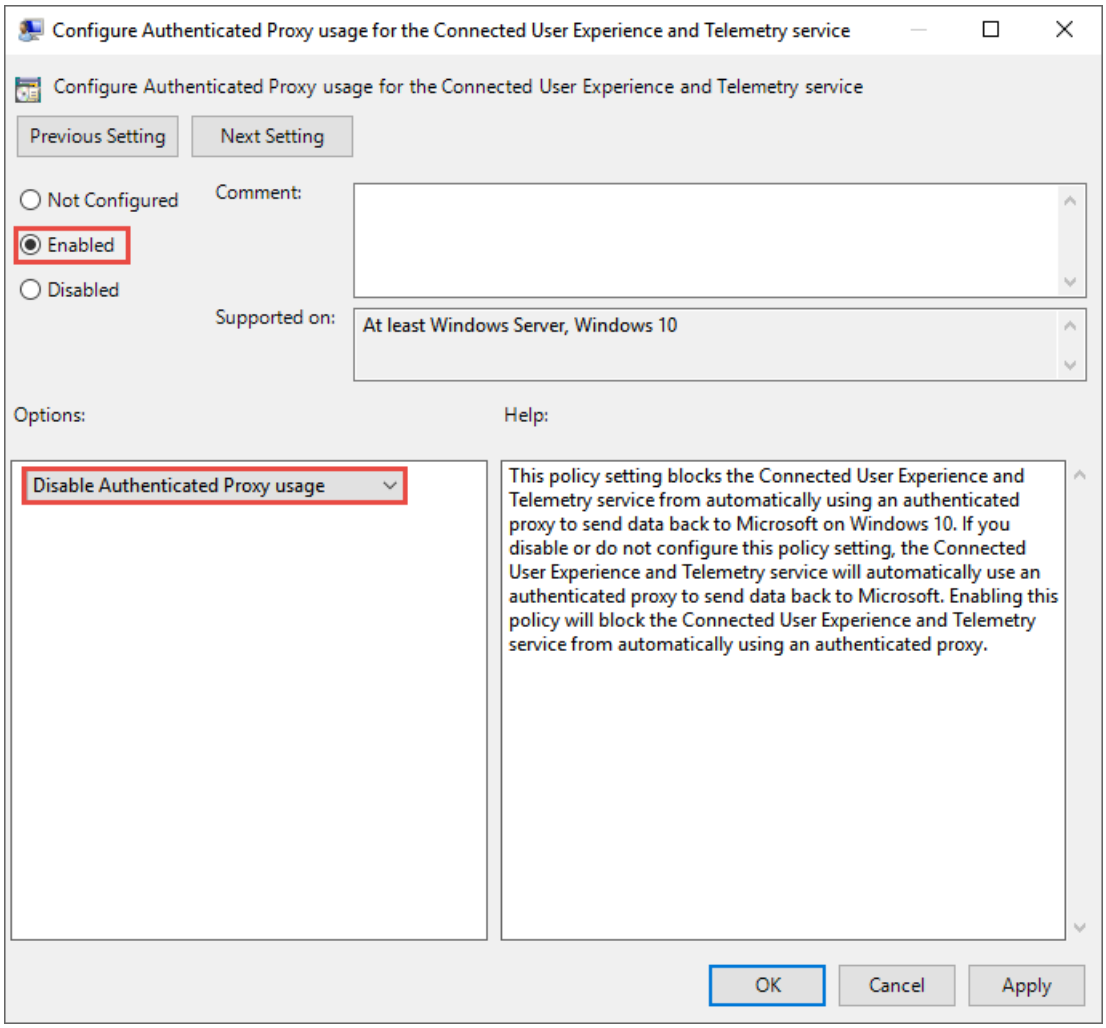

3. Open **Administrative Templates > Windows Components > Data Collection and Preview Builds > Configure connected user experiences and telemetry**:

Configure the proxy

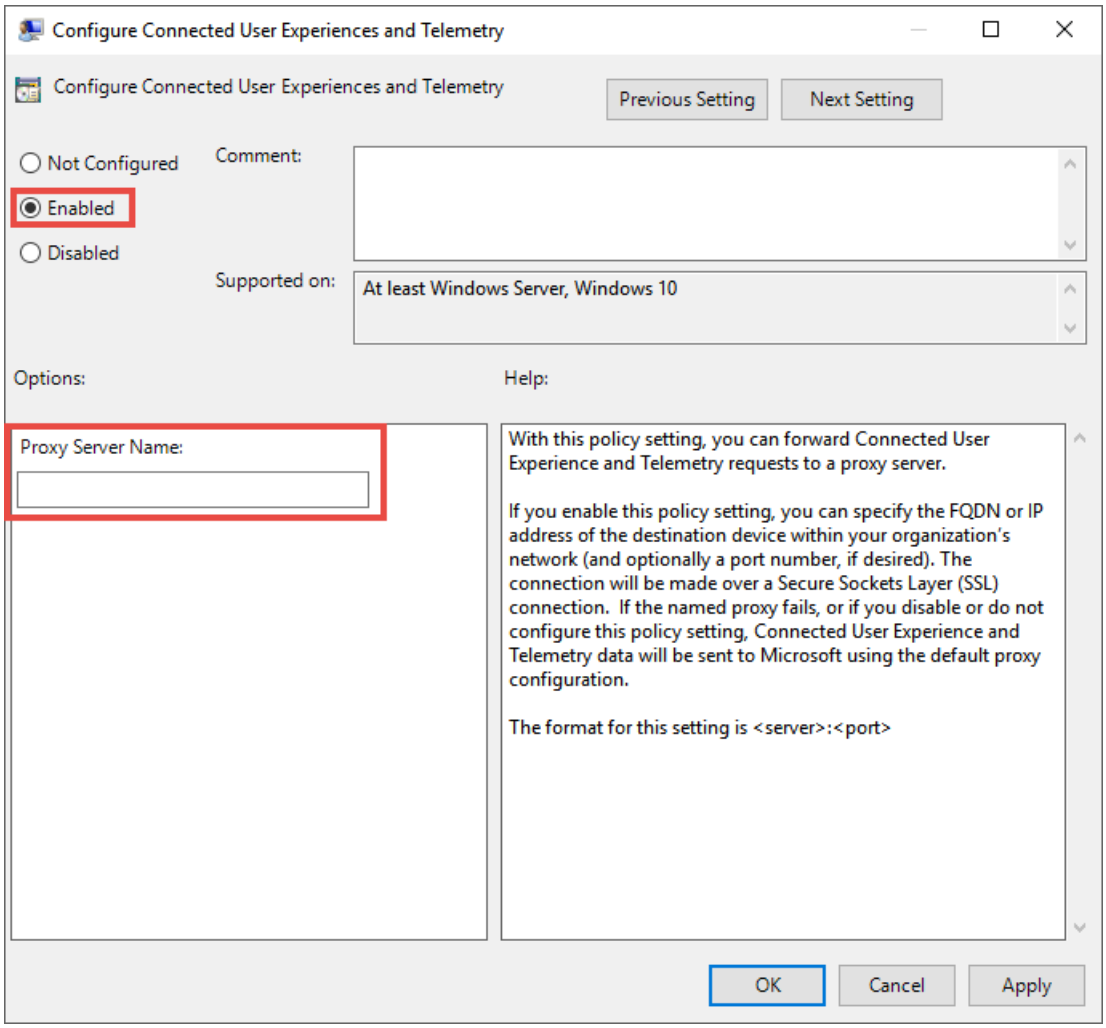

The policy sets two registry values `TelemetryProxyServer` as REG_SZ and `DisableEnterpriseAuthProxy` as REG_DWORD under the registry key `HKLM\Software\Policies\Microsoft\Windows\DataCollection`.

The registry value TelemetryProxyServer is in this format <server name or ip>:<port>. For example: **10.0.0.6:8080**

The registry value `DisableEnterpriseAuthProxy` should be set to 1.

Configure the proxy server manually using "netsh" command

Use netsh to configure a system-wide static proxy.

> **NOTE**
> This will affect all applications including Windows services which use WinHTTP with default proxy. - Laptops that are changing topology (for example: from office to home) will malfunction with netsh. Use the registry-based static proxy configuration.

1. Open an elevated command-line:
   a. Go to **Start** and type **cmd**
   b. Right-click **Command prompt** and select **Run as administrator**.
2. Enter the following command and press **Enter**:

   ```
   netsh winhttp set proxy <proxy>:<port>
   ```

   For example: **netsh winhttp set proxy 10.0.0.6:8080**

3. To reset the winhttp proxy, enter the following command and press **Enter**:

```
netsh winhttp reset proxy
```

See Netsh Command Syntax, Contexts, and Formatting to learn more.

Enable access to Endpoint DLP cloud service URLs in the proxy server

If a proxy or firewall is blocking all traffic by default and allowing only specific domains through, add the domains listed in the downloadable sheet to the allowed domains list.

This downloadable spreadsheet lists the services and their associated URLs that your network must be able to connect to. You should ensure that there are no firewall or network filtering rules that would deny access to these URLs, or you may need to create an allow rule specifically for them.

If a proxy or firewall has HTTPS scanning (SSL inspection) enabled, exclude the domains listed in the above table from HTTPS scanning. If a proxy or firewall is blocking anonymous traffic, as Endpoint DLP is connecting from system context, make sure anonymous traffic is permitted in the previously listed URLs.

Verify client connectivity to Microsoft cloud service URLs

Verify the proxy configuration completed successfully, that WinHTTP can discover and communicate through the proxy server in your environment, and that the proxy server allows traffic to the Defender for Endpoint service URLs.

1. Download the MDATP Client Analyzer tool to the PC where Endpoint DLP is running on.
2. Extract the contents of MDATPClientAnalyzer.zip on the device.
3. Open an elevated command-line:
   a. Go to **Start** and type **cmd**.
   b. Right-click **Command prompt** and select **Run as administrator**.
4. Enter the following command and press **Enter**:

```
HardDrivePath\MDATPClientAnalyzer.cmd
```

Replace *HardDrivePath* with the path where the MDATPClientAnalyzer tool was downloaded to, for example

C:\Work\tools\MDATPClientAnalyzer\MDATPClientAnalyzer.cmd

5. Extract the **MDATPClientAnalyzerResult.zip*** file created by tool in the folder used in the *HardDrivePath*.

6. Open **MDATPClientAnalyzerResult.txt** and verify that you have performed the proxy configuration steps to enable server discovery and access to the service URLs. The tool checks the connectivity of Defender for Endpoint service URLs that Defender for Endpoint client is configured to interact with. It then prints the results into the **MDATPClientAnalyzerResult.txt** file for each URL that can potentially be used to communicate with the Defender for Endpoint services. For example:

   Testing URL : https://xxx.microsoft.com/xxx
   1 - Default proxy: Succeeded (200)
   2 - Proxy auto discovery (WPAD): Succeeded (200)
   3 - Proxy disabled: Succeeded (200)
   4 - Named proxy: Doesn't exist
   5 - Command line proxy: Doesn't exist

If at least one of the connectivity options returns a (200) status, then the Defender for Endpoint client can communicate with the tested URL properly using this connectivity method.

However, if the connectivity check results indicate a failure, an HTTP error is displayed (see HTTP Status Codes). You can then use the URLs in the table shown in Enable access to Endpoint DLP cloud service URLs in the proxy server. The URLs you'll use will depend on the region selected during the onboarding procedure. [!NOTE] The Connectivity Analyzer tool is not compatible with ASR rule Block process creations originating from PSExec and WMI commands. You will need to temporarily disable this rule to run the connectivity tool.

[!NOTE] When the TelemetryProxyServer is set, in Registry or via Group Policy, Defender for Endpoint will fall back to direct if it can't access the defined proxy. Related topics • Onboard Windows 10 devices • Troubleshoot Microsoft Endpoint DLP onboarding issues

See also

- Learn about Endpoint data loss prevention
- Using Endpoint data loss prevention
- Overview of data loss prevention
- Create, test, and tune a DLP policy
- Get started with Activity explorer
- Microsoft Defender for Endpoint
- Onboarding tools and methods for Windows 10 machines
- Microsoft 365 subscription
- Azure AD joined devices
- Download the new Microsoft Edge based on Chromium

Onboarding tools and methods for Windows 10 devices

2/18/2021 • 2 minutes to read • Edit Online

Applies to:

- Microsoft 365 Endpoint data loss prevention (DLP)

Devices in your organization must be configured so that the Microsoft 365 Endpoint data loss prevention service can get sensor data from them. There are various methods and deployment tools that you can use to configure the devices in your organization.

The following deployment tools and methods are supported:

- group policy
- Microsoft Endpoint Configuration Manager
- Mobile Device Management (including Microsoft Intune)
- local script

In this section

TOPIC	DESCRIPTION
Onboard Windows 10 devices using Group Policy	Use Group Policy to deploy the configuration package on devices.
Onboard Windows devices using Microsoft Endpoint Configuration Manager	You can use either use Microsoft Endpoint Configuration Manager (current branch) version 1606 or Microsoft Endpoint Configuration Manager (current branch) version 1602 or earlier to deploy the configuration package on devices.
Onboard Windows 10 devices using Mobile Device Management tools	Use Mobile Device Management tools or Microsoft Intune to deploy the configuration package on device.
Onboard Windows 10 devices using a local script	Learn how to use the local script to deploy the configuration package on endpoints.
Onboard non-persistent virtual desktop infrastructure (VDI) devices	Learn how to use the configuration package to configure VDI devices.

Onboard Windows 10 devices using Group Policy

11/2/2020 • 3 minutes to read • Edit Online

Applies to:

- Microsoft 365 Endpoint data loss prevention (DLP)
- Group Policy

> **NOTE**
> To use Group Policy (GP) updates to deploy the package, you must be on Windows Server 2008 R2 or later.

> For Windows Server 2019, you may need to replace NT AUTHORITY\Well-Known-System-Account with NT AUTHORITY\SYSTEM of the XML file that the Group Policy preference creates.

Onboard devices using Group Policy

1. Open the GP configuration package .zip file (*DeviceComplianceOnboardingPackage.zip*) that you downloaded from the service onboarding wizard. You can also get the package from Microsoft Compliance center
2. In the navigation pane, select **Settings** > **Device Onboarding**.
3. In the **Deployment method** field, select **Group policy**.
4. Click **Download package** and save the .zip file.
5. Extract the contents of the .zip file to a shared, read-only location that can be accessed by the device. You should have a folder called *OptionalParamsPolicy* and the file *DeviceComplianceLocalOnboardingScript.cmd*.
6. Open the Group Policy Management Console (GPMC), right-click the Group Policy Object (GPO) you want to configure and click **Edit**.
7. In the **Group Policy Management Editor**, go to **Computer configuration**, then **Preferences**, and then **Control panel settings**.
8. Right-click **Scheduled tasks**, point to **New**, and then click **Immediate Task (At least Windows 7)**.
9. In the **Task** window that opens, go to the **General** tab. Under **Security options** click **Change User or Group** and type SYSTEM and then click **Check Names** then **OK**. NT AUTHORITY\SYSTEM appears as the user account the task will run as.
10. Select **Run whether user is logged on or not** and check the **Run with highest privileges** check box.
11. Go to the **Actions** tab and click **New...** Ensure that **Start a program** is selected in the **Action** field. Enter the file name and location of the shared *WindowsDefenderATPOnboardingScript.cmd* file.
12. Click **OK** and close any open GPMC windows.

Offboard devices using Group Policy

For security reasons, the package used to Offboard devices will expire 30 days after the date it was downloaded. Expired offboarding packages sent to a device will be rejected. When downloading an offboarding package you will be notified of the packages expiry date and it will also be included in the package name.

> **NOTE**
> Onboarding and offboarding policies must not be deployed on the same device at the same time, otherwise this will cause unpredictable collisions.

1. Get the offboarding package from Microsoft Compliance center.
2. In the navigation pane, select **Settings** > **//Device onboarding** > **Offboarding**.
3. In the **Deployment method** field, select **Group policy**.
4. Click **Download package** and save the .zip file.
5. Extract the contents of the .zip file to a shared, read-only location that can be accessed by the device. You should have a file named *DeviceComplianceOffboardingScript_valid_until_YYYY-MM-DD.cmd*.
6. Open the Group Policy Management Console (GPMC), right-click the Group Policy Object (GPO) you want to configure and click **Edit**.
7. In the **Group Policy Management Editor**, go to **Computer configuration,** then **Preferences**, and then **Control panel settings**.
8. Right-click **Scheduled tasks**, point to **New**, and then click **Immediate task**.
9. In the **Task** window that opens, go to the **General** tab. Choose the local SYSTEM user account (BUILTIN\SYSTEM) under **Security options**.
10. Select **Run whether user is logged on or not** and check the **Run with highest privileges** check-box.
11. Go to the **Actions** tab and click **New...**. Ensure that **Start a program** is selected in the **Action** field. Enter the file name and location of the shared *DeviceComplianceOffboardingScript_valid_until_YYYY-MM-DD.cmd* file.
12. Click **OK** and close any open GPMC windows.

> **IMPORTANT**
> Offboarding causes the device to stop sending sensor data to the portal but data from the device.

Monitor device configuration

With Group Policy there isn't an option to monitor deployment of policies on the devices. Monitoring can be done directly on the portal, or by using the different deployment tools.

Monitor devices using the portal

1. Go to Microsoft Compliance center.
2. Click **Devices** list.
3. Verify that devices are appearing.

NOTE

It can take several days for devices to start showing on the **Devices list**. This includes the time it takes for the policies to be distributed to the device, the time it takes before the user logs on, and the time it takes for the endpoint to start reporting.

Related topics

- Onboard Windows 10 devices using Microsoft Endpoint Configuration Manager
- Onboard Windows 10 devices using Mobile Device Management tools
- Onboard Windows 10 devices using a local script
- Onboard non-persistent virtual desktop infrastructure (VDI) devices
- Run a detection test on a newly onboarded Microsoft Defender ATP devices
- Troubleshoot Microsoft Defender Advanced Threat Protection onboarding issues

Onboard Windows 10 devices using Mobile Device Management tools

11/2/2020 • 2 minutes to read • Edit Online

Applies to:

- Microsoft 365 Endpoint data loss prevention (DLP)

You can use mobile device management (MDM) solutions to configure devices. Microsoft 365 Endpoint data loss prevention supports MDMs by providing OMA-URIs to create policies to manage devices.

Before you begin

If you're using Microsoft Intune, you must have the device MDM Enrolled. Otherwise, settings will not be applied successfully.

For more information on enabling MDM with Microsoft Intune, see Device enrollment (Microsoft Intune).

Onboard devices using Microsoft Intune

Follow the instructions from Intune.

> **NOTE**
> - The **Health Status for onboarded devices** policy uses read-only properties and can't be remediated.

Offboard and monitor devices using Mobile Device Management tools

For security reasons, the package used to Offboard devices will expire 30 days after the date it was downloaded. Expired offboarding packages sent to a device will be rejected. When downloading an offboarding package you will be notified of the packages expiry date and it will also be included in the package name.

> **NOTE**
> Onboarding and offboarding policies must not be deployed on the same device at the same time, otherwise this will cause unpredictable collisions.

1. Get the offboarding package from Microsoft Compliance center.
2. In the navigation pane, select **Settings** > **Device onboarding** > **Offboarding**.
3. In the **Deployment method** field, select **Mobile Device Management / Microsoft Intune**.
4. Click **Download package**, and save the .zip file.
5. Extract the contents of the .zip file to a shared, read-only location that can be accessed by the network administrators who will deploy the package. You should have a file named *DeviceCompliance_valid_until_YYYY-MM-DD.offboarding*.
6. Use the Microsoft Intune custom configuration policy to deploy the following supported OMA-URI

settings.

OMA-URI: ./Device/Vendor/MSFT/WindowsAdvancedThreatProtection/Offboarding
Date type: String
Value: [Copy and paste the value from the content of the DeviceCompliance_valid_until_YYYY-MM-DD.offboarding file]

For more information on Microsoft Intune policy settings see, Windows 10 policy settings in Microsoft Intune.

> **NOTE**
> The **Health Status for offboarded devices** policy uses read-only properties and can't be remediated.

> **IMPORTANT**
> Offboarding causes the device to stop sending sensor data to the portal but data from the device, including reference to any alerts it has had will be retained for up to 6 months.

Related topics

- Onboard Windows 10 devices using Group Policy
- Onboard Windows 10 devices using Microsoft Endpoint Configuration Manager
- Onboard Windows 10 devices using a local script
- Onboard non-persistent virtual desktop infrastructure (VDI) devices
- Troubleshoot Microsoft Defender Advanced Threat Protection onboarding issues

Onboard Windows 10 devices using Configuration Manager

11/2/2020 • 6 minutes to read • Edit Online

Applies to:

- Microsoft 365 Endpoint data loss prevention (DLP)
- System Center 2012 R2 Configuration Manager

Onboard devices using System Center Configuration Manager

1. Open the Configuration Manager configuration package .zip file (*DeviceComplianceOnboardingPackage.zip*) that you downloaded from the service onboarding wizard. You can also get the package from Microsoft Compliance center.
2. In the navigation pane, select **Settings** > **Device Onboarding** > **Onboarding**.
3. In the **Deployment method** field, select **Microsoft Endpoint Configuration Manager 2012/2012 R2/1511/1602**.
4. Select **Download package**, and save the .zip file.
5. Extract the contents of the .zip file to a shared, read-only location that can be accessed by the network administrators who will deploy the package. You should have a file named *DeviceComplianceOnboardingScript.cmd*.
6. Deploy the package by following the steps in the Packages and Programs in System Center 2012 R2 Configuration Manager article.
7. Choose a predefined device collection to deploy the package to.

> **NOTE**
> Microsoft 365 Endpoint data loss prevention doesn't support onboarding during the Out-Of-Box Experience (OOBE) phase. Make sure users complete OOBE after running Windows installation or upgrading.

> **TIP**
> After onboarding the device, you can choose to run a detection test to verify that an device is properly onboarded to the service. For more information, see Run a detection test on a newly onboarded Microsoft Defender ATP device.
>
> Note that it is possible to create a detection rule on a Configuration Manager application to continuously check if a device has been onboarded. An application is a different type of object than a package and program. If a device is not yet onboarded (due to pending OOBE completion or any other reason), Configuration Manager will retry to onboard the device until the rule detects the status change.
>
> This behavior can be accomplished by creating a detection rule checking if the "OnboardingState" registry value (of type REG_DWORD) = 1. This registry value is located under "HKLM\SOFTWARE\Microsoft\Windows Advanced Threat Protection\Status". For more information, see Configure Detection Methods in System Center 2012 R2 Configuration Manager.

Configure sample collection settings

For each device, you can set a configuration value to state whether samples can be collected from the device

when a request is made through Microsoft Defender Security Center to submit a file for deep analysis.

> **NOTE**
> These configuration settings are typically done through Configuration Manager.

You can set a compliance rule for configuration item in Configuration Manager to change the sample share setting on a device.

This rule should be a *remediating* compliance rule configuration item that sets the value of a registry key on targeted devices to make sure they're complaint.

The configuration is set through the following registry key entry:

```
Path: “HKLM\SOFTWARE\Policies\Microsoft\Windows Advanced Threat Protection”
Name: "AllowSampleCollection"
Value: 0 or 1
```

Where:
Key type is a D-WORD.
Possible values are:

- 0 - doesn't allow sample sharing from this device
- 1 - allows sharing of all file types from this device

The default value in case the registry key doesn't exist is 1.

For more information about System Center Configuration Manager Compliance, see Introduction to compliance settings in System Center 2012 R2 Configuration Manager.

Other recommended configuration settings

After onboarding devices to the service, it's important to take advantage of the included threat protection capabilities by enabling them with the following recommended configuration settings.

Device collection configuration

If you're using Endpoint Configuration Manager, version 2002 or later, you can choose to broaden the deployment to include servers or down-level clients.

Next generation protection configuration

The following configuration settings are recommended:

Scan

- Scan removable storage devices such as USB drives: Yes

Real-time Protection

- Enable Behavioral Monitoring: Yes
- Enable protection against Potentially Unwanted Applications at download and prior to installation: Yes

Cloud Protection Service

- Cloud Protection Service membership type: Advanced membership

Attack surface reduction Configure all available rules to Audit.

> **NOTE**
> Blocking these activities may interrupt legitimate business processes. The best approach is setting everything to audit, identifying which ones are safe to turn on, and then enabling those settings on endpoints which do not have false positive detections.

Network protection

Prior to enabling network protection in audit or block mode, ensure that you've installed the antimalware platform update, which can be obtained from the support page.

Controlled folder access

Enable the feature in audit mode for at least 30 days. After this period, review detections and create a list of applications that are allowed to write to protected directories.

For more information, see Evaluate controlled folder access.

Offboard devices using Configuration Manager

For security reasons, the package used to Offboard devices will expire 30 days after the date it was downloaded. Expired offboarding packages sent to a device will be rejected. When downloading an offboarding package, you will be notified of the packages expiry date and it will also be included in the package name.

> **NOTE**
> Onboarding and offboarding policies must not be deployed on the same device at the same time, otherwise this will cause unpredictable collisions.

Offboard devices using Microsoft Endpoint Configuration Manager current branch

If you use Microsoft Endpoint Configuration Manager current branch, see Create an offboarding configuration file.

Offboard devices using System Center 2012 R2 Configuration Manager

1. Get the offboarding package from Microsoft Compliance center:
2. In the navigation pane, select **Settings** > **Device onboarding** > **Offboarding**.
3. Select Windows 10 as the operating system.
4. In the **Deployment method** field, select **Microsoft Endpoint Configuration Manager 2012/2012 R2/1511/1602**.
5. Select **Download package**, and save the .zip file.
6. Extract the contents of the .zip file to a shared, read-only location that can be accessed by the network administrators who will deploy the package. You should have a file named *DeviceComplianceOffboardingScript_valid_until_YYYY-MM-DD.cmd*.
7. Deploy the package by following the steps in the Packages and Programs in System Center 2012 R2 Configuration Manager article.
8. Choose a predefined device collection to deploy the package to.

IMPORTANT

Offboarding causes the device to stop sending sensor data to the portal but data from the device, including reference to any alerts it has had will be retained for up to 6 months.

Monitor device configuration

If you're using Microsoft Endpoint Configuration Manager current branch, use the built-in Microsoft Defender ATP dashboard in the Configuration Manager console. For more information, see Microsoft Defender Advanced Threat Protection - Monitor.

If you're using System Center 2012 R2 Configuration Manager, monitoring consists of two parts:

1. Confirming the configuration package has been correctly deployed and is running (or has successfully run) on the devices in your network.
2. Checking that the devices are compliant with the Microsoft 365 Endpoint data loss prevention service (this ensures the device can complete the onboarding process and can continue to report data to the service).

Confirm the configuration package has been correctly deployed

1. In the Configuration Manager console, click **Monitoring** at the bottom of the navigation pane.
2. Select **Overview** and then **Deployments**.
3. Select on the deployment with the package name.
4. Review the status indicators under **Completion Statistics** and **Content Status**.

   If there are failed deployments (devices with **Error**, **Requirements Not Met**, or **Failed statuses**), you may need to troubleshoot the devices. For more information, see, Troubleshoot Microsoft Defender Advanced Threat Protection onboarding issues.

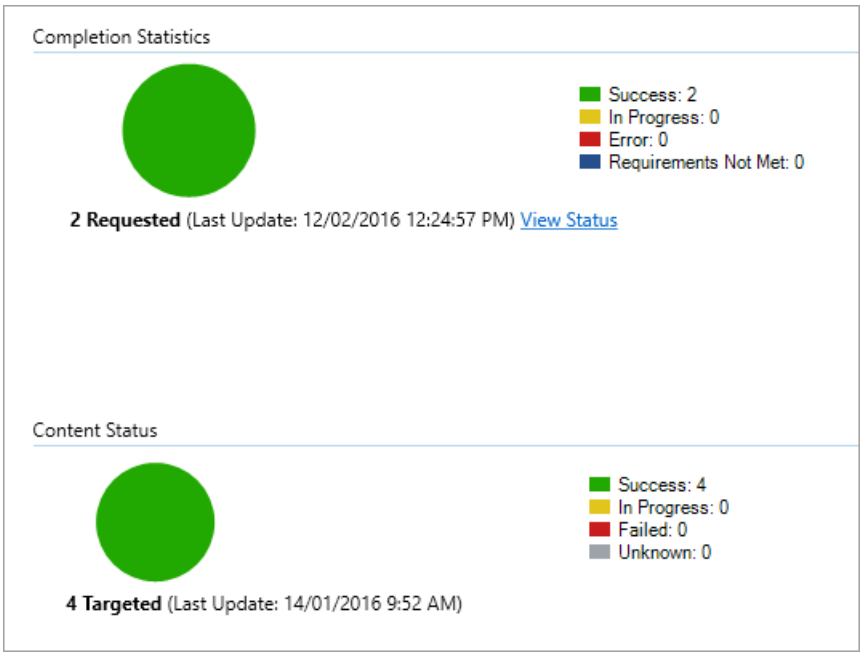

Check that the devices are compliant with the Microsoft 365 Endpoint data loss prevention service

You can set a compliance rule for configuration item in System Center 2012 R2 Configuration Manager to monitor your deployment.

> **NOTE**
>
> This procedure and registry entry applies to Endpoint DLP as well as Advanced Threat Protection.

This rule should be a *non-remediating* compliance rule configuration item that monitors the value of a registry key on targeted devices.

Monitor the following registry key entry:

```
Path: “HKLM\SOFTWARE\Microsoft\Windows Advanced Threat Protection\Status”
Name: “OnboardingState”
Value: “1”
```

For more information, see Introduction to compliance settings in System Center 2012 R2 Configuration Manager.

Related topics

- Onboard Windows 10 devices using Group Policy
- Onboard Windows 10 devices using Mobile Device Management tools
- Onboard Windows 10 devices using a local script
- Onboard non-persistent virtual desktop infrastructure (VDI) devices
- Run a detection test on a newly onboarded Microsoft Defender ATP device
- Troubleshoot Microsoft Defender Advanced Threat Protection onboarding issues

Onboard Windows 10 devices using a local script

11/2/2020 • 2 minutes to read • Edit Online

Applies to:

- Microsoft 365 Endpoint data loss prevention (DLP)

You can also manually onboard individual devices to Microsoft 365 Endpoint data loss prevention. You might want to do this first when testing the service before you commit to onboarding all devices in your network.

> **IMPORTANT**
>
> This script has been optimized for use on up to 10 devices.
>
> To deploy at scale, use other deployment options. For example, you can deploy an onboarding script to more than 10 devices in production with the script available in Onboard Windows 10 devices using Group Policy.

Onboard devices

1. Open the GP configuration package .zip file (*DeviceComplianceOnboardingPackage.zip*) that you downloaded from the service onboarding wizard. You can also get the package from Microsoft Compliance center
2. In the navigation pane, select **Settings** > **Device onboarding**.
3. In the **Deployment method** field, select **Local Script**.
4. Click **Download package** and save the .zip file.
5. Extract the contents of the configuration package to a location on the device you want to onboard (for example, the Desktop). You should have a file named *DeviceOnboardingScript.cmd*.
6. Open an elevated command-line prompt on the device and run the script:
7. Go to **Start** and type **cmd**.
8. Right-click **Command prompt** and select **Run as administrator**.

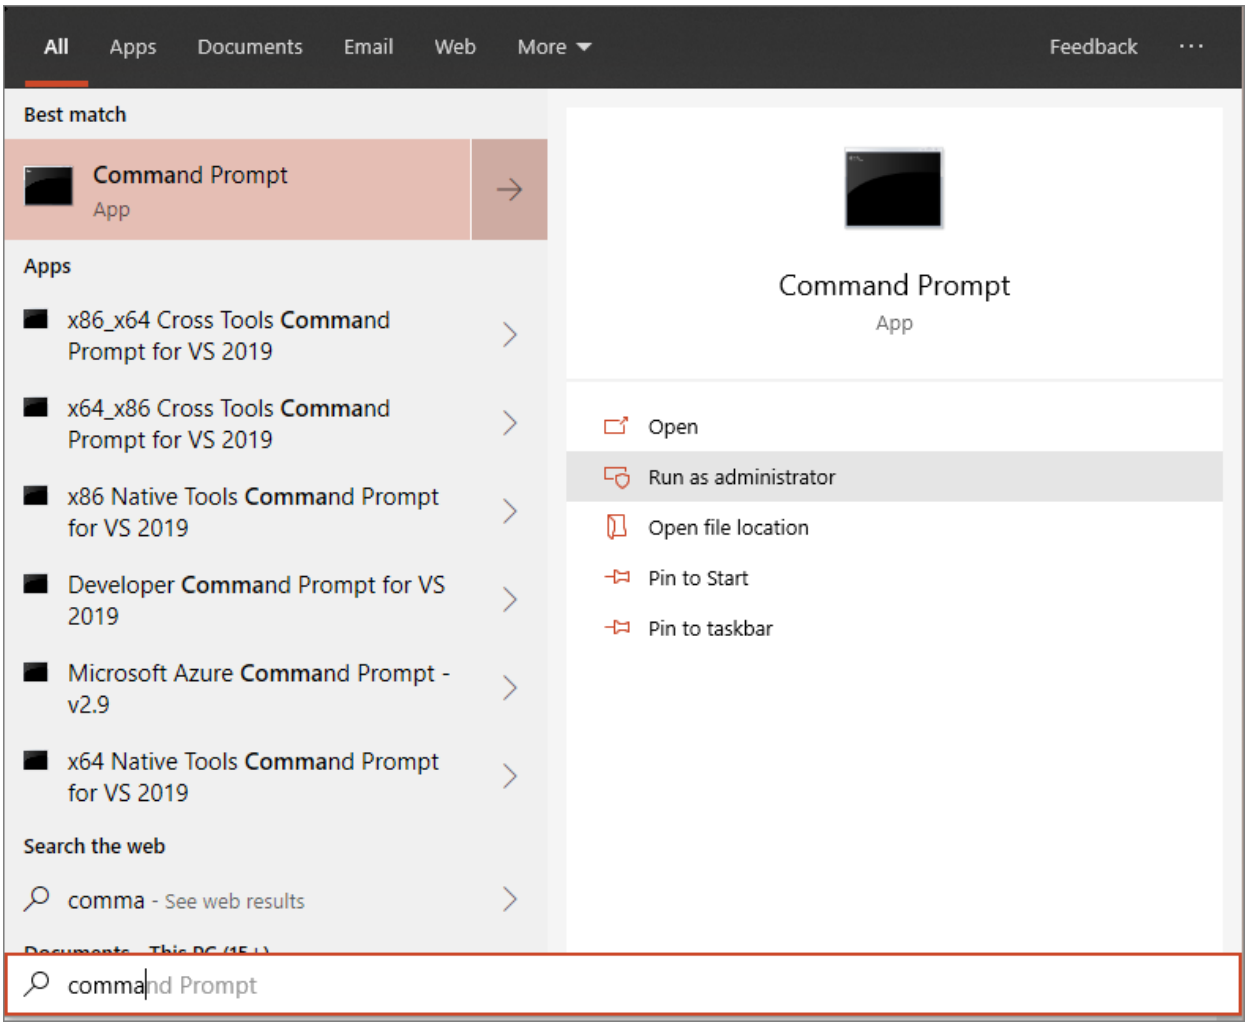

9. Type the location of the script file. If you copied the file to the desktop, type: *%userprofile%\Desktop\WindowsDefenderATPOnboardingScript.cmd*
10. Press the **Enter** key or click **OK**.

For information on how you can manually validate that the device is compliant and correctly reports sensor data see, Troubleshoot Microsoft Defender Advanced Threat Protection onboarding issues.

Offboard devices using a local script

For security reasons, the package used to Offboard devices will expire 30 days after the date it was downloaded. Expired offboarding packages sent to an device will be rejected. When downloading an offboarding package you will be notified of the packages expiry date and it will also be included in the package name.

> **NOTE**
> Onboarding and offboarding policies must not be deployed on the same device at the same time, otherwise this will cause unpredictable collisions.

1. Get the offboarding package from Microsoft Compliance center
2. In the navigation pane, select **Settings** > **Device offboarding**.
3. In the **Deployment method** field, select **Local Script**.
4. Click **Download package** and save the .zip file.
5. Extract the contents of the .zip file to a shared, read-only location that can be accessed by the devices. You

should have a file named *DeviceComplianceOffboardingScript_valid_until_YYYY-MM-DD.cmd*.

6. Open an elevated command-line prompt on the device and run the script:
7. Go to **Start** and type **cmd**.
8. Right-click **Command prompt** and select **Run as administrator**.

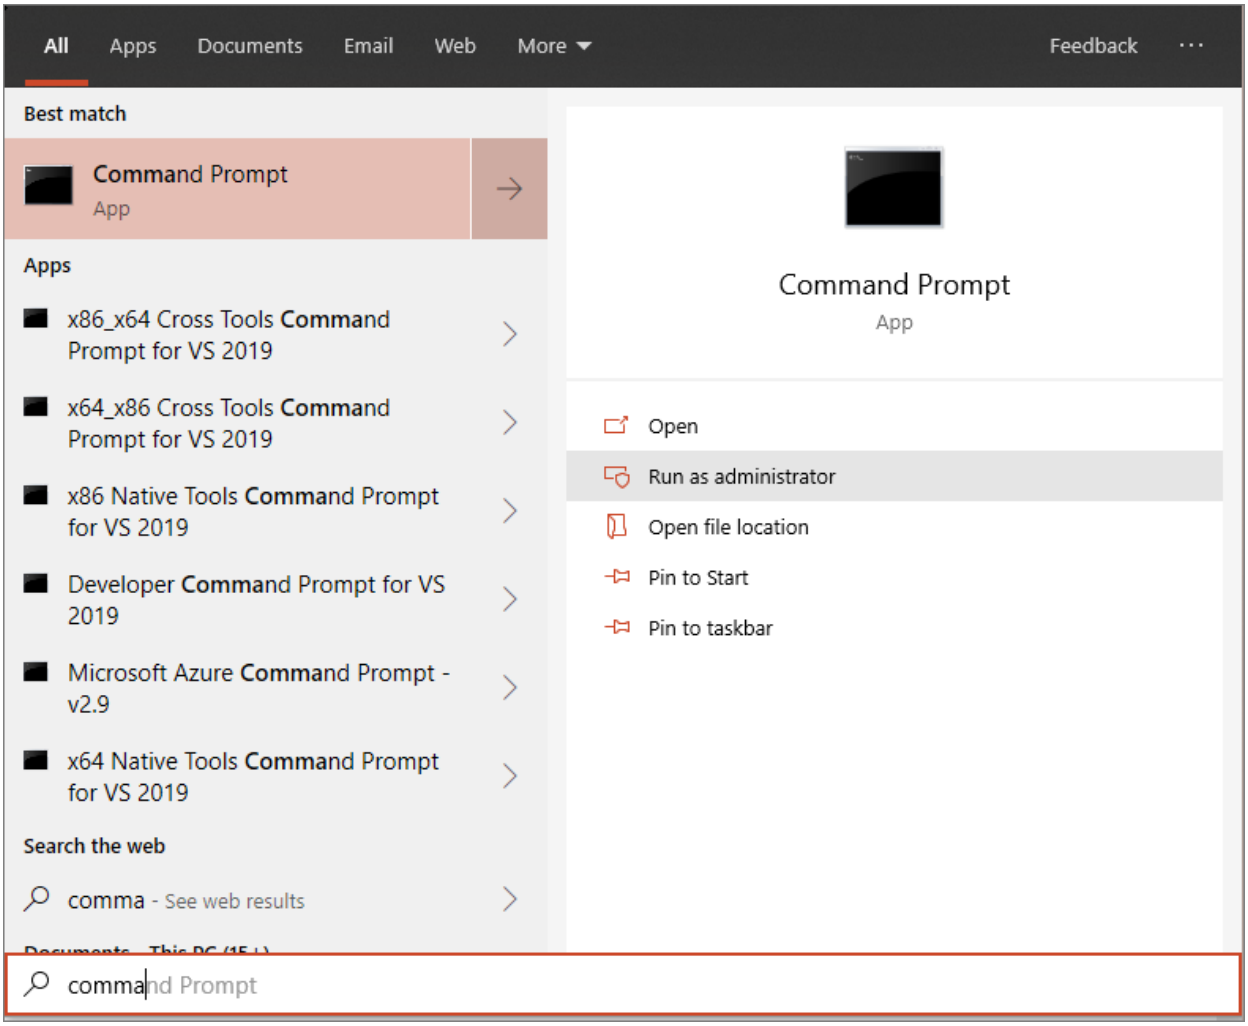

9. Type the location of the script file. If you copied the file to the desktop, type:
%userprofile%\Desktop\WindowsDefenderATPOffboardingScript_valid_until_YYYY-MM-DD.cmd
10. Press the **Enter** key or click **OK**.

> **IMPORTANT**
> Offboarding causes the device to stop sending sensor data to the portal.

Monitor device configuration

You can follow the different verification steps in the [Troubleshoot onboarding issues] ((https://docs.microsoft.com/windows/security/threat-protection/microsoft-defender-atp/troubleshoot-onboarding) to verify that the script completed successfully and the agent is running.

Monitoring can also be done directly on the portal, or by using the different deployment tools.

Monitor devices using the portal

1. Go to Microsoft 365 Compliance center.
2. Choose **Settings** > **Device onboarding** > **Devices**.

3. Verify that devices are appearing.

Related topics

- Onboard Windows 10 devices using Group Policy
- Onboard Windows 10 devices using Microsoft Endpoint Configuration Manager
- Onboard Windows 10 devices using Mobile Device Management tools
- Onboard non-persistent virtual desktop infrastructure (VDI) devices
- Run a detection test on a newly onboarded Microsoft Defender ATP device
- Troubleshoot Microsoft Defender Advanced Threat Protection onboarding issues

Onboard non-persistent virtual desktop infrastructure (VDI) devices

11/2/2020 • 3 minutes to read • Edit Online

Applies to:

- Microsoft 365 Endpoint data loss prevention (DLP)
- Virtual desktop infrastructure (VDI) devices

> **WARNING**
> Microsoft 365 Endpoint data loss prevention support for Windows Virtual Desktop supports single session scenarios. Multi-session scenarios on Windows Virtual Desktop are currently not supported.

Onboard VDI devices

Microsoft 365 Endpoint data loss prevention supports non-persistent VDI session onboarding.

> **NOTE**
> To onboard non-persistent VDI sessions, VDI devices must be on Windows 10 1809 or higher.

There might be associated challenges when onboarding VDIs. The following are typical challenges for this scenario:

- Instant early onboarding of a short-lived sessions, which must be onboarded to Microsoft 365 Endpoint data loss prevention prior to the actual provisioning.
- The device name is typically reused for new sessions.

VDI devices can appear in the Microsoft 365 Compliance center as either:

- Single entry for each device.
  Note that in this case, the *same* device name must be configured when the session is created, for example using an unattended answer file.
- Multiple entries for each device - one for each session.

The following steps will guide you through onboarding VDI devices and will highlight steps for single and multiple entries.

> **WARNING**
> For environments where there are low resource configurations, the VDI boot procedure might slow the Microsoft 365 Endpoint data loss prevention onboarding.

1. Open the VDI configuration package .zip file (*DeviceCompliancePackage.zip*) that you downloaded from the service onboarding wizard.
2. In the navigation pane, select **Settings** > **Device onboarding** > **Onboarding**.

3. In the **Deployment method** field, select **VDI onboarding scripts for non-persistent endpoints**.
4. Click **Download package** and save the .zip file.
5. Copy the files from the DeviceCompliancePackage folder extracted from the .zip file into the `golden/master` image under the path `C:\WINDOWS\System32\GroupPolicy\Machine\Scripts\Startup`.
6. If you are not implementing a single entry for each device, copy DeviceComplianceOnboardingScript.cmd.
7. If you are implementing a single entry for each device, copy both Onboard-NonPersistentMachine.ps1 and DeviceComplianceOnboardingScript.cmd.

   > **NOTE**
   > If you don't see the `C:\WINDOWS\System32\GroupPolicy\Machine\Scripts\Startup` folder, it might be hidden. You'll need to choose the **Show hidden files and folders** option from File Explorer.

8. Open a Local Group Policy Editor window and navigate to **Computer Configuration** > **Windows Settings** > **Scripts** > **Startup**.

   > **NOTE**
   > Domain Group Policy may also be used for onboarding non-persistent VDI devices.

9. Depending on the method you'd like to implement, follow the appropriate steps:

   For single entry for each device

   Select the **PowerShell Scripts** tab, then click **Add** (Windows Explorer will open directly in the path where you copied the onboarding script earlier). Navigate to onboarding PowerShell script `Onboard-NonPersistentMachine.ps1`.

   For multiple entries for each device:

   Select the **Scripts** tab, then click **Add** (Windows Explorer will open directly in the path where you copied the onboarding script earlier). Navigate to the onboarding bash script `DeviceComplianceOnboardingScript.cmd`.

10. Test your solution:
    a. Create a pool with one device.
    b. Logon to device.
    c. Logoff from device.
    d. Logon to device with another user.
    e. **For single entry for each device**: Check only one entry in Microsoft Defender Security Center.
    For multiple entries for each device: Check multiple entries in Microsoft Defender Security Center.
11. Click **Devices list** on the Navigation pane.
12. Use the search function by entering the device name and select **Device** as search type.

Updating non-persistent virtual desktop infrastructure (VDI) images

As a best practice, we recommend using offline servicing tools to patch golden/master images.
For example, you can use the below commands to install an update while the image remains offline:

```
DISM /Mount-image /ImageFile:"D:\Win10-1909.vhdx" /index:1 /MountDir:"C:\Temp\OfflineServicing"
DISM /Image:"C:\Temp\OfflineServicing" /Add-Package /Packagepath:"C:\temp\patch\windows10.0-kb4541338-x64.msu"
DISM /Unmount-Image /MountDir:"C:\Temp\OfflineServicing" /commit
```

For more information on DISM commands and offline servicing, please refer to the articles below:

- Modify a Windows image using DISM
- DISM Image Management Command-Line Options
- Reduce the Size of the Component Store in an Offline Windows Image

If offline servicing is not a viable option for your non-persistent VDI environment, the following steps should be taken to ensure consistency and sensor health:

1. After booting the master image for online servicing or patching, run an offboarding script to turn off the Microsoft 365 Endpoint data loss prevention sensor. For more information, see Offboard devices using a local script.
2. Ensure the sensor is stopped by running the command below in a CMD window:

   ```
   sc query sense
   ```

3. Service the image as needed.
4. Run the below commands using PsExec.exe (which can be downloaded from https://download.sysinternals.com/files/PSTools.zip) to cleanup the cyber folder contents that the sensor may have accumulated since boot:

   ```
   PsExec.exe -s cmd.exe
   cd "C:\ProgramData\Microsoft\Windows Defender Advanced Threat Protection\Cyber"
   del *.* /f /s /q
   REG DELETE “HKLM\SOFTWARE\Microsoft\Windows Advanced Threat Protection" /v senseGuid /f
   exit
   ```

5. Re-seal the golden/master image as you normally would.

Related topics

- Onboard Windows 10 devices using Group Policy
- Onboard Windows 10 devices using Microsoft Endpoint Configuration Manager
- Onboard Windows 10 devices using Mobile Device Management tools
- Onboard Windows 10 devices using a local script
- Troubleshoot Microsoft Defender Advanced Threat Protection onboarding issues

Data loss prevention and Microsoft Teams

2/18/2021 • 6 minutes to read • Edit Online

> **NOTE**
> Data loss prevention capabilities were recently added to Microsoft Teams chat and channel messages for users licensed for Office 365 E5/A5, Microsoft 365 E5/A5, Microsoft 365 Information Protection and Governance or Office 365 Advanced Compliance. Office 365 and Microsoft 365 E3 include DLP protection for SharePoint Online, OneDrive, and Exchange Online. This also includes files that are shared through Teams because Teams uses SharePoint Online and OneDrive to share files. Support for DLP protection in Teams Chat requires E5. To learn more about licensing requirements, see Microsoft 365 Tenant-Level Services Licensing Guidance.

> **IMPORTANT**
> DLP for Teams is only supported when the user has a mailbox that is in Exchange Online

Overview of DLP for Microsoft Teams

Recently, data loss prevention (DLP) capabilities were extended to include Microsoft Teams chat and channel messages, **including private channel messages**.

If your organization has DLP, you can now define policies that prevent people from sharing sensitive information in a Microsoft Teams channel or chat session. Here are some examples of how this protection works:

- **Example 1: Protecting sensitive information in messages**. Suppose that someone attempts to share sensitive information in a Teams chat or channel with guests (external users). If you have a DLP policy defined to prevent this, messages with sensitive information that are sent to external users are deleted. This happens automatically, and within seconds, according to how your DLP policy is configured.

  > **NOTE**
  > DLP for Microsoft Teams blocks sensitive content when shared with Microsoft Teams users who have:
  > - guest access in teams and channels; or
  > - external access in meetings and chat sessions.
  >
  > DLP for external chat sessions will only work if both the sender and the receiver are in Teams Only mode and using Microsoft Teams native federation. DLP for Teams does not block messages in interop with Skype for Business or non-native federated chat sessions.

- **Example 2: Protecting sensitive information in documents**. Suppose that someone attempts to share a document with guests in a Microsoft Teams channel or chat, and the document contains sensitive information. If you have a DLP policy defined to prevent this, the document won't open for those users. Note that in this case, your DLP policy must include SharePoint and OneDrive in order for protection to be in place. (This is an example of DLP for SharePoint that shows up in Microsoft Teams, and therefore requires that users are licensed for Office 365 DLP (included in Office 365 E3), but does not require users to be licensed for Office 365 Advanced Compliance.)

Policy tips help educate users

Similar to how DLP works in Exchange, Outlook, Outlook on the web, SharePoint Online, OneDrive for Business sites, and Office desktop clients, policy tips appear when an action conflicts with a DLP policy. Here's an example of a policy tip:

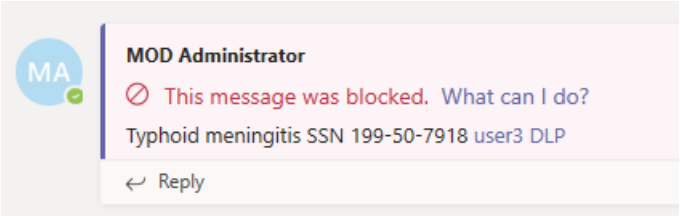

In this case, the sender attempted to share a social security number in a Microsoft Teams channel. The **What can I do?** link opens a dialog box that provides options for the sender to resolve the issue. Notice that in this case, the sender can opt to override the policy, or notify an admin to review and resolve it.

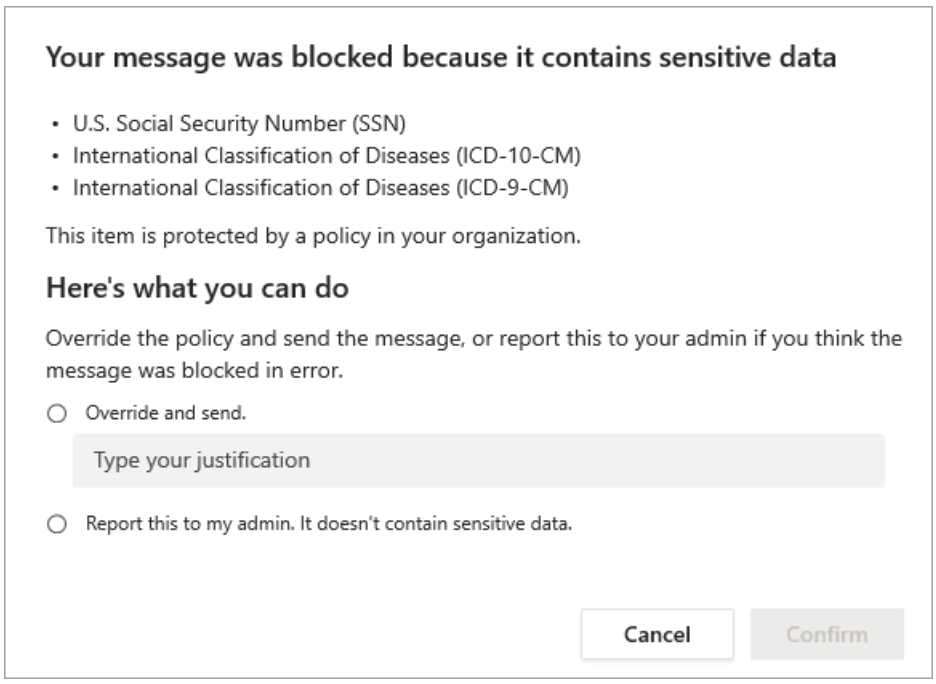

In your organization, you can choose to allow users to override a DLP policy. And, when you configure your DLP policies, you can use the default policy tips, or customize policy tips for your organization.

Returning to our example, where a sender shared a social security number in a Teams channel, here's what the recipient saw:

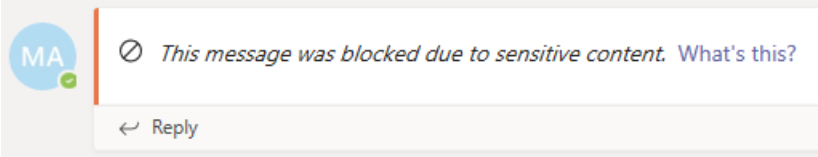

The **What's this?** link opens an article about DLP policies, which helps explain why the message was blocked.

To customize policy tips

To perform this task, you must be assigned a role that has permissions to edit DLP policies. To learn more, see Permissions.

1. Go to the Security & Compliance Center (https://protection.office.com) and sign in.
2. Choose **Data loss prevention** > **Policy**.
3. Select a policy, and next to **Policy settings**, choose **Edit**.
4. Either create a new rule, or edit an existing rule for the policy.

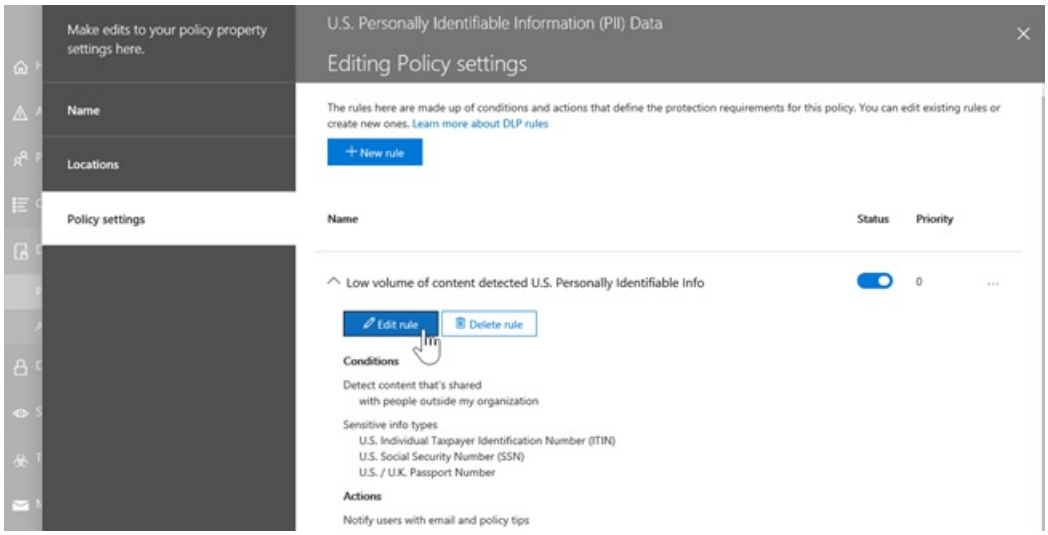

5. On the **User notifications** tab, select **Customize the email text** and/or **Customize the policy tip text** options.

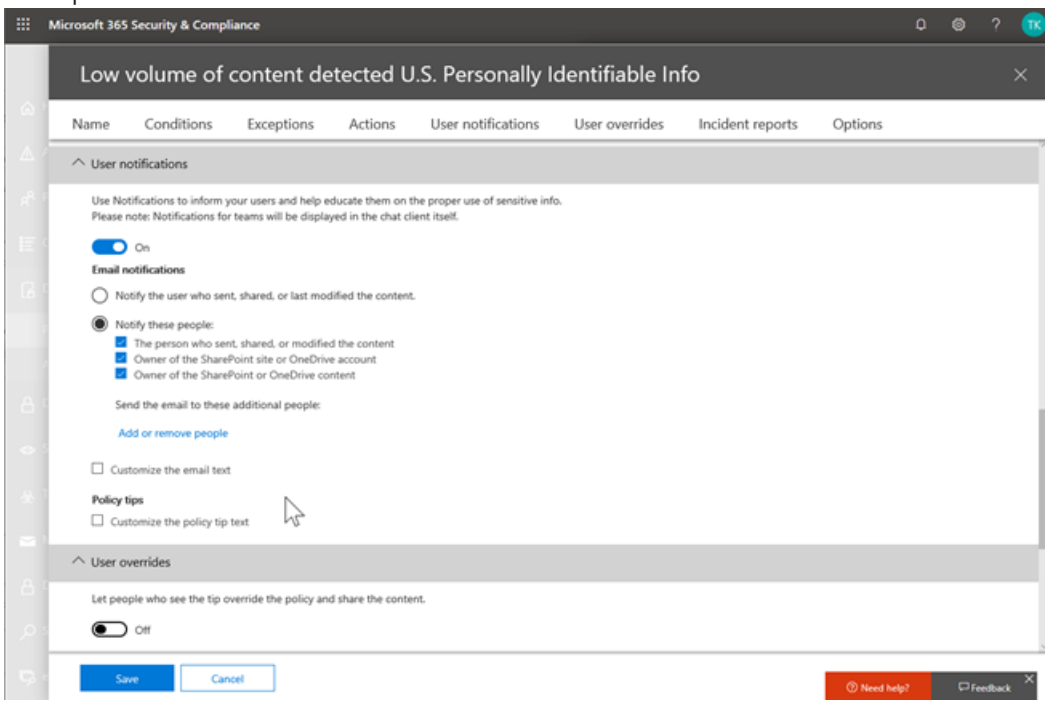

6. Specify the text you want to use for email notifications and/or policy tips, and then choose **Save**.
7. On the **Policy settings** tab, choose **Save**.

Allow approximately one hour for your changes to work their way through your data center and sync to user accounts.

Add Microsoft Teams as a location to existing DLP policies

To perform this task, you must be assigned a role that has permissions to edit DLP policies. To learn more, see Permissions.

1. Go to the Security & Compliance Center (https://protection.office.com) and sign in.
2. Choose **Data loss prevention** > **Policy**.
3. Select a policy, and look at the values under **Locations**. If you see **Teams chat and channel messages**, you're all set. If you don't, click **Edit**.

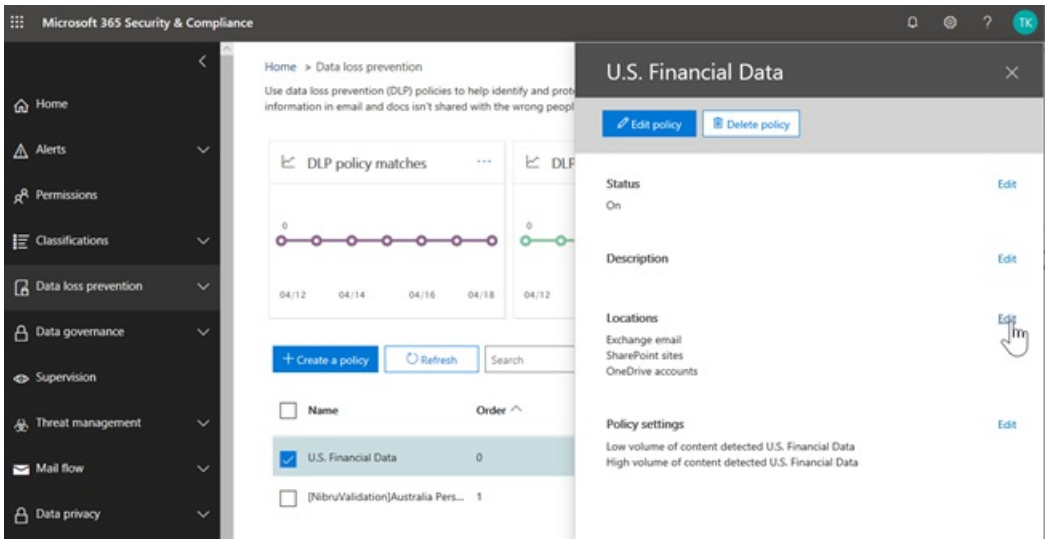

4. In the **Status** column, turn the policy on for **Teams chat and channel messages**.

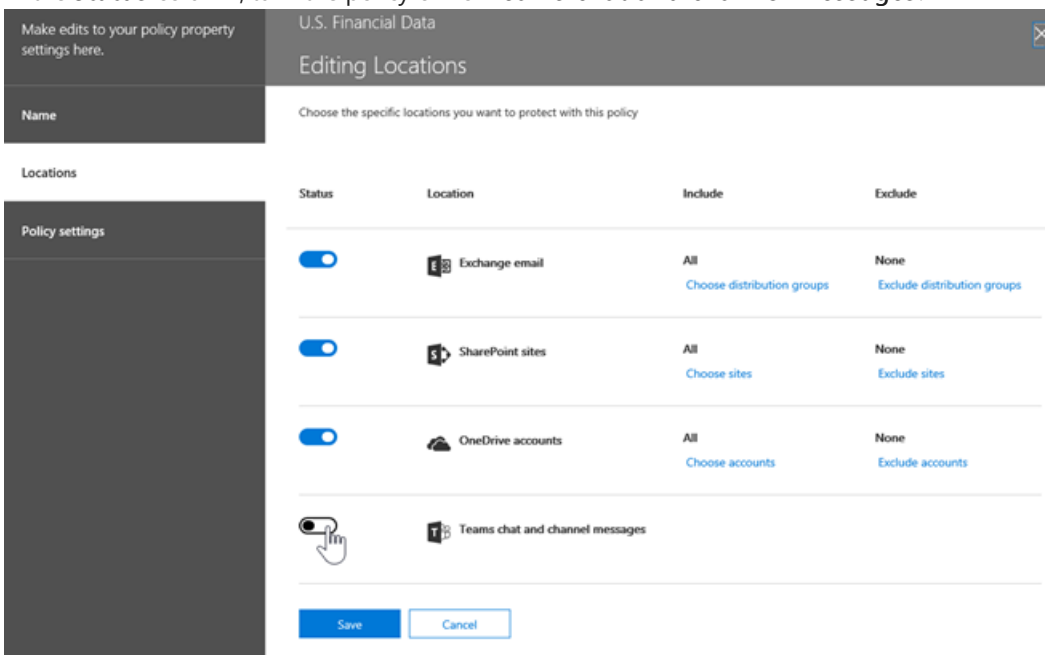

5. On the **Choose locations** tab, keep the default setting of all accounts, or select **Let me choose specific locations**. You can specify:

   a. up to 1000 individual accounts to include or exclude
   b. distribution lists and security groups to include or exclude. **This is a public preview feature.**

6. Then choose **Next**.

7. Click **Save**.

Allow approximately one hour for your changes to work their way through your data center and sync to user accounts.

Define a new DLP policy for Microsoft Teams

To perform this task, you must be assigned a role that has permissions to edit DLP policies. To learn more, see Permissions.

1. Go to the Security & Compliance Center (https://protection.office.com) and sign in.

2. Choose **Data loss prevention** > **Policy** > **+ Create a policy**.

3. Choose a template, and then choose **Next**.

In our example, we chose the U.S. Personally Identifiable Information Data template.

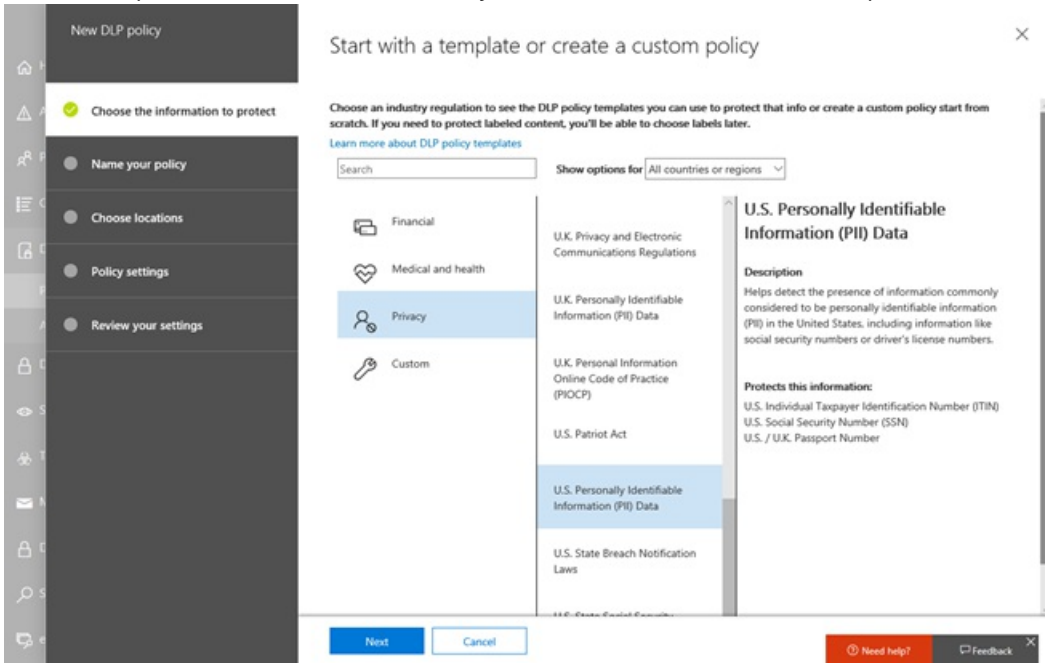

4. On the **Name your policy** tab, specify a name and description for the policy, and then choose **Next**.

5. On the **Choose locations** tab, keep the default setting of all accounts, or select **Let me choose specific locations**. You can specify:

   a. up to 1000 individual accounts to include or exclude
   b. distribution lists and security groups to include or exclude. **This is a public preview feature.**

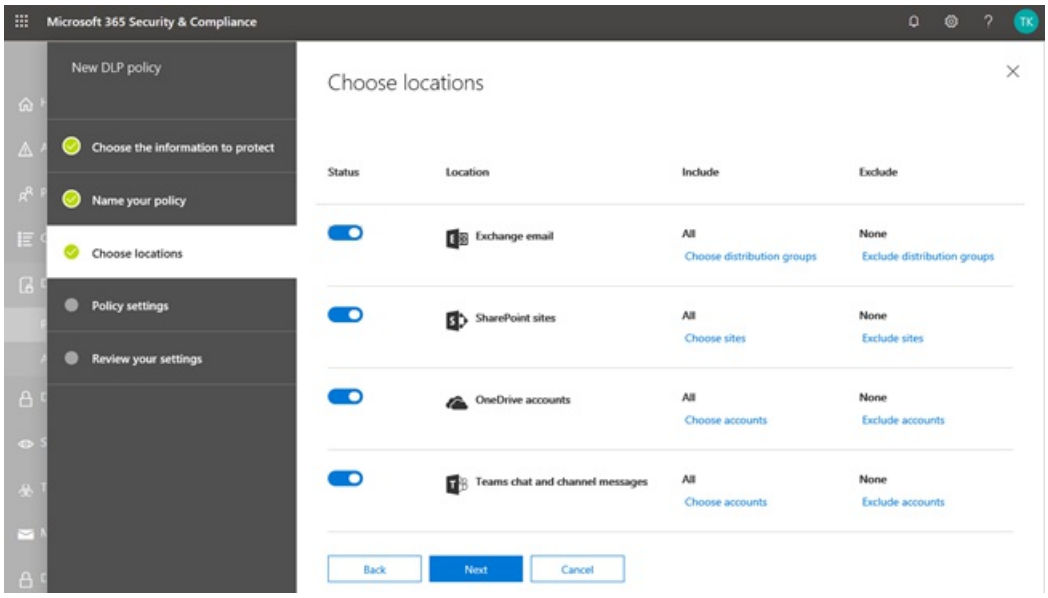

> **NOTE**
> If you want to make sure documents that contain sensitive information are not shared inappropriately in Teams, make sure **SharePoint sites** and **OneDrive accounts** are turned on, along with **Teams chat and channel messages**.

6. On the **Policy settings** tab, under **Customize the type of content you want to protect**, keep the default simple settings, or choose **Use advanced settings**, and then choose **Next**. If you choose advanced settings, you can create or edit rules for your policy. (To get help with this, see Simple settings vs. advanced settings.)

7. On the **Policy settings** tab, under **What do you want to do if we detect sensitive info?**, review the

settings. (Here's where you can choose to keep default policy tips and email notifications, or customize them.)

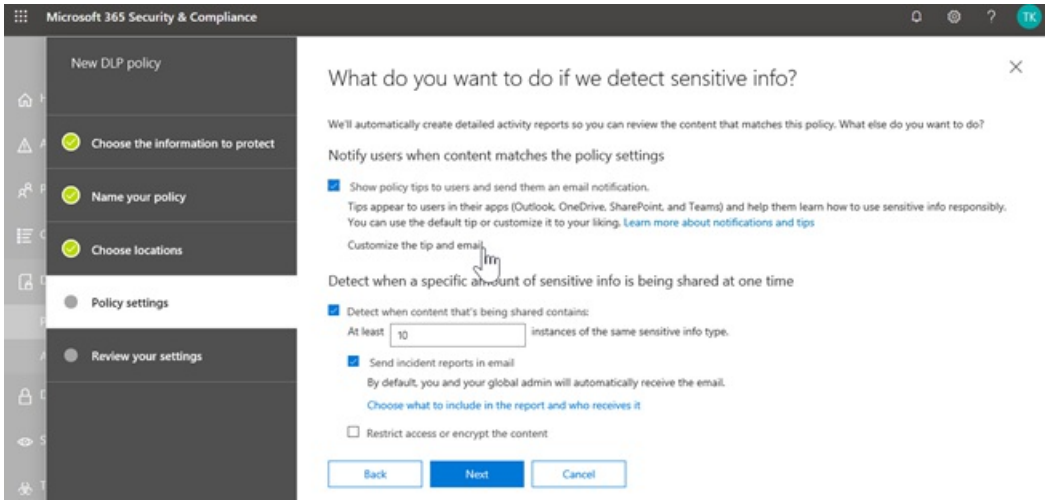

When you're finished reviewing or editing settings, choose **Next**.

8. On the **Policy settings** tab, under **Do you want to turn on the policy or test things out first?**, choose whether to turn the policy on, test it first, or keep it turned off for now, and then choose **Next**.

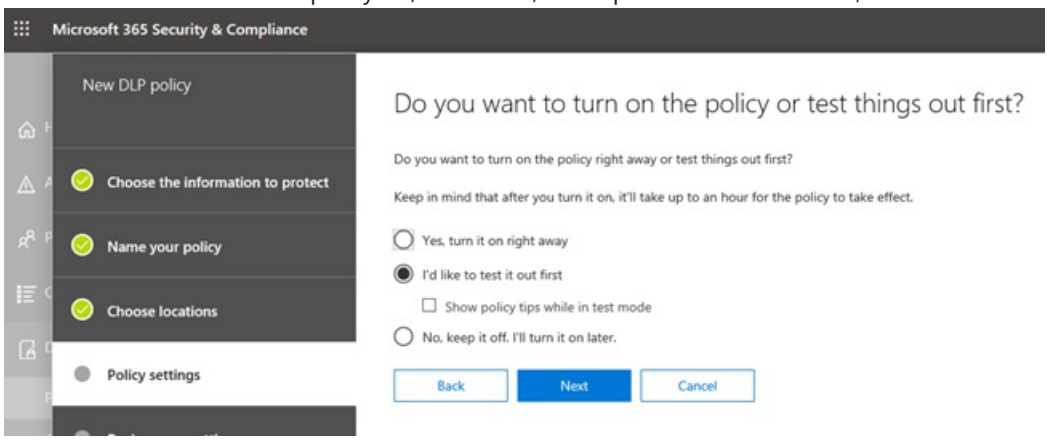

9. On the **Review your settings** tab, review the settings for your new policy. Choose **Edit** to make changes. When you're finished, choose **Create**.

Allow approximately one hour for your new policy to work its way through your data center and sync to user accounts.

Prevent external access to sensitive documents

To ensure that SharePoint documents that contain sensitive information cannot be accessed by external guests either from SharePoint or Teams by default, select the following:

- You can ensure that documents are protected until DLP scans and marks them as safe to share by marking new files as sensitive by default
- Recommended DLP policy structure
  - **Conditions**
    - Content contains any of these sensitive information types: [Select all that applies]
    - Content is shared from Microsoft 365 with people outside my organization

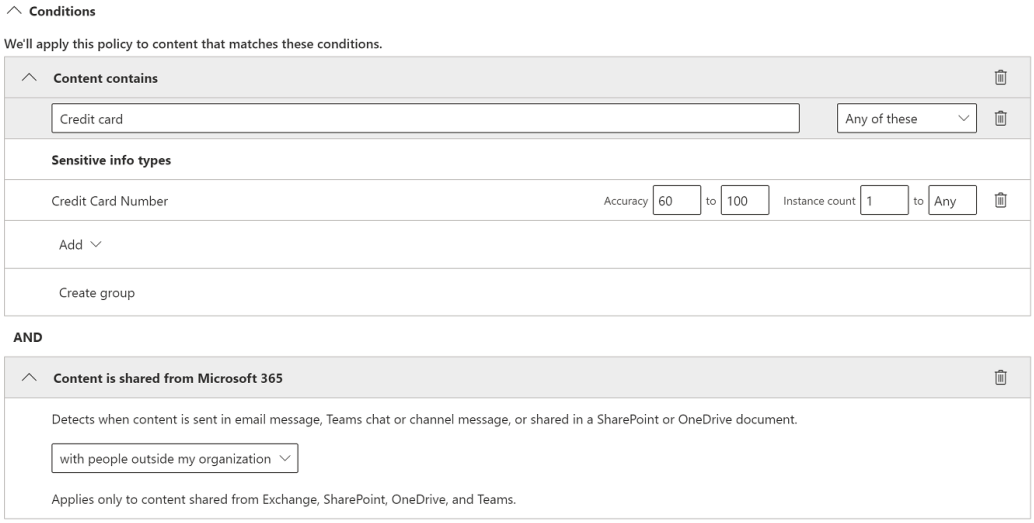

- **Actions**
  - Restrict access to the content for external users
  - Notify users with email and policy tips
  - Send incident reports to the Administrator

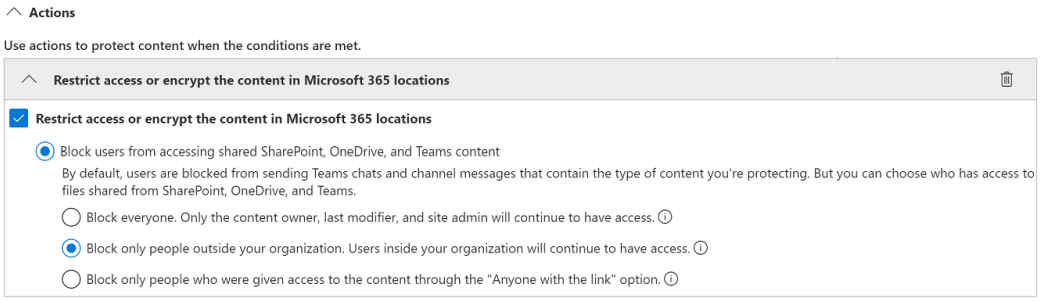

DLP policy in action when attempting to share a document in SharePoint that contains sensitive information with an external guest:

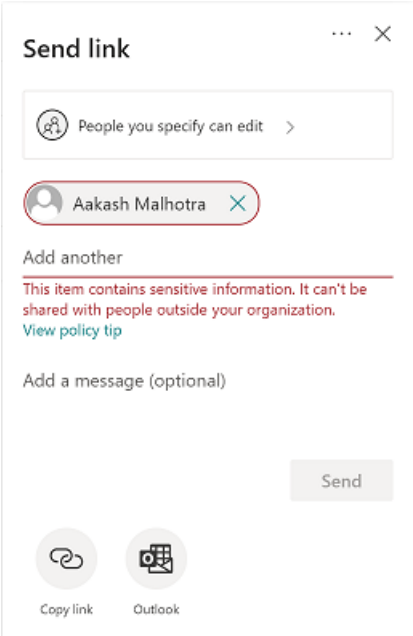

DLP policy in action when guest attempts to open a document in Teams with block external:

This link is only available to internal users

This link is not available to you.

TECHNICAL DETAILS

GO BACK TO SITE

Related articles

Create, test, and tune a DLP policy

Send email notifications and show policy tips for DLP policies

Using Endpoint data loss prevention

2/18/2021 • 7 minutes to read • Edit Online

This article walks you through three scenarios where you create and modify a DLP policy that uses devices as a location.

DLP settings

Before you get started you should set up your DLP settings which are applied to all DLP policies for devices. You must configure these if you intend to create policies that enforce:

- cloud egress restrictions
- unallowed apps restrictions

Or

- If you want to exclude noisy file paths from monitoring

File path exclusions

You may want to exclude certain paths from DLP monitoring, DLP alerting, and DLP policy enforcement on your devices because they are too noisy or don't contain files you are interested in. Files in those locations will not be audited and any files that are created or modified in those locations will not be subject to DLP policy enforcement. You can configure path exclusions in DLP settings.

You can use this logic to construct your exclusion paths:

- Valid file path that ends with '\', which means only files directly under folder.
  For example: C:\Temp\
- Valid file path that ends with '*', which means only files under sub-folders, besides the files directly under the folder.
  For example: C:\Temp*
- Valid file path that ends without '\' or '*', which means all files directly under folder and all sub-folders.
  For example: C:\Temp
- A path with wildcard between '\' from each side.
  For example: C:\Users*\Desktop\
- A path with wildcard between '\' from each side and with '(number)' to give exact number of subfolders.
  For example: C:\Users*(1)\Downloads\
- A path with SYSTEM environment variables.
  For example: %SystemDrive%\Test*
- A mix of all the above.

For example: %SystemDrive%\Users*\Documents*(2)\Sub\

Unallowed apps

When a policy's **Access by unallowed apps and browsers** setting is turned on and users attempt to use these apps to access a protected file, the activity will be allowed, blocked, or blocked but users can override the restriction. All activity is audited and available to review in activity explorer.

> **IMPORTANT**
> Do not include the path to the executable, but only the executable name (such as browser.exe).

Browser and domain restrictions

Restrict sensitive files that match your policies from being shared with unrestricted cloud service domains.

Service domains

You can control whether sensitive files protected by your policies can be uploaded to specific service domains from Microsoft Edge.

If the list mode is set to **Block**, then user will not be able to upload sensitive items to those domains. When an upload action is blocked because an item matches a DLP policy, DLP will either generate a warning or block the upload of the sensitive item.

If the list mode is set to **Allow**, then users will be able to upload sensitive items ***only*** to those domains, and upload access to all other domains is not allowed.

Unallowed browsers

You add browsers, identified by their executable names, that will be blocked from accessing files that match the conditions of an enforced a DLP policy where the upload to cloud services restriction is set to block or block override. When these browsers are blocked from accessing a file, the end users will see a toast notification asking them to open the file through Edge Chromium.

Business justification in policy tips

You can control how users interact with the business justification option in DLP policy tip notifications. This option appears when users perform an activity that's protected by the **Block with override** setting in a DLP policy. You can choose from one the following options:

- By default, users can select either a built-in justification, or enter their own text.
- Users can only select a built-in justification.
- Users can only enter their own justification.

Tying DLP settings together

With Endpoint DLP and Edge Chromium Web browser, you can restrict unintentional sharing of sensitive items to unallowed cloud apps and services. Edge Chromium understands when an item is restricted by an Endpoint DLP policy and enforces access restrictions.

When you use Endpoint DLP as a location in a properly configured DLP policy and the Edge Chromium browser, the unallowed browsers that you've defined in these settings will be prevented from accessing the sensitive items that match your DLP policy controls. Instead, users will be redirected to use Edge Chromium and Edge Chromium, with its understanding of DLP imposed restrictions, can block or restrict activities when the conditions in the DLP policy are met.

To use this restriction you'll need to configure three important pieces:

1. Specify the places – services, domains, IP addresses – that you want to prevent sensitive items from being shared to.

2. Add the browsers that aren't allowed to access certain sensitive items when a DLP policy match occurs.
3. Configure DLP policies to define the kinds of sensitive items for which upload should be restricted to these places by turning on **Upload to cloud services** and **Access from unallowed browser**.

You can continue to add new services, apps, and policies to extend and augment your restrictions to meet your business needs and protect sensitive data.

This configuration will help ensure your data remains safe while also avoiding unnecessary restrictions that prevent or restrict users from accessing and sharing non-sensitive items.

Endpoint DLP policy scenarios

To help familiarize you with Endpoint DLP features and how they surface in DLP policies, we've put together some scenarios for you to follow. All the Endpoint DLP content will be folded in to the main DLP content set when Endpoint DLP becomes generally available.

> **IMPORTANT**
> These Endpoint DLP scenarios are not the official procedures for creating and tuning DLP policies. Refer to the below topics when you need to work with DLP policies in general situations:
>
> - Overview of data loss prevention
> - Get started with the default DLP policy
> - Create a DLP policy from a template
> - Create, test, and tune a DLP policy

Scenario 1: Create a policy from a template, audit only

These scenarios require that you already have devices onboarded and reporting into Activity explorer. If you haven't onboarded devices yet, see Get started with Endpoint data loss prevention.

1. Open the Data loss prevention page.
2. Choose **Create policy**.
3. For this scenario, choose **Privacy**, then **U.S. Personally Identifiable Information (PII) Data** and choose **Next**.
4. Toggle the **Status** field to off for all locations except **Devices**. Choose **Next**.
5. Accept the default **Review and customize settings from the template** selection and choose **Next**.
6. Accept the default **Protection actions** values and choose **Next**.
7. Select **Audit or restrict activities on Windows devices** and leave the actions set to **Audit only**. Choose **Next**.
8. Accept the default **I'd like to test it out first** value and choose **Show policy tips while in test mode**. Choose **Next**.
9. Review your settings and choose **Submit**.
10. The new DLP policy will appear in the policy list.
11. Check Activity explorer for data from the monitored endpoints. Set the location filter for devices and add the policy, then filter by policy name to see the impact of this policy. See, Get started with activity explorer if needed.
12. Attempt to share a test that contains content that will trigger the U.S. Personally Identifiable Information

(PII) Data condition with someone outside your organization. This should trigger the policy.

13. Check Activity explorer for the event.

Scenario 2: Modify the existing policy, set an alert

1. Open the Data loss prevention page.
2. Choose the **U.S. Personally Identifiable Information (PII) Data** policy that you created in scenario 1.
3. Choose **edit policy**.
4. Go to the **Advanced DLP rules** page and edit the **Low volume of content detected U.S. Personally Identifiable Inf**.
5. Scroll down to the **Incident reports** section and set **Send an alert to admins when a rule match occurs** to **On**. Email alerts will be automatically sent to the administrator and anyone else you add to the list of recipients.

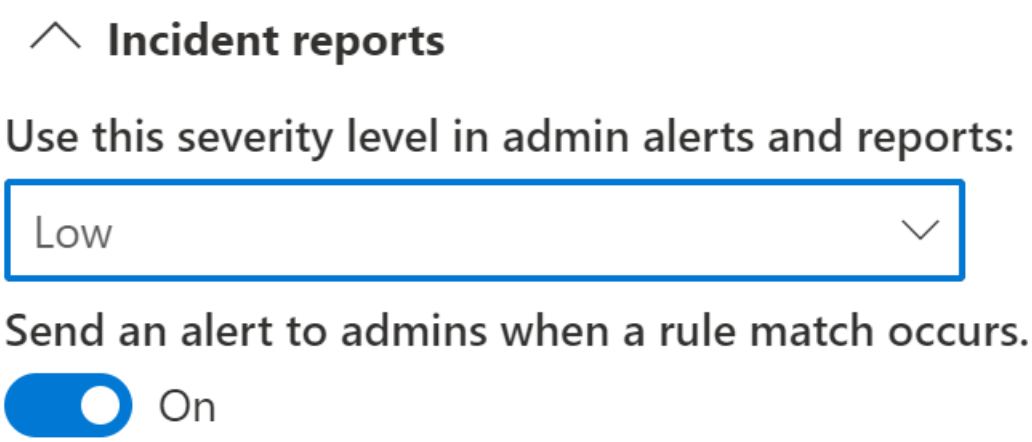

6. For the purposes of this scenario, choose **Send alert every time an activity matches the rule**.
7. Choose **Save**.
8. Retain all your previous settings by choosing **Next** and then **Submit** the policy changes.
9. Attempt to share a test that contains content that will trigger the U.S. Personally Identifiable Information (PII) Data condition with someone outside your organization. This should trigger the policy.
10. Check Activity explorer for the event.

Scenario 3: Modify the existing policy, block the action with allow override

1. Open the Data loss prevention page.
2. Choose the **U.S. Personally Identifiable Information (PII) Data** policy that you created in scenario 1.
3. Choose **edit policy**.
4. Go to the **Advanced DLP rules** page and edit the **Low volume of content detected U.S. Personally Identifiable Inf**.
5. Scroll down to the **Audit or restrict activities on Windows device** section and for each activity set the corresponding action to **Block with override**.

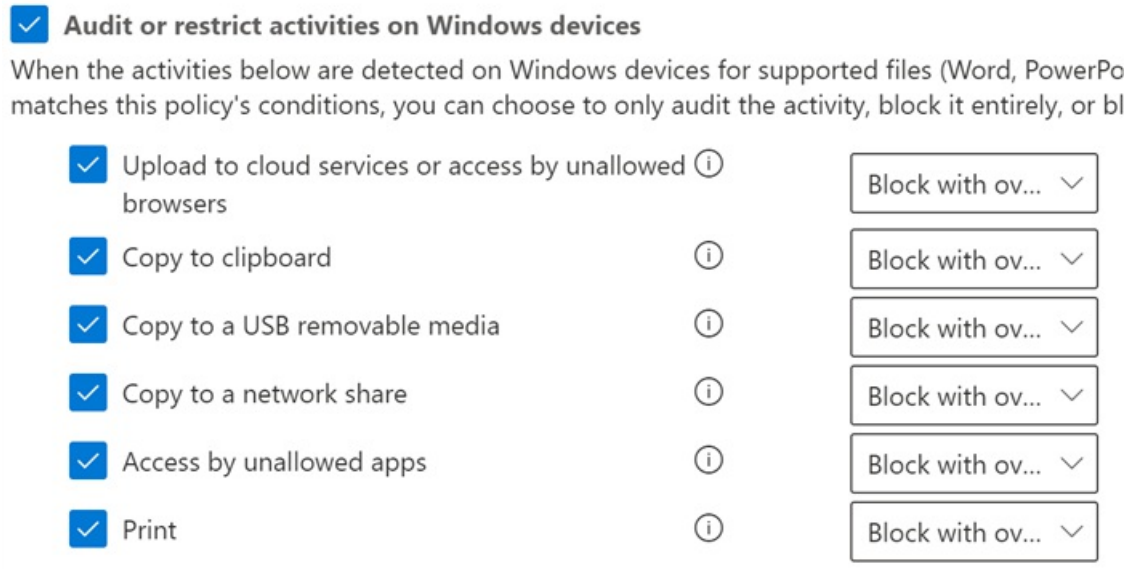

6. Choose **Save**.
7. Repeat steps 4-7 for the **High volume of content detected U.S. Personally Identifiable Inf**.
8. Retain all your previous settings by choosing **Next** and then **Submit** the policy changes.
9. Attempt to share a test that contains content that will trigger the U.S. Personally Identifiable Information (PII) Data condition with someone outside your organization. This should trigger the policy.

   You'll see a popup like this on the client device:

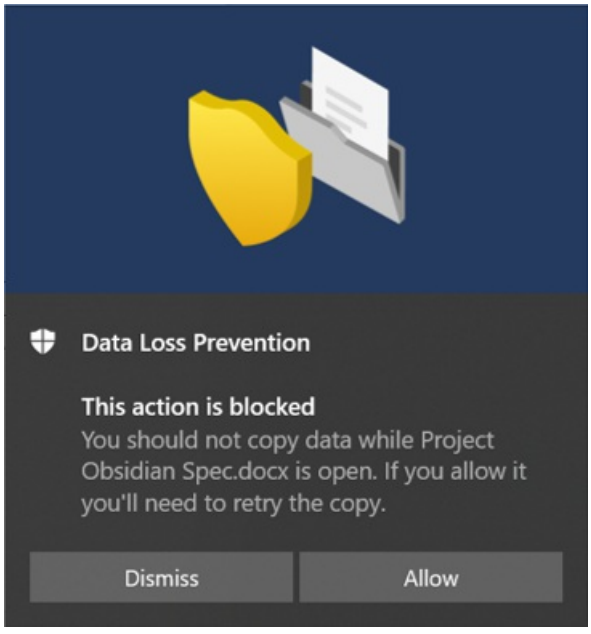

10. Check Activity explorer for the event.

See also

- Learn about Endpoint data loss prevention
- Get started with Endpoint data loss prevention
- Overview of data loss prevention
- Create, test, and tune a DLP policy
- Get started with Activity explorer
- Microsoft Defender for Endpoint
- Onboarding tools and methods for Windows 10 machines

- Microsoft 365 subscription
- Azure Active Directory (AAD) joined
- Download the new Microsoft Edge based on Chromium
- Get started with the default DLP policy
- Create a DLP policy from a template

Send email notifications and show policy tips for DLP policies

2/18/2021 • 15 minutes to read • Edit Online

You can use a data loss prevention (DLP) policy to identify, monitor, and protect sensitive information across Office 365. You want people in your organization who work with this sensitive information to stay compliant with your DLP policies, but you don't want to block them unnecessarily from getting their work done. This is where email notifications and policy tips can help.

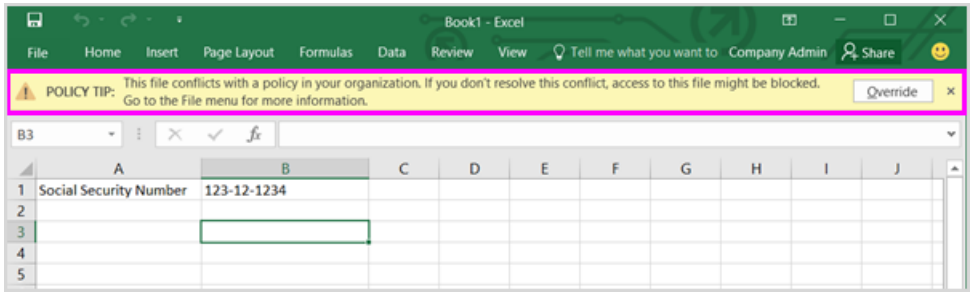

A policy tip is a notification or warning that appears when someone is working with content that conflicts with a DLP policy—for example, content like an Excel workbook on a OneDrive for Business site that contains personally identifiable information (PII) and is shared with an external user.

You can use email notifications and policy tips to increase awareness and help educate people about your organization's policies. You can also give people the option to override the policy, so that they're not blocked if they have a valid business need or if the policy is detecting a false positive.

In the Security & Compliance Center, when you create a DLP policy, you can configure the user notifications to:

- Send an email notification to the people you choose that describes the issue.
- Display a policy tip for content that conflicts with the DLP policy:
  - For email in Outlook on the web and Outlook 2013 and later, the policy tip appears at the top of a message above the recipients while the message is being composed.
  - For documents in a OneDrive for Business account or SharePoint Online site, the policy tip is indicated by a warning icon that appears on the item. To view more information, you can select an item and then choose **Information** in the upper-right corner of the page to open the details pane.
  - For Excel, PowerPoint, and Word documents that are stored on a OneDrive for Business site or SharePoint Online site that's included in the DLP policy, the policy tip appears on the Message Bar and the Backstage view (**File** menu > **Info**).

Add user notifications to a DLP policy

When you create a DLP policy, you can enable **User notifications**. When user notifications are enabled, Microsoft 365 sends out both email notifications and policy tips. You can customize who notification emails are sent to, the email text and the policy tip text.

1. Go to https://protection.office.com.
2. Sign in using your work or school account. You're now in the Security & Compliance Center.

3. In the Security & Compliance Center > left navigation > **Data loss prevention** > **Policy** > **+ Create a policy**.

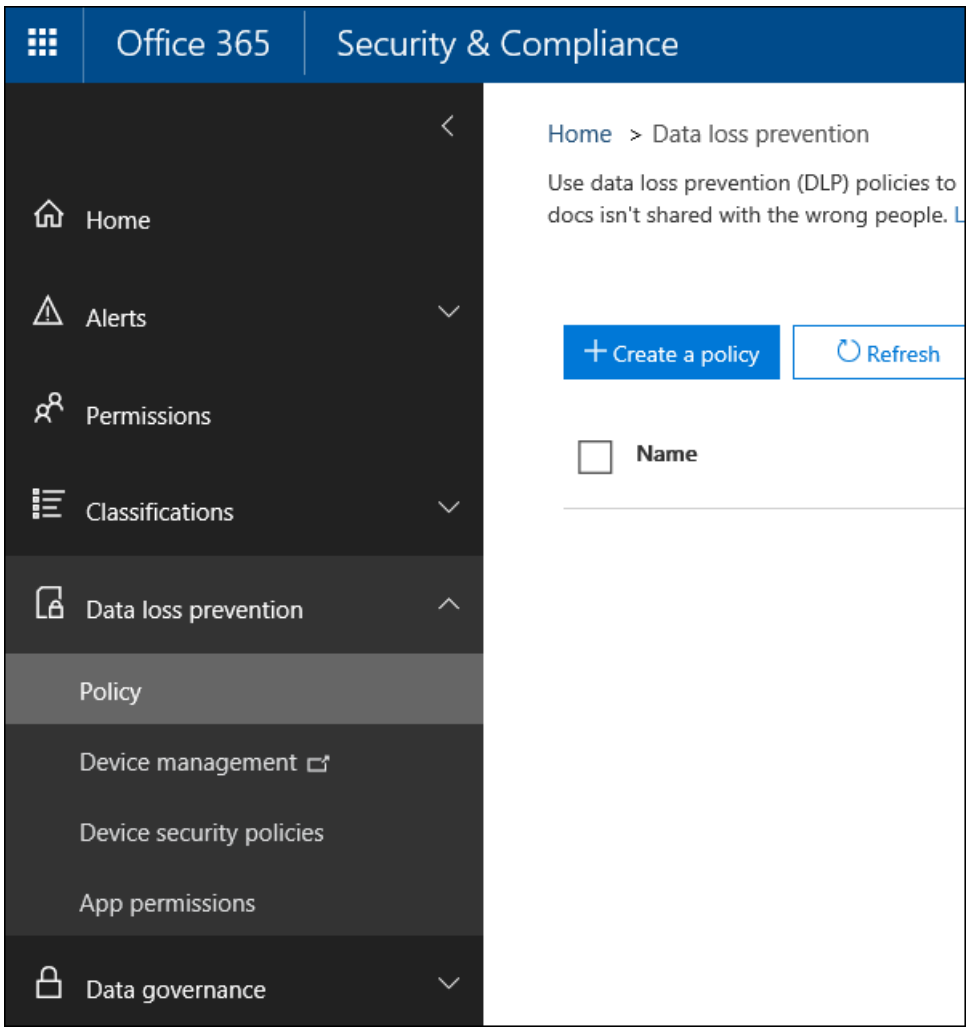

4. Choose the DLP policy template that protects the types of sensitive information that you need > **Next**.

   To start with an empty template, choose **Custom** > **Custom policy** > **Next**.

5. Name the policy > **Next**.

6. To choose the locations that you want the DLP policy to protect, do one of the following:

   - Choose **All locations in Office 365** > **Next**.
   - Choose **Let me choose specific locations** > **Next**.

   To include or exclude an entire location such as all Exchange email or all OneDrive accounts, switch the **Status** of that location on or off.

   To include only specific SharePoint sites or OneDrive accounts, switch the **Status** to on, and then click the links under **Include** to choose specific sites or accounts.

7. Choose **Use advanced settings** > **Next**.

8. Choose **+ New rule**.

9. In the rule editor, under **User notifications**, switch the status on.

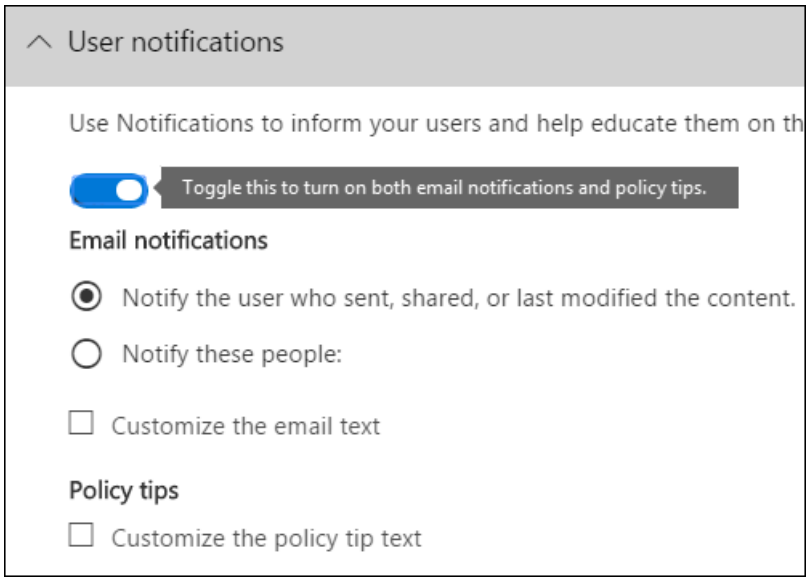

NOTE

DLP policies apply to all documents that match the policy, whether those documents are new or existing. However, an email notification is only generated when new content matches an existing DLP policy. Existing content is protected, but will not generate a user notification via email.

Options for configuring email notifications

For each rule in a DLP policy, you can:

- Send the notification to the people you choose. These people can include the owner of the content, the person who last modified the content, the owner of the site where the content is stored, or a specific user.
- Customize the text that's included in the notification by using HTML or tokens. See the section below for more information.

NOTE

Email notifications can be sent only to individual recipients—not groups or distribution lists. Only new content will trigger an email notification. Editing existing content will trigger policy tips, but not an email notification.

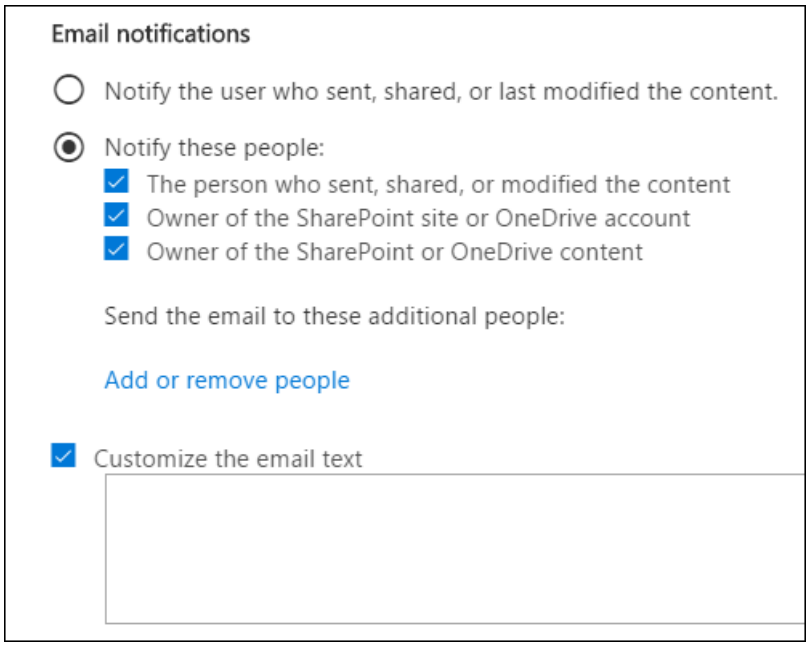

Default email notification

Notifications have a Subject line that begins with the action taken, such as "Notification", "Message Blocked" for email, or "Access Blocked" for documents. If the notification is about a document, the notification message body includes a link that takes you to the site where the document's stored and opens the policy tip for the document, where you can resolve any issues (see the section below about policy tips). If the notification is about a message, the notification includes as an attachment the message that matches a DLP policy.

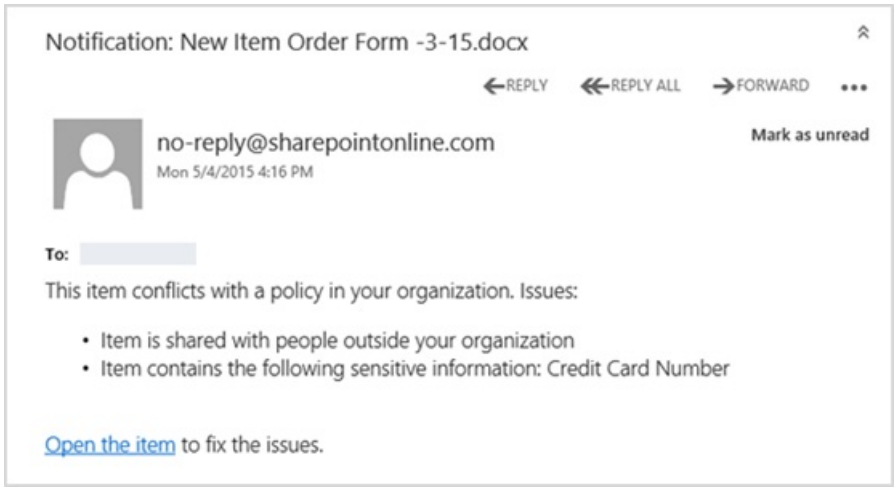

By default, notifications display text similar to the following for an item on a site. The notification text is configured separately for each rule, so the text that's displayed differs depending on which rule is matched.

IF THE DLP POLICY RULE DOES THIS...	THEN THE DEFAULT NOTIFICATION FOR SHAREPOINT OR ONEDRIVE FOR BUSINESS DOCUMENTS SAYS THIS...	THEN THE DEFAULT NOTIFICATION FOR OUTLOOK MESSAGES SAYS THIS...
Sends a notification but doesn't allow override	This item conflicts with a policy in your organization.	Your email message conflicts with a policy in your organization.
Blocks access, sends a notification, and allows override	This item conflicts with a policy in your organization. If you don't resolve this conflict, access to this file might be blocked.	Your email message conflicts with a policy in your organization. The message wasn't delivered to all recipients.
Blocks access and sends a notification	This item conflicts with a policy in your organization. Access to this item is blocked for everyone except its owner, last modifier, and the primary site collection administrator.	Your email message conflicts with a policy in your organization. The message wasn't delivered to all recipients.

Custom email notification

You can create a custom email notification instead of sending the default email notification to your end users or admins. The custom email notification supports HTML and has a 5,000-character limit. You can use HTML to include images, formatting, and other branding in the notification.

You can also use the following tokens to help customize the email notification. These tokens are variables that are replaced by specific information in the notification that's sent.

TOKEN	DESCRIPTION
%%AppliedActions%%	The actions applied to the content.
%%ContentURL%%	The URL of the document on the SharePoint Online site or OneDrive for Business site.

TOKEN	DESCRIPTION
%%MatchedConditions%%	The conditions that were matched by the content. Use this token to inform people of possible issues with the content.

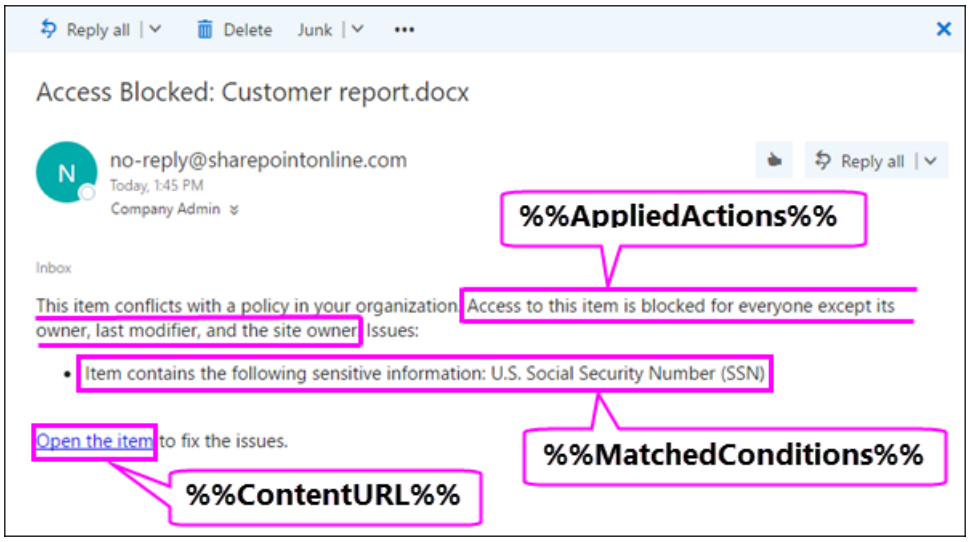

Options for configuring policy tips

For each rule in a DLP policy, you can configure policy tips to:

- Simply notify the person that the content conflicts with a DLP policy, so that they can take action to resolve the conflict. You can use the default text (see the tables below) or enter custom text about your organization's specific policies.
- Allow the person to override the DLP policy. Optionally, you can:
  - Require the person to enter a business justification for overriding the policy. This information is logged and you can view it in the DLP reports in the **Reports** section of the Security & Compliance Center.
  - Allow the person to report a false positive and override the DLP policy. This information is also logged for reporting, so that you can use false positives to fine tune your rules.

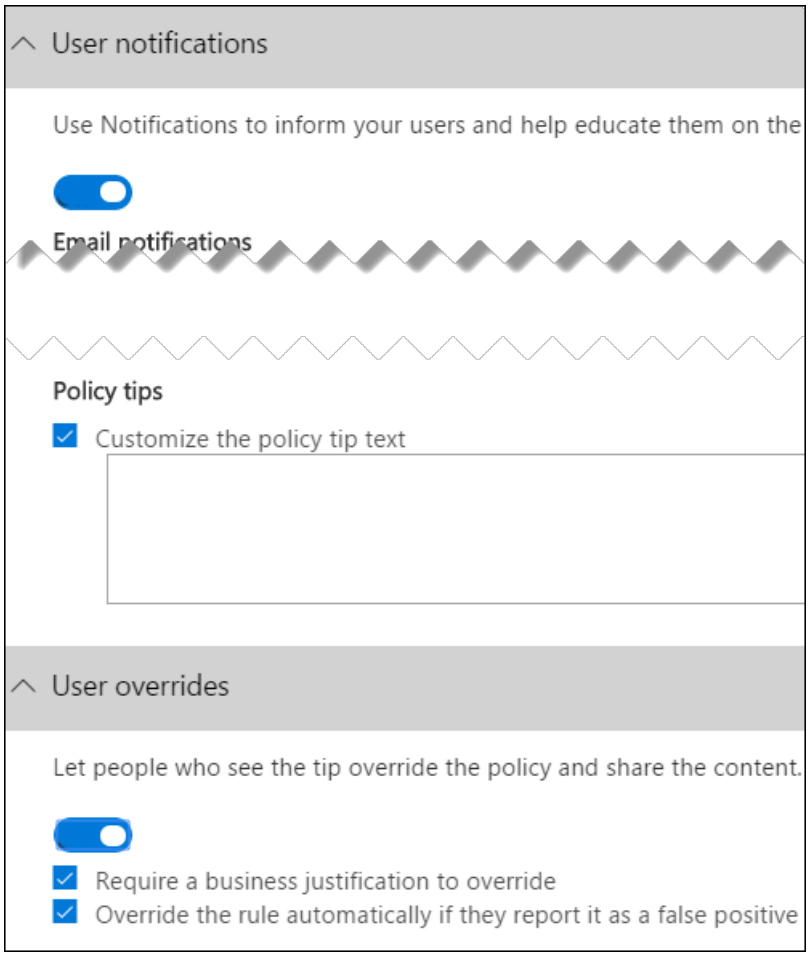

For example, you may have a DLP policy applied to OneDrive for Business sites that detects personally identifiable information (PII), and this policy has three rules:

1. First rule: If fewer than five instances of this sensitive information are detected in a document, and the document is shared with people inside the organization, the **Send a notification** action displays a policy tip. For policy tips, no override options are necessary because this rule is simply notifying people and not blocking access.

2. Second rule: If greater than five instances of this sensitive information are detected in a document, and the document is shared with people inside the organization, the **Block access to content** action restricts the permissions for the file, and the **Send a notification** action allows people to override the actions in this rule by providing a business justification. Your organization's business sometimes requires internal people to share PII data, and you don't want your DLP policy to block this work.

3. Third rule: If greater than five instances of this sensitive information are detected in a document, and the document is shared with people outside the organization, the **Block access to content** action restricts the permissions for the file, and the **Send a notification** action does not allow people to override the actions in this rule because the information is shared externally. Under no circumstances should people in your organization be allowed to share PII data outside the organization.

Here are some fine points to understand about using a policy tip to override a rule:

- The option to override is per rule, and it overrides all of the actions in the rule (except sending a notification, which can't be overridden).

- It's possible for content to match several rules in a DLP policy, but only the policy tip from the most restrictive, highest-priority rule will be shown. For example, a policy tip from a rule that blocks access to content will be shown over a policy tip from a rule that simply sends a notification. This prevents people from seeing a cascade of policy tips.

- If the policy tips in the most restrictive rule allow people to override the rule, then overriding this rule also overrides any other rules that the content matched.

Policy tips on OneDrive for Business sites and SharePoint Online sites

When a document on a OneDrive for Business site or SharePoint Online site matches a rule in a DLP policy, and that rule uses policy tips, the policy tips display special icons on the document:

1. If the rule sends a notification about the file, the warning icon appears.
2. If the rule blocks access to the document, the blocked icon appears.

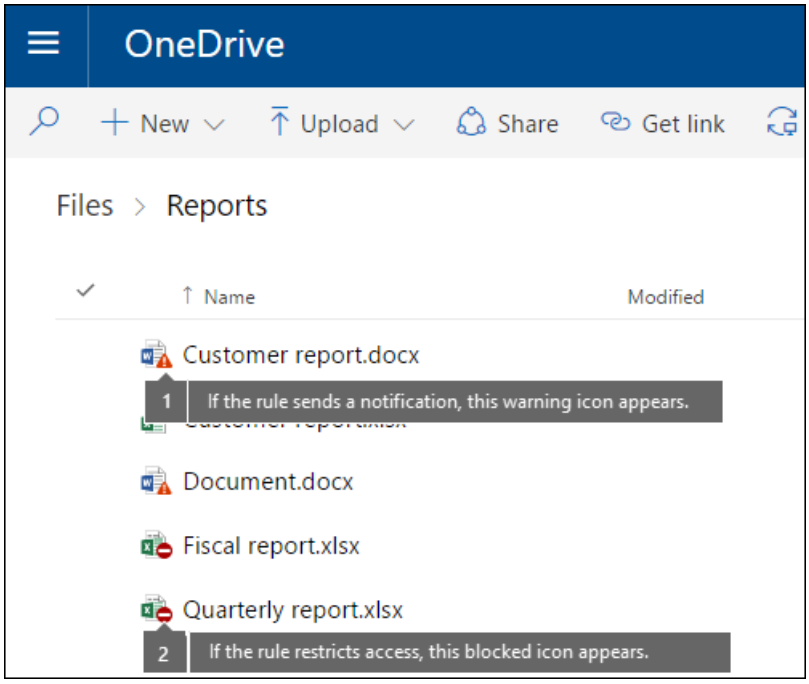

To take action on a document, you can select an item > choose **Information** in the upper-right corner of the page to open the details pane > **View policy tip**.

The policy tip lists the issues with the content, and if the policy tips are configured with these options, you can choose **Resolve**, and then **Override** the policy tip or **Report** a false positive.

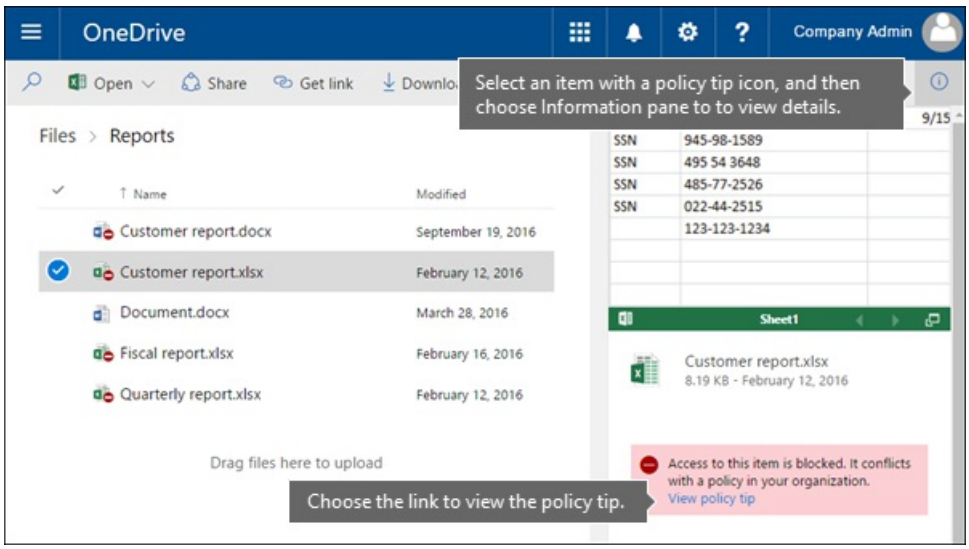

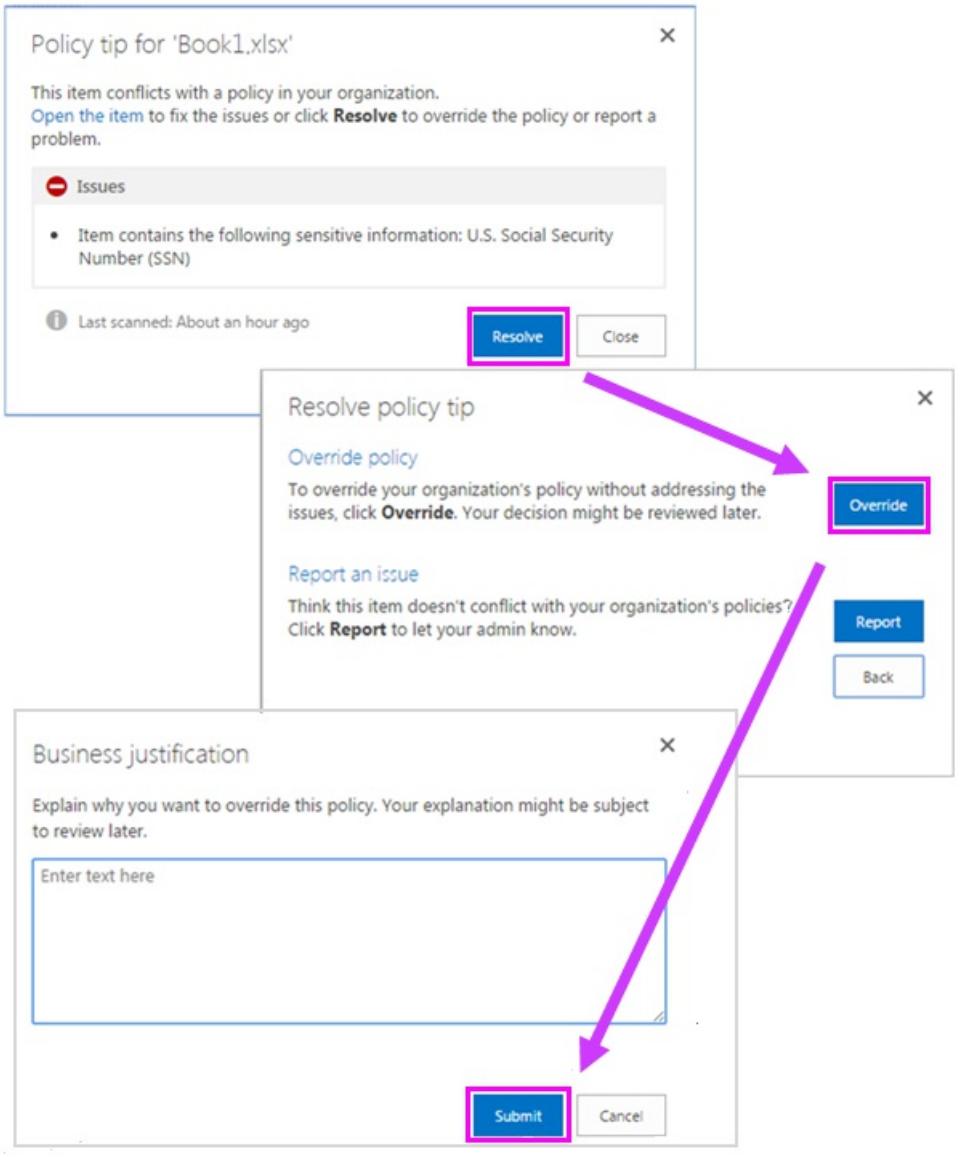

DLP policies are synced to sites and contented is evaluated against them periodically and asynchronously, so there may be a short delay between the time you create the DLP policy and the time you begin to see policy tips. There may be a similar delay from when you resolve or override a policy tip to when the icon on the document on the site goes away.

Default text for policy tips on sites

By default, policy tips display text similar to the following for an item on a site. The notification text is configured separately for each rule, so the text that's displayed differs depending on which rule is matched.

IF THE DLP POLICY RULE DOES THIS...	THEN THE DEFAULT POLICY TIP SAYS THIS...
Sends a notification but doesn't allow override	This item conflicts with a policy in your organization.
Blocks access, sends a notification, and allows override	This item conflicts with a policy in your organization. If you don't resolve this conflict, access to this file might be blocked.
Blocks access and sends a notification	This item conflicts with a policy in your organization. Access to this item is blocked for everyone except its owner, last modifier, and the primary site collection administrator.

Custom text for policy tips on sites

You can customize the text for policy tips separately from the email notification. Unlike custom text for email

notifications (see above section), custom text for policy tips does not accept HTML or tokens. Instead, custom text for policy tips is plain text only with a 256-character limit.

Policy tips in Outlook on the web and Outlook 2013 and later

When you compose a new email in Outlook on the web and Outlook 2013 and later, you'll see a policy tip if you add content that matches a rule in a DLP policy, and that rule uses policy tips. The policy tip appears at the top of the message, above the recipients, while the message is being composed.

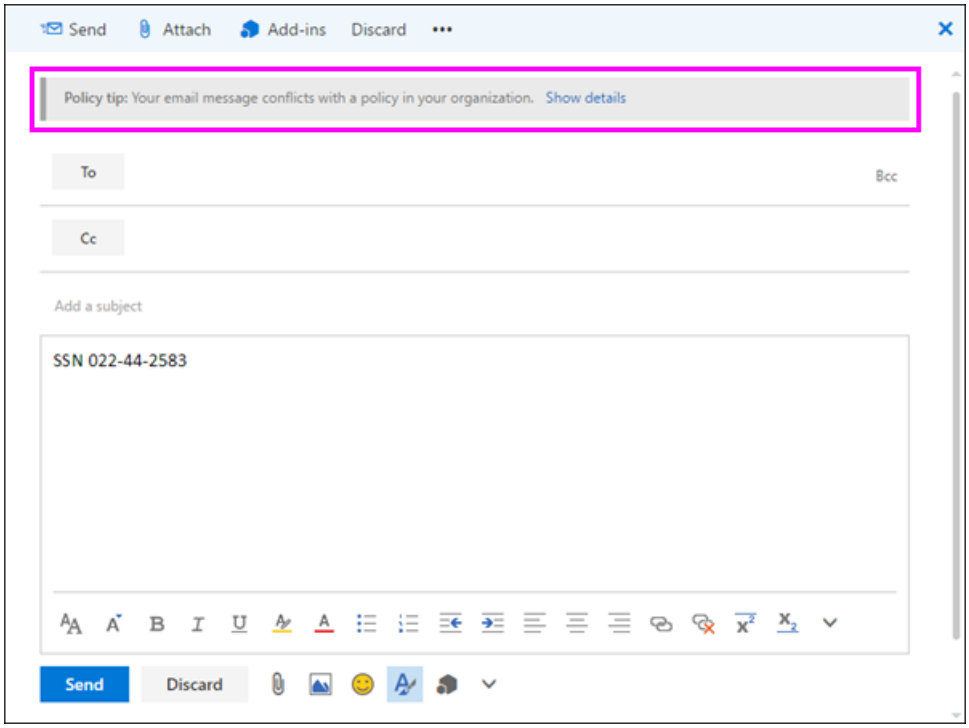

Policy tips work whether the sensitive information appears in the message body, subject line, or even a message attachment as shown here.

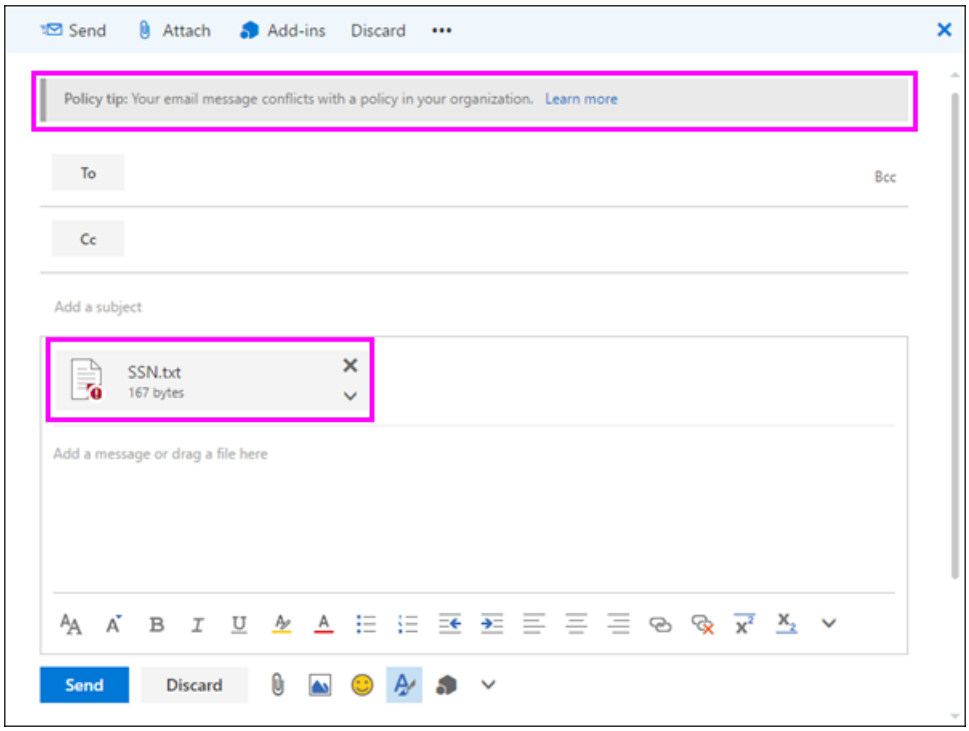

If the policy tips are configured to allow override, you can choose **Show Details** > **Override** > enter a business justification or report a false positive > **Override**.

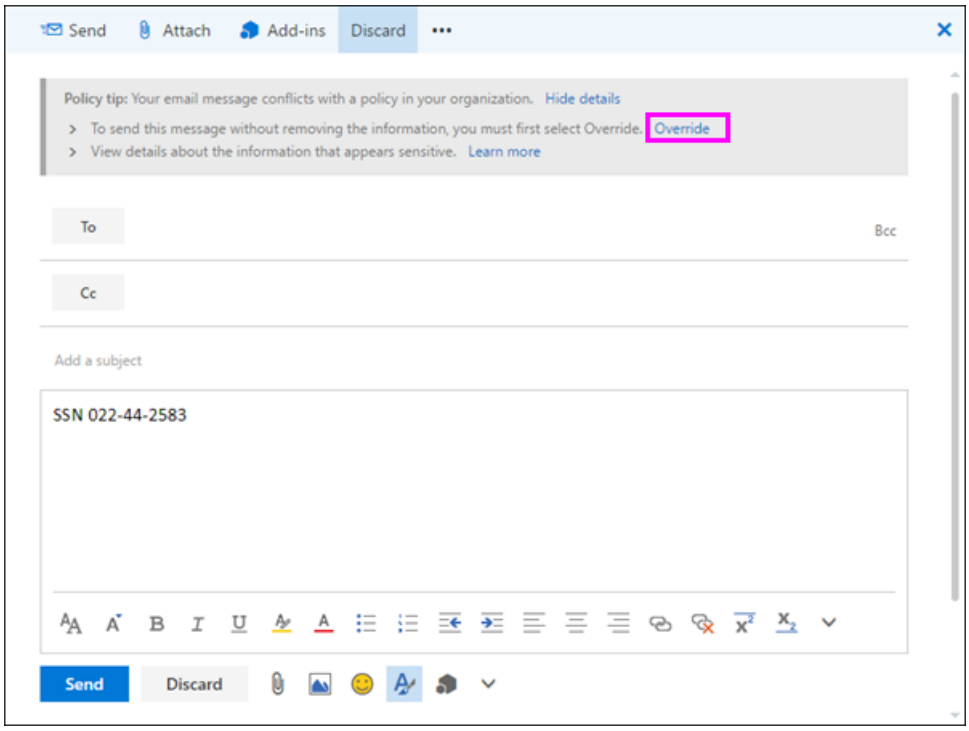

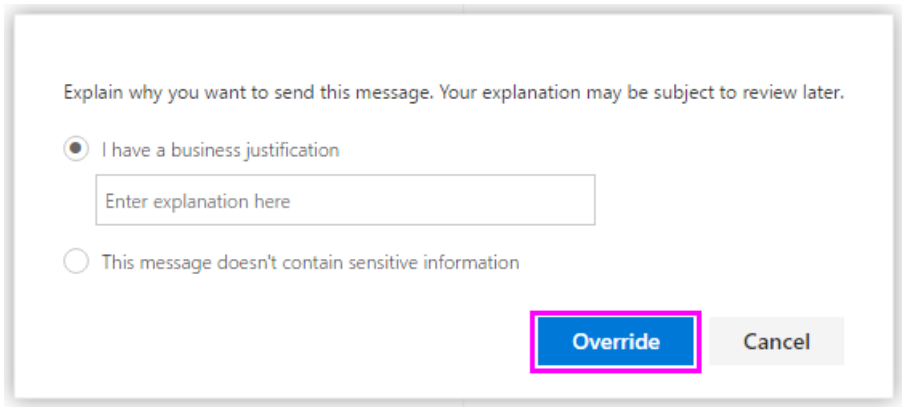

Note that when you add sensitive information to an email, there may be latency between when the sensitive information is added and when the policy tip appears.

Outlook 2013 and later supports showing policy tips for only some conditions

Currently, Outlook 2013 and later supports showing policy tips only for these conditions:

- Content contains
- Content is shared

Note that all of these conditions work in Outlook, where they will match content and enforce protective actions on content. But showing policy tips to users is not yet supported.

Policy tips in the Exchange admin center vs. the Security & Compliance Center

Policy tips can work either with DLP policies and mail flow rules created in the Exchange admin center, or with DLP policies created in the Security & Compliance Center, but not both. This is because these policies are stored in different locations, but policy tips can draw only from a single location.

If you've configured policy tips in the Exchange admin center, any policy tips that you configure in the Security & Compliance Center won't appear to users in Outlook on the web and Outlook 2013 and later until you turn off the tips in the Exchange admin center. This ensures that your current Exchange mail flow rules (also known as transport rules) will continue to work until you choose to switch over to the Security & Compliance Center.

Note that while policy tips can draw only from a single location, email notifications are always sent, even if you're using DLP policies in both the Security & Compliance Center and the Exchange admin center.

Default text for policy tips in email

By default, policy tips display text similar to the following for email.

IF THE DLP POLICY RULE DOES THIS...	THEN THE DEFAULT POLICY TIP SAYS THIS...
Sends a notification but doesn't allow override	Your email conflicts with a policy in your organization.
Blocks access, sends a notification, and allows override	Your email conflicts with a policy in your organization.
Blocks access and sends a notification	Your email conflicts with a policy in your organization.

Policy tips in Excel, PowerPoint, and Word

When people work with sensitive content in the desktop versions of Excel, PowerPoint, and Word, policy tips can notify them in real time that the content conflicts with a DLP policy. This requires that:

- The Office document is stored on a OneDrive for Business site or SharePoint Online site.
- The site is included in a DLP policy that's configured to use policy tips.

Office desktop programs automatically sync DLP policies directly from Office 365, and then scan your documents to ensure that they don't conflict with your DLP policies and display policy tips in real time.

> **NOTE**
> Office desktop apps scan documents themselves to determine if DLP policy tips should be shown; they do not show policy tips that SharePoint Online sites or OneDrive for Business sites have already determined should be shown on a file. As a result, you may not always see a DLP policy tip in the desktop apps that you see in the SharePoint Online sites or OneDrive for Business sites. In contrast, the Office applications on the web only show DLP policy tips that SharePoint Online sites or OneDrive for Business sites have already determined should be shown.

Depending on how you configure the policy tips in the DLP policy, people can choose to simply ignore the policy tip, override the policy with or without a business justification, or report a false positive.

Policy tips appear on the Message Bar.

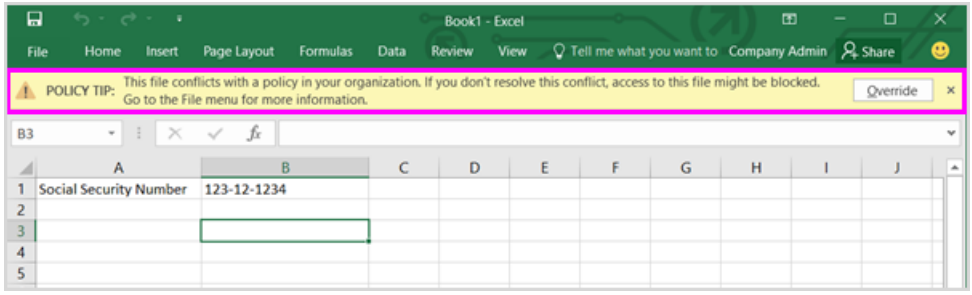

And policy tips also appear in the Backstage view (on the **File** tab).

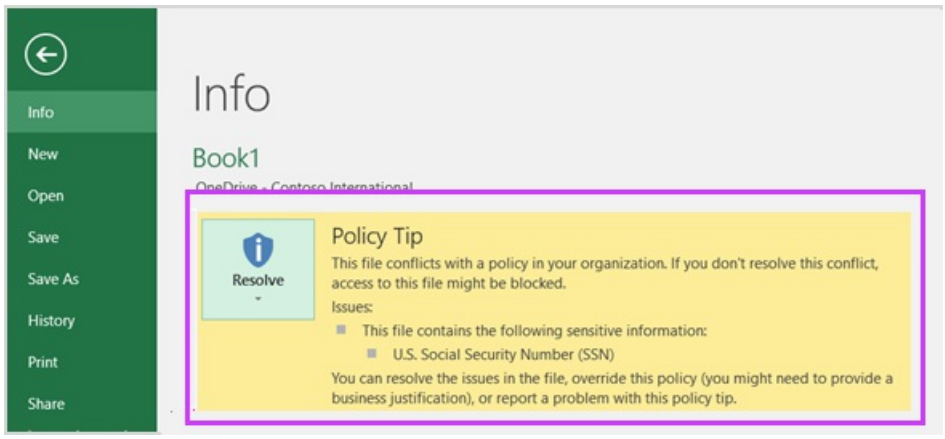

If policy tips in the DLP policy are configured with these options, you can choose **Resolve** to **Override** a policy tip or **Report** a false positive.

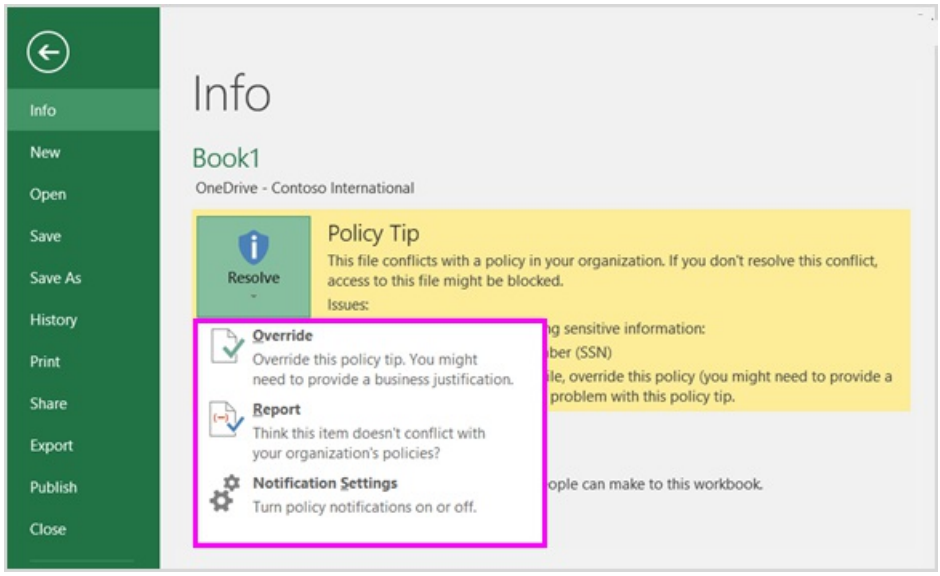

In each of these Office desktop programs, people can choose to turn off policy tips. If turned off, policy tips that are simple notifications will not appear on the Message Bar or Backstage view (on the **File** tab). However, policy tips about blocking and overriding will still appear, and they will still receive the email notification. In addition, turning off policy tips does not exempt the document from any DLP policies that have been applied to it.

Default text for policy tips in Excel 2016, PowerPoint 2016, and Word 2016

By default, policy tips display text similar to the following on the Message Bar and Backstage view of an open document. The notification text is configured separately for each rule, so the text that's displayed differs depending on which rule is matched.

IF THE DLP POLICY RULE DOES THIS...	THEN THE DEFAULT POLICY TIP SAYS THIS...
Sends a notification but doesn't allow override	This file conflicts with a policy in your organization. Go to the **File** menu for more information.
Blocks access, sends a notification, and allows override	This file conflicts with a policy in your organization. If you don't resolve this conflict, access to this file might be blocked. Go to the **File** menu for more information.
Blocks access and sends a notification	This file conflicts with a policy in your organization. If you don't resolve this conflict, access to this file might be blocked. Go to the **File** menu for more information.

Custom text for policy tips in Excel, PowerPoint, and Word

You can customize the text for policy tips separately from the email notification. Unlike custom text for email notifications (see above section), custom text for policy tips does not accept HTML or tokens. Instead, custom text for policy tips is plain text only with a 256-character limit.

More information

- Overview of data loss prevention policies
- Create a DLP policy from a template
- Create a DLP policy to protect documents with FCI or other properties
- What the DLP policy templates include
- Sensitive information type entity definitions

What the DLP policy templates include

11/2/2020 • 28 minutes to read • Edit Online

Data loss prevention (DLP) in the Security & Compliance Center includes ready-to-use policy templates that address common compliance requirements, such as helping you to protect sensitive information subject to the U.S. Health Insurance Act (HIPAA), U.S. Gramm-Leach-Bliley Act (GLBA), or U.S. Patriot Act. This topic lists all of the policy templates, what types of sensitive information they look for, and what the default conditions and actions are. This topic does not include every detail of how each policy template is configured; instead, the topic presents with you enough information to help you decide which template is the best starting point for your scenario. Remember, you can customize these policy templates to meet your specific requirements.

Australia Financial Data

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
Australia Financial: Scan content shared outside - low count	Content contains sensitive information: SWIFT Code — Min count 1, Max count 9 Australia Tax File Number — Min count 1, Max count 9 Australia Bank Account Number — Min count 1, Max count 9 Credit Card Number — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
Australia Financial: Scan content shared outside - high count	Content contains sensitive information: SWIFT Code — Min count 10, Max count any Australia Tax File Number — Min count 10, Max count any Australia Bank Account Number — Min count 10, Max count any Credit Card Number — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

Australia Health Records Act (HRIP Act)

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
Australia HRIP: Scan content shared outside - low count	Content contains sensitive information: Australia Tax File Number — Min count 1, Max count 9 Australia Medical Account Number — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
Australia HRIP: Scan content shared outside - high count	Content contains sensitive information: Australia Tax File Number — Min count 10, Max count any Australia Medical Account Number — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

Australia Personally Identifiable Information (PII) Data

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
Australia PII: Scan content shared outside - low count	Content contains sensitive information: Australia Tax File Number — Min count 1, Max count 9 Australia Driver's License Number — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
Australia PII: Scan content shared outside - high count	Content contains sensitive information: Australia Tax File Number — Min count 10, Max count any Australia Driver's License Number — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

Australia Privacy Act

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
Australia Privacy: Scan content shared outside - low count	Content contains sensitive information: Australia Driver's License Number — Min count 1, Max count 9 Australia Passport Number — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
Australia Privacy: Scan content shared outside - high count	Content contains sensitive information: Australia Driver's License Number — Min count 10, Max count any Australia Passport Number — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

Canada Financial Data

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
Canada Financial Data: Scan content shared outside - low count	Content contains sensitive information: Credit Card Number — Min count 1, Max count 9 Canada Bank Account Number — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
Canada Financial Data: Scan content shared outside - high count	Content contains sensitive information: Credit Card Number — Min count 10, Max count any Canada Bank Account Number — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

Canada Health Information Act (HIA)

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
Canada HIA: Scan content shared outside - low count	Content contains sensitive information: Canada Passport Number — Min count 1, Max count 9 Canada Social Insurance Number — Min count 1, Max count 9 Canada Health Service Number — Min count 1, Max count 9 Canada Personal Health Identification Number (PHIN) — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
Canada HIA: Scan content shared outside - high count	Content contains sensitive information: Canada Passport Number — Min count 10, Max count any Canada Social Insurance Number — Min count 10, Max count any Canada Health Service Number — Min count 10, Max count any Canada Personal Health Identification Number (PHIN) — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

Canada Personal Health Act (PHIPA) - Ontario

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
Canada PHIPA: Scan content shared outside - low count	Content contains sensitive information: Canada Passport Number — Min count 1, Max count 9 Canada Social Insurance Number — Min count 1, Max count 9 Canada Health Service Number — Min count 1, Max count 9 Canada Personal Health Identification Number (PHIN) — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
Canada PHIPA: Scan content shared outside - high count	Content contains sensitive information: Canada Passport Number — Min count 10, Max count any Canada Social Insurance Number — Min count 10, Max count any Canada Health Service Number — Min count 10, Max count any Canada Personal Health Identification Number (PHIN) — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

Canada Personal Health Information Act (PHIA) - Manitoba

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
Canada PHIA: Scan content shared outside - low count	Content contains sensitive information: Canada Social Insurance Number — Min count 1, Max count 9 Canada Health Service Number — Min count 1, Max count 9 Canada Personal Health Identification Number (PHIN) — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
Canada PHIA: Scan content shared outside - high count	Content contains sensitive information: Canada Social Insurance Number — Min count 10, Max count any Canada Health Service Number — Min count 10, Max count any Canada Personal Health Identification Number (PHIN) — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

Canada Personal Information Protection Act (PIPA)

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
Canada PIPA: Scan content shared outside - low count	Content contains sensitive information: Canada Passport Number — Min count 1, Max count 9 Canada Social Insurance Number — Min count 1, Max count 9 Canada Health Service Number — Min count 1, Max count 9 Canada Personal Health Identification Number (PHIN) — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
Canada PIPA: Scan content shared outside - high count	Content contains sensitive information: Canada Passport Number — Min count 10, Max count any Canada Social Insurance Number — Min count 10, Max count any Canada Health Service Number — Min count 10, Max count any Canada Personal Health Identification Number (PHIN) — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

Canada Personal Information Protection Act (PIPEDA)

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
Canada PIPEDA: Scan content shared outside - low count	Content contains sensitive information: Canada Driver's License Number — Min count 1, Max count 9 Canada Bank Account Number — Min count 1, Max count 9 Canada Passport Number — Min count 1, Max count 9 Canada Social Insurance Number — Min count 1, Max count 9 Canada Health Service Number — Min count 1, Max count 9 Canada Personal Health Identification Number (PHIN) — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
Canada PIPEDA: Scan content shared outside - high count	Content contains sensitive information: Canada Driver's License Number — Min count 10, Max count any Canada Bank Account Number — Min count 10, Max count any Canada Passport Number — Min count 10, Max count any Canada Social Insurance Number — Min count 10, Max count any Canada Health Service Number — Min count 10, Max count any Canada Personal Health Identification Number (PHIN) — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

Canada Personally Identifiable Information (PII) Data

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
Canada PII: Scan content shared outside - low count	Content contains sensitive information: Canada Driver's License Number — Min count 1, Max count 9 Canada Bank Account Number — Min count 1, Max count 9 Canada Passport Number — Min count 1, Max count 9 Canada Social Insurance Number — Min count 1, Max count 9 Canada Health Service Number — Min count 1, Max count 9 Canada Personal Health Identification Number (PHIN) — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
Canada PII: Scan content shared outside - high count	Content contains sensitive information: Canada Driver's License Number — Min count 10, Max count any Canada Bank Account Number — Min count 10, Max count any Canada Passport Number — Min count 10, Max count any Canada Social Insurance Number — Min count 10, Max count any Canada Health Service Number — Min count 10, Max count any Canada Personal Health Identification Number (PHIN) — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

France Data Protection Act

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
France DPA: Scan content shared outside - low count	Content contains sensitive information: France National ID Card (CNI) — Min count 1, Max count 9 France Social Security Number (INSEE) — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
France DPA: Scan content shared outside - high count	Content contains sensitive information: France National ID Card (CNI) — Min count 10, Max count any France Social Security Number (INSEE) — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

France Financial Data

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
France Financial: Scan content shared outside - low count	Content contains sensitive information: Credit Card Number — Min count 1, Max count 9 EU Debit Card Number — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
France Financial: Scan content shared outside - high count	Content contains sensitive information: Credit Card Number — Min count 10, Max count any EU Debit Card Number — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

France Personally Identifiable Information (PII) Data

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
France PII: Scan content shared outside - low count	Content contains sensitive information: France Social Security Number (INSEE) — Min count 1, Max count 9 France Driver's License Number — Min count 1, Max count 9 France Passport Number — Min count 1, Max count 9 France National ID Card (CNI) — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
France PII: Scan content shared outside - high count	Content contains sensitive information: France Social Security Number (INSEE) — Min count 10, Max count any France Driver's License Number — Min count 10, Max count any France Passport Number — Min count 10, Max count any France National ID Card (CNI) — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

General Data Protection Regulation (GDPR)

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
Low volume EU Sensitive content found	Content contains sensitive information: EU Debit Card Number — Min count 1, Max count 9 EU Driver's License Number — Min count 1, Max count 9 EU National Identification Number — Min count 1, Max count 9 EU Passport Number — Min count 1, Max count 9 EU Social Security Number (SSN) or Equivalent ID — Min count 1, Max count 9 EU Tax Identification Number (TIN) — Min count 1, Max count 9 Content is shared with: People outside my organization	Send incident reports to Administrator

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
High volume of EU Sensitive content found	Content contains sensitive information: EU Debit Card Number — Min count 1, Max count 9 EU Driver's License Number — Min count 1, Max count 9 EU National Identification Number — Min count 1, Max count 9 EU Passport Number — Min count 1, Max count 9 EU Social Security Number (SSN) or Equivalent ID — Min count 1, Max count 9 EU Tax Identification Number (TIN) — Min count 1, Max count 9 Content is shared with: People outside my organization	Restrict access to the content for external users Notify users with email and policy tips Allow override Require business justification Send incident reports to Administrator

Germany Financial Data

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
Germany Financial Data: Scan content shared outside - low count	Content contains sensitive information: Credit Card Number — Min count 1, Max count 9 EU Debit Card Number — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
Germany Financial Data: Scan content shared outside - high count	Content contains sensitive information: Credit Card Number — Min count 10, Max count any EU Debit Card Number — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

Germany Personally Identifiable Information (PII) Data

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
Germany PII: Scan content shared outside - low count	Content contains sensitive information: German Driver's License Number — Min count 1, Max count 9 German Passport Number — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
Germany PII: Scan content shared outside - high count	Content contains sensitive information: German Driver's License Number — Min count 10, Max count any German Passport Number — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

Israel Financial Data

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
Israel Financial Data: Scan content shared outside - low count	Content contains sensitive information: Israel Bank Account Number — Min count 1, Max count 9 SWIFT Code — Min count 1, Max count 9 Credit Card Number — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
Israel Financial Data: Scan content shared outside - high count	Content contains sensitive information: Israel Bank Account Number — Min count 10, Max count any SWIFT Code — Min count 10, Max count any Credit Card Number — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

Israel Personally Identifiable Information (PII) Data

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
Israel PII: Scan content shared outside - low count	Content contains sensitive information: Israel National ID — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
Israel PII: Scan content shared outside - high count	Content contains sensitive information: Israel National ID — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

Israel Protection of Privacy

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
Israel Privacy: Scan content shared outside - low count	Content contains sensitive information: Israel National ID — Min count 1, Max count 9 Israel Bank Account Number — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
Israel Privacy: Scan content shared outside - high count	Content contains sensitive information: Israel National ID — Min count 10, Max count any Israel Bank Account Number — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

Japan Financial Data

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
Japan Financial: Scan content shared outside - low count	Content contains sensitive information: Japan Bank Account Number — Min count 1, Max count 9 Credit Card Number — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
Japan Financial: Scan content shared outside - high count	Content contains sensitive information: Japan Bank Account Number — Min count 10, Max count any Credit Card Number — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

Japan Personally Identifiable Information (PII) Data

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
Japan PII: Scan content shared outside - low count	Content contains sensitive information: Japan Resident Registration Number — Min count 1, Max count 9 Japan Social Insurance Number (SIN) — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
Japan PII: Scan content shared outside - high count	Content contains sensitive information: Japan Resident Registration Number — Min count 10, Max count any Japan Social Insurance Number (SIN) — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

Japan Protection of Personal Information

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
Japan PPI: Scan content shared outside - low count	Content contains sensitive information: Japan Resident Registration Number — Min count 1, Max count 9 Japan Social Insurance Number (SIN) — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
Japan PPI: Scan content shared outside - high count	Content contains sensitive information: Japan Resident Registration Number — Min count 10, Max count any Japan Social Insurance Number (SIN) — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

PCI Data Security Standard (PCI DSS)

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
PCI DSS: Scan content shared outside - low count	Content contains sensitive information: Credit Card Number — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
PCI DSS: Scan content shared outside - high count	Content contains sensitive information: Credit Card Number — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

Saudi Arabia - Anti-Cyber Crime Law

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
Saudi Arabia ACC: Scan content shared outside - low count	Content contains sensitive information: SWIFT Code — Min count 1, Max count 9 International Banking Account Number (IBAN) — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
Saudi Arabia ACC: Scan content shared outside - high count	Content contains sensitive information: SWIFT Code — Min count 10, Max count any International Banking Account Number (IBAN) — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

Saudi Arabia Financial Data

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
Saudi Arabia Financial: Scan content shared outside - low count	Content contains sensitive information: Credit Card Number — Min count 1, Max count 9 SWIFT Code — Min count 1, Max count 9 International Banking Account Number (IBAN) — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
Saudi Arabia Financial: Scan content shared outside - high count	Content contains sensitive information: Credit Card Number — Min count 10, Max count any SWIFT Code — Min count 10, Max count any International Banking Account Number (IBAN) — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

Saudi Arabia Personally Identifiable Information (PII) Data

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
Saudi Arabia PII: Scan content shared outside - low count	Content contains sensitive information: Saudi Arabia National ID — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
Saudi Arabia PII: Scan content shared outside - high count	Content contains sensitive information: Saudi Arabia National ID — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

U.K. Access to Medical Reports Act

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
U.K. AMRA: Scan content shared outside - low count	Content contains sensitive information: U.K. National Health Service Number — Min count 1, Max count 9 U.K. National Insurance Number (NINO) — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
U.K. AMRA: Scan content shared outside - high count	Content contains sensitive information: U.K. National Health Service Number — Min count 10, Max count any U.K. National Insurance Number (NINO) — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

U.K. Data Protection Act

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
U.K. DPA: Scan content shared outside - low count	Content contains sensitive information: U.K. National Insurance Number (NINO) — Min count 1, Max count 9 U.S. / U.K. Passport Number — Min count 1, Max count 9 SWIFT Code — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
U.K. DPA: Scan content shared outside - high count	Content contains sensitive information: U.K. National Insurance Number (NINO) — Min count 10, Max count any U.S. / U.K. Passport Number — Min count 10, Max count any SWIFT Code — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

U.K. Financial Data

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
U.K. Financial: Scan content shared outside - low count	Content contains sensitive information: Credit Card Number — Min count 1, Max count 9 EU Debit Card Number — Min count 1, Max count 9 SWIFT Code —Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
U.K. Financial: Scan content shared outside - high count	Content contains sensitive information: Credit Card Number — Min count 10, Max count any EU Debit Card Number — Min count 10, Max count any SWIFT Code — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

U.K. Personal Information Online Code of Practice (PIOCP)

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
U.K. PIOCP: Scan content shared outside - low count	Content contains sensitive information: U.K. National Insurance Number (NINO) — Min count 1, Max count 9 U.K. National Health Service Number — Min count 1, Max count 9 SWIFT Code — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
U.K. PIOCP: Scan content shared outside - high count	Content contains sensitive information: U.K. National Insurance Number (NINO) — Min count 10, Max count any U.K. National Health Service Number — Min count 10, Max count any SWIFT Code — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

U.K. Personally Identifiable Information (PII) Data

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
U.K. PII: Scan content shared outside - low count	Content contains sensitive information: U.K. National Insurance Number (NINO) — Min count 1, Max count 9 U.S. / U.K. Passport Number — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
U.K. PII: Scan content shared outside - high count	Content contains sensitive information: U.K. National Insurance Number (NINO) — Min count 10, Max count any U.S. / U.K. Passport Number — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

U.K. Privacy and Electronic Communications Regulations

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
U.K. PECR: Scan content shared outside - low count	Content contains sensitive information: SWIFT Code — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
U.K. PECR: Scan content shared outside - high count	Content contains sensitive information: SWIFT Code — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

U.S. Federal Trade Commission (FTC) Consumer Rules

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
U.S. FTC Rules: Scan content shared outside - low count	Content contains sensitive information: Credit Card Number — Min count 1, Max count 9 U.S. Bank Account Number — Min count 1, Max count 9 ABA Routing Number — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
U.S. FTC Rules: Scan content shared outside - high count	Content contains sensitive information: Credit Card Number — Min count 10, Max count any U.S. Bank Account Number — Min count 10, Max count any ABA Routing Number — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

U.S. Financial Data

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
U.S. Financial: Scan content shared outside - low count	Content contains sensitive information: Credit Card Number — Min count 1, Max count 9 U.S. Bank Account Number — Min count 1, Max count 9 ABA Routing Number — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
U.S. Financial: Scan content shared outside - high count	Content contains sensitive information: Credit Card Number — Min count 10, Max count any U.S. Bank Account Number — Min count 10, Max count any ABA Routing Number — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

U.S. Gramm-Leach-Bliley Act (GLBA)

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
U.S. GLBA: Scan content shared outside - low count	Content contains sensitive information: Credit Card Number — Min count 1, Max count 9 U.S. Bank Account Number — Min count 1, Max count 9 U.S. Individual Taxpayer Identification Number (ITIN) — Min count 1, Max count 9 U.S. Social Security Number (SSN) — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
U.S. GLBA: Scan content shared outside - high count	Content contains sensitive information: Credit Card Number — Min count 10, Max count any U.S. Bank Account Number — Min count 10, Max count any U.S. Individual Taxpayer Identification Number (ITIN) — Min count 10, Max count any U.S. Social Security Number (SSN) — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

U.S. Health Insurance Act (HIPAA)

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
Content matches U.S. HIPAA	Contains any of the following sensitive information: U.S. Social Security Number (SSN) — Min count 1, Max count any Drug Enforcement Agency (DEA) Number — Min count 1, Max count any **AND** Content contains any of these terms: International Classification of Diseases (ICD-9-CM) — Min count 1, Max count any International Classification of Diseases (ICD-10-CM) — Min count 1, Max count any Content is shared with: People outside my organization	Send a notification

U.S. Patriot Act

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
U.S. Patriot Act: Scan content shared outside - low count	Content contains sensitive information: Credit Card Number — Min count 1, Max count 9 U.S. Bank Account Number — Min count 1, Max count 9 U.S. Individual Taxpayer Identification Number (ITIN) — Min count 1, Max count 9 U.S. Social Security Number (SSN) — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
U.S. Patriot Act: Scan content shared outside - high count	Content contains sensitive information: Credit Card Number — Min count 10, Max count any U.S. Bank Account Number — Min count 10, Max count any U.S. Individual Taxpayer Identification Number (ITIN) — Min count 10, Max count any U.S. Social Security Number (SSN) — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

U.S. Personally Identifiable Information (PII) Data

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
U.S. PII: Scan content shared outside - low count	Content contains sensitive information: U.S. Individual Taxpayer Identification Number (ITIN) — Min count 1, Max count 9 U.S. Social Security Number (SSN) — Min count 1, Max count 9 U.S. / U.K. Passport Number — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
U.S. PII: Scan content shared outside - high count	Content contains sensitive information: U.S. Individual Taxpayer Identification Number (ITIN) — Min count 10, Max count any U.S. Social Security Number (SSN) — Min count 10, Max count any U.S. / U.K. Passport Number — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

U.S. State Breach Notification Laws

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
U.S. State Breach: Scan content shared outside - low count	Content contains sensitive information: Credit Card Number — Min count 1, Max count 9 U.S. Bank Account Number — Min count 1, Max count 9 U.S. Driver's License Number — Min count 1, Max count 9 U.S. Social Security Number (SSN) — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
U.S. State Breach: Scan content shared outside - high count	Content contains sensitive information: Credit Card Number — Min count 10, Max count any U.S. Bank Account Number — Min count 10, Max count any U.S. Driver's License Number — Min count 10, Max count any U.S. Social Security Number (SSN) — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

U.S. State Social Security Number Confidentiality Laws

RULE NAME	CONDITIONS (INCLUDING SENSITIVE INFORMATION TYPES)	ACTIONS
U.S. SSN Laws: Scan content shared outside - low count	Content contains sensitive information: U.S. Social Security Number (SSN) — Min count 1, Max count 9 Content is shared with: People outside my organization	Send a notification
U.S. SSN Laws: Scan content shared outside - high count	Content contains sensitive information: U.S. Social Security Number (SSN) — Min count 10, Max count any Content is shared with: People outside my organization	Block access to content Send a notification Allow override Require business justification Send incident report

Create a DLP policy to protect documents with FCI or other properties

2/18/2021 • 7 minutes to read • Edit Online

Microsoft 365 data loss prevention (DLP) policies can use classification properties or item properties to identify sensitive items. For example you can use:

- Windows Server File Classification infrastructure (FCI) properties
- SharePoint document properties
- third-party system document properties

For example, your organization might use Windows Server FCI to identify items with personal data such as social security numbers, and then classify the document by setting the **Personally Identifiable Information** property to **High**, **Moderate**, **Low**, **Public**, or **Not PII** based on the type and number of occurrences of personal data found in the document.

In Microsoft 365, you can create a DLP policy that identifies documents that have that property set to specific values, such as **High** and **Medium**, and then takes an action such as blocking access to those files. The same policy can have another rule that takes a different action if the property is set to **Low**, such as sending an email notification. In this way, DLP integrates with Windows Server FCI and can help protect Office documents uploaded or shared to Microsoft 365 from Windows Server-based file servers.

A DLP policy simply looks for a specific property name/value pair. Any document property can be used, as long as the property has a corresponding managed property for SharePoint search. For example, a SharePoint site collection might use a content type named **Trip Report** with a required field named **Customer**. Whenever a person creates a trip report, they must enter the customer name. This property name/value pair can also be used in a DLP policy—for example, if you want a rule that blocks access to the document for guests when the **Customer** field contains **Contoso**.

If you want to apply your DLP policy to content with specific Microsoft 365 labels, you should not follow the steps here. Instead, learn how to Using a retention label as a condition in a DLP policy.

Before you create the DLP policy

Before you can use a Windows Server FCI property or other property in a DLP policy, you need to create a managed property in the SharePoint admin center. Here's why.

In SharePoint Online and OneDrive for Business, the search index is built up by crawling the content on your sites. The crawler picks up content and metadata from the documents in the form of crawled properties. The search schema helps the crawler decide what content and metadata to pick up. Examples of metadata are the author and the title of a document. However, to get the content and metadata from the documents into the

search index, the crawled properties must be mapped to managed properties. Only managed properties are kept in the index. For example, a crawled property related to author is mapped to a managed property related to author.

> **NOTE**
>
> Be sure to use a managed property name and not a crawled property name when creating DLP rules using the `ContentPropertyContainsWords` condition.

This is important because DLP uses the search crawler to identify and classify sensitive information on your sites, and then store that sensitive information in a secure portion of the search index. When you upload a document to Office 365, SharePoint automatically creates crawled properties based on the document properties. But to use an FCI or other property in a DLP policy, that crawled property needs to be mapped to a managed property so that content with that property is kept in the index.

For more information on search and managed properties, see Manage the search schema in SharePoint Online.

Step 1: Upload a document with the needed property to Office 365

You first need to upload a document with the property that you want to reference in your DLP policy. Microsoft 365 will detect the property and automatically create a crawled property from it. In the next step, you'll create a managed property, and then map the managed property to this crawled property.

Step 2: Create a managed property

1. Sign in to the Microsoft 365 admin center.
2. In the left navigation, choose **Admin centers** > **SharePoint**. You're now in the SharePoint admin center.
3. In the left navigation, choose **search** > on the **search administration** page > **Manage Search Schema**.

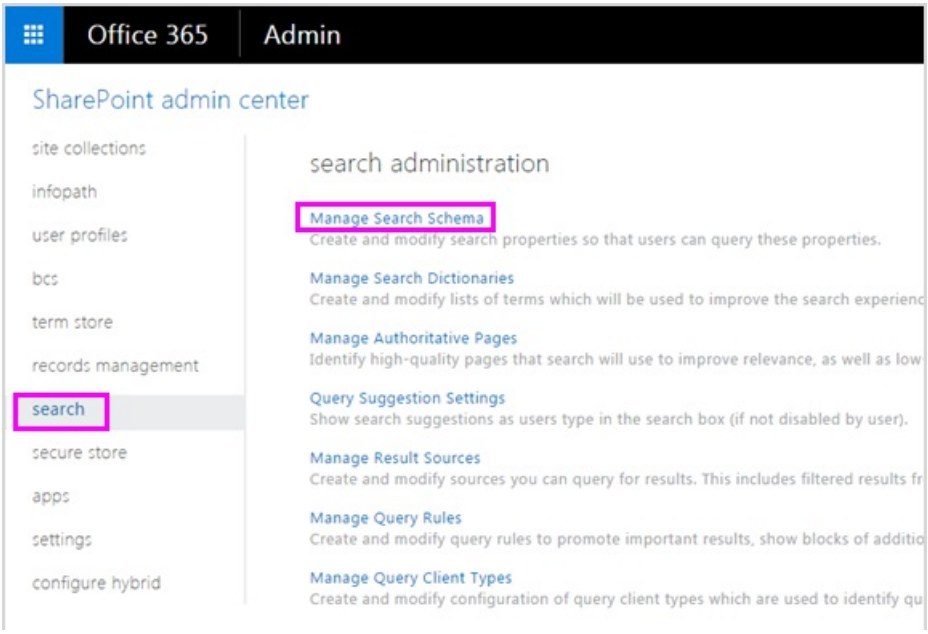

4. On the **Managed Properties** page > **New Managed Property**.

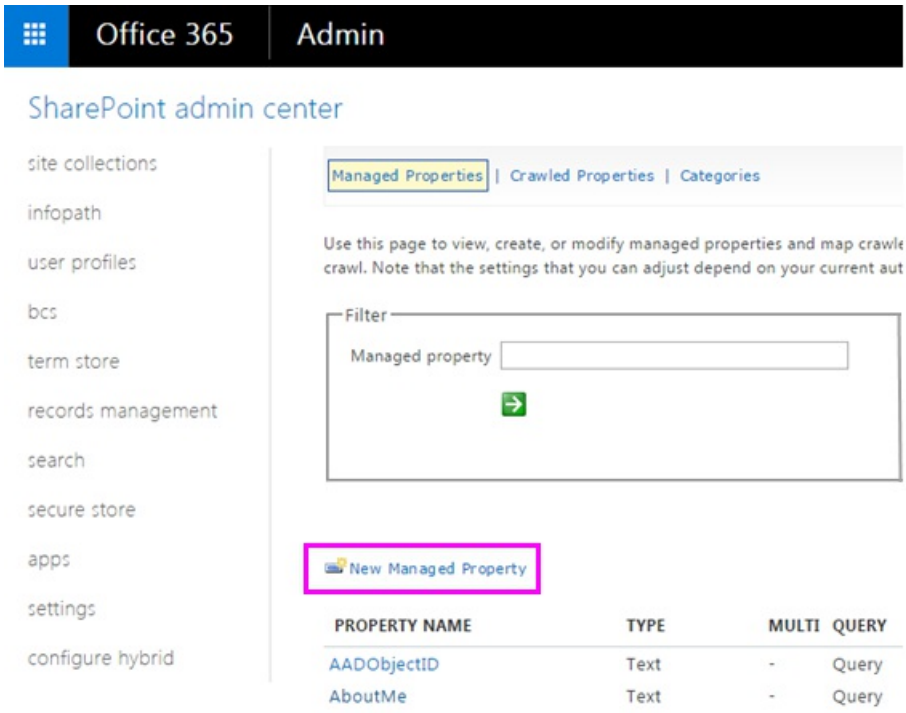

5. Enter a name and description for the property. This name is what will appear in your DLP policies.
6. For **Type**, choose **Text**.
7. Under **Main characteristics**, select **Queryable** and **Retrievable**.
8. Under **Mappings to crawled properties** > **Add a mapping**.
9. In the **crawled property selection** dialog box > find and select the crawled property that corresponds to the Windows Server FCI property or other property that you will use in your DLP policy > **OK**.

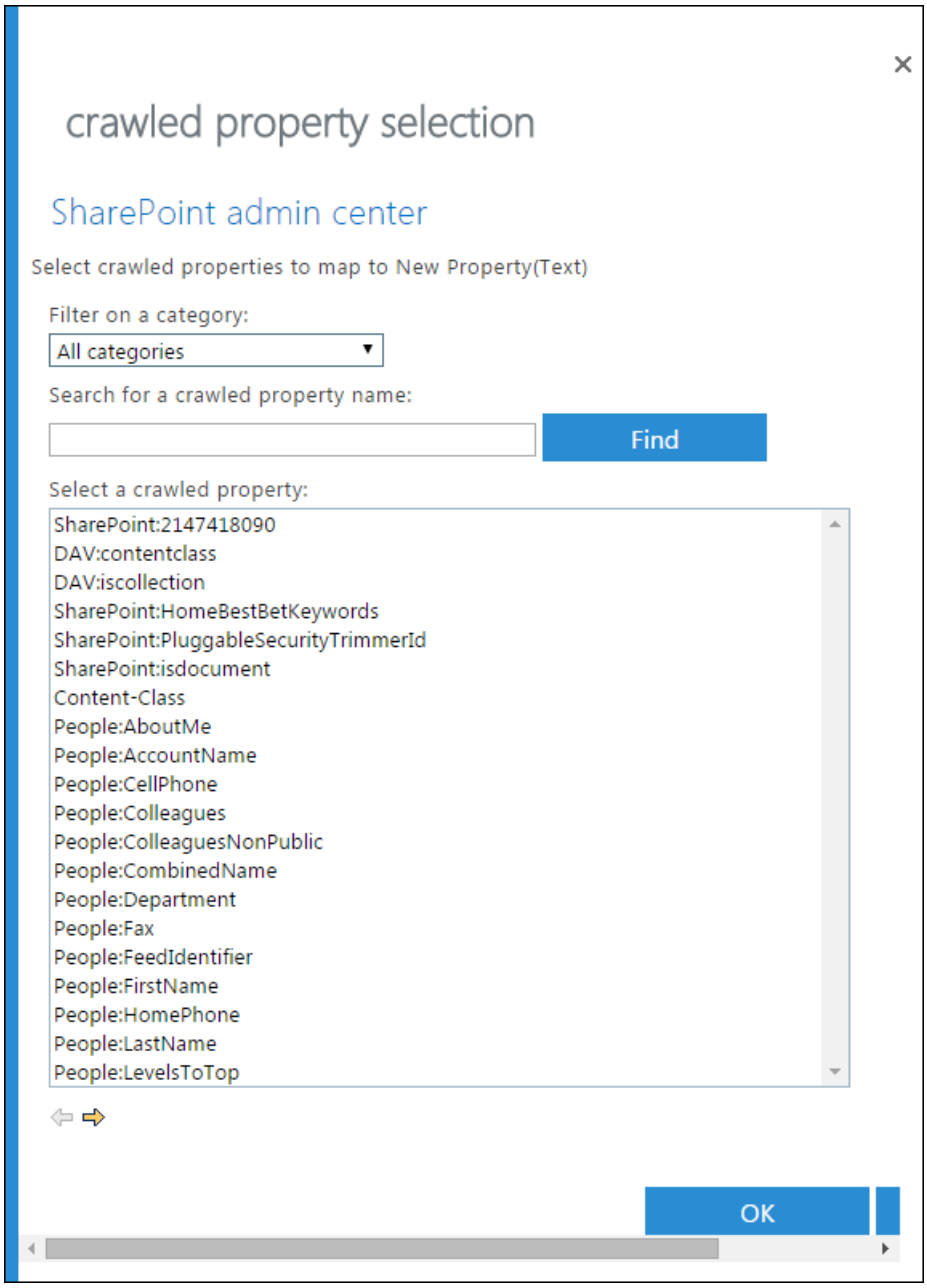

10. At the bottom of the page > **OK**.

Create a DLP policy that uses an FCI property or other property

In this example, an organization is using FCI on its Windows Server-based file servers; specifically, they're using the FCI classification property named **Personally Identifiable Information** with possible values of **High**, **Moderate**, **Low**, **Public**, and **Not PII**. Now they want to use their existing FCI classification in their DLP policies in Office 365.

First, they follow the steps above to create a managed property in SharePoint Online, which maps to the crawled property created automatically from the FCI property.

Next, they create a DLP policy with two rules that both use the condition **Document properties contain any of these values**:

- **FCI PII content - High, Moderate** The first rule restricts access to the document if the FCI classification property **Personally Identifiable Information** equals **High** or **Moderate** and the document is shared with people outside the organization.
- **FCI PII content - Low** The second rule sends a notification to the document owner if the FCI classification property **Personally Identifiable Information** equals **Low** and the document is shared

with people outside the organization.

Create the DLP policy by using PowerShell

The condition **Document properties contain any of these values** is temporarily not available in the UI of the Security & Compliance Center, but you can still use this condition by using PowerShell. You can use the `New\Set\Get-DlpCompliancePolicy` cmdlets to work with a DLP policy, and use the `New\Set\Get-DlpComplianceRule` cmdlets with the `ContentPropertyContainsWords` parameter to add the condition **Document properties contain any of these values**.

For more information on these cmdlets, see Security & Compliance Center cmdlets.

1. Connect to the Security & Compliance Center using remote PowerShell
2. Create the policy by using `New-DlpCompliancePolicy`.

This PowerShell creates a DLP policy that applies to all locations.

```
New-DlpCompliancePolicy -Name FCI_PII_policy -ExchangeLocation All -SharePointLocation All -OneDriveLocation
All -Mode Enable
```

3. Create the two rules described above by using `New-DlpComplianceRule`, where one rule is for the **Low** value, and another rule is for the **High** and **Moderate** values.

   Here is a PowerShell example that creates these two rules. The property name/value pairs are enclosed in quotation marks, and a property name may specify multiple values separated by commas with no spaces, like `"<Property1>:<Value1>,<Value2>","<Property2>:<Value3>,<Value4>"....`

   ```
   New-DlpComplianceRule -Name FCI_PII_content-High,Moderate -Policy FCI_PII_policy -AccessScope
   NotInOrganization -BlockAccess $true -ContentPropertyContainsWords "Personally Identifiable
   Information:High,Moderate" -Disabled $falseNew-DlpComplianceRule -Name FCI_PII_content-Low -Policy
   FCI_PII_policy -AccessScope NotInOrganization -BlockAccess $false -ContentPropertyContainsWords
   "Personally Identifiable Information:Low" -Disabled $false -NotifyUser Owner
   ```

   Windows Server FCI includes many built-in properties, including **Personally Identifiable Information** used in this example. The possible values for each property can be different for every organization. The **High**, **Moderate**, and **Low** values used here are only an example. For your organization, you can view the Windows Server FCI classification properties with their possible values in the file Server Resource Manager on the Windows Server-based file server. For more information, see Create a classification property.

When you finish, your policy should have two new rules that both use the **Document properties contain any of these values** condition. This condition won't appear in the UI, though the other conditions, actions, and settings will appear.

One rule blocks access to content where the **Personally Identifiable Information** property equals **High** or **Moderate**. A second rule sends a notification about content where the **Personally Identifiable Information** property equals **Low**.

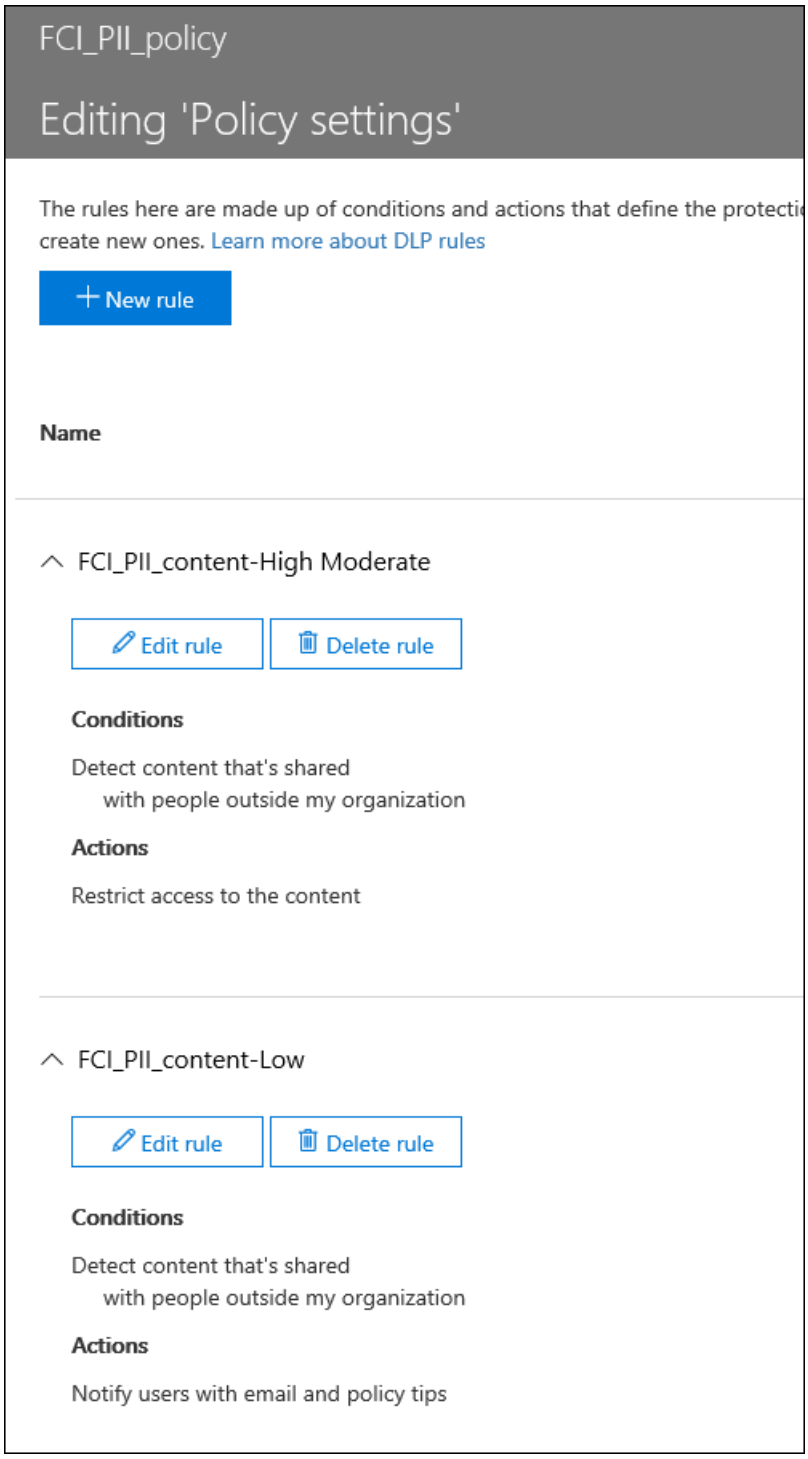

After you create the DLP policy

Doing the steps in the previous sections will create a DLP policy that will quickly detect content with that property, but only if that content is newly uploaded (so that the content's indexed), or if that content is old but just edited (so that the content's re-indexed).

To detect content with that property everywhere, you may want to manually request that your library, site, or site collection be re-indexed, so that the DLP policy is aware of all the content with that property. In SharePoint Online, content is automatically crawled based on a defined crawl schedule. The crawler picks up content that has changed since the last crawl and updates the index. If you need your DLP policy to protect content before the next scheduled crawl, you can take these steps.

Caution

Re-indexing a site can cause a massive load on the search system. Don't re-index your site unless your scenario absolutely requires it.

For more information, see Manually request crawling and re-indexing of a site, a library or a list.

Reindex a site (optional)

1. On the site, choose **Settings** (gear icon in upper right) > **Site Settings**.
2. Under **Search**, choose **Search and offline availability** > **Reindex site**.

More information

- Overview of data loss prevention policies
- Create a DLP policy from a template
- Send notifications and show policy tips for DLP policies
- What the DLP policy templates include
- Sensitive information type entity definitions

View the reports for data loss prevention

11/2/2020 • 3 minutes to read • Edit Online

After you create your data loss prevention (DLP) policies, you'll want to verify that they're working as you intended and helping you to stay compliant. With the DLP reports in the Security & Compliance Center, you can quickly view:

- **DLP policy matches** This report shows the count of DLP policy matches over time. You can filter the report by date, location, policy, or action. You can use this report to:
  - Tune or refine your DLP policies as you run them in test mode. You can view the specific rule that matched the content.
  - Focus on specific time periods and understand the reasons for spikes and trends.
  - Discover business processes that violate your organization's DLP policies.
  - Understand any business impact of the DLP policies by seeing what actions are being applied to content.
  - Verify compliance with a specific DLP policy by showing any matches for that policy.
  - View a list of top users and repeat users who are contributing to incidents in your organization.
  - View a list of the top types of sensitive information in your organization.
- **DLP incidents** This report also shows policy matches over time, like the policy matches report. However, the policy matches report shows matches at a rule level; for example, if an email matched three different rules, the policy matches report shows three different line items. By contrast, the incidents report shows matches at an item level; for example, if an email matched three different rules, the incidents report shows a single line item for that piece of content.

  Because the report counts are aggregated differently, the policy matches report is better for identifying matches with specific rules and fine tuning DLP policies. The incidents report is better for identifying specific pieces of content that are problematic for your DLP policies.
- **DLP false positives and overrides** If your DLP policy allows users to override it or report a false positive, this report shows a count of such instances over time. You can filter the report by date, location, or policy. You can use this report to:
  - Tune or refine your DLP policies by seeing which policies incur a high number of false positives.
  - View the justifications submitted by users when they resolve a policy tip by overriding the policy.
  - Discover where DLP policies conflict with valid business processes by incurring a high number of user overrides.

All DLP reports can show data from the most recent four-month time period. The most recent data can take up to 24 hours to appear in the reports.

You can find these reports in the Security & Compliance Center > **Reports** > **Dashboard**.

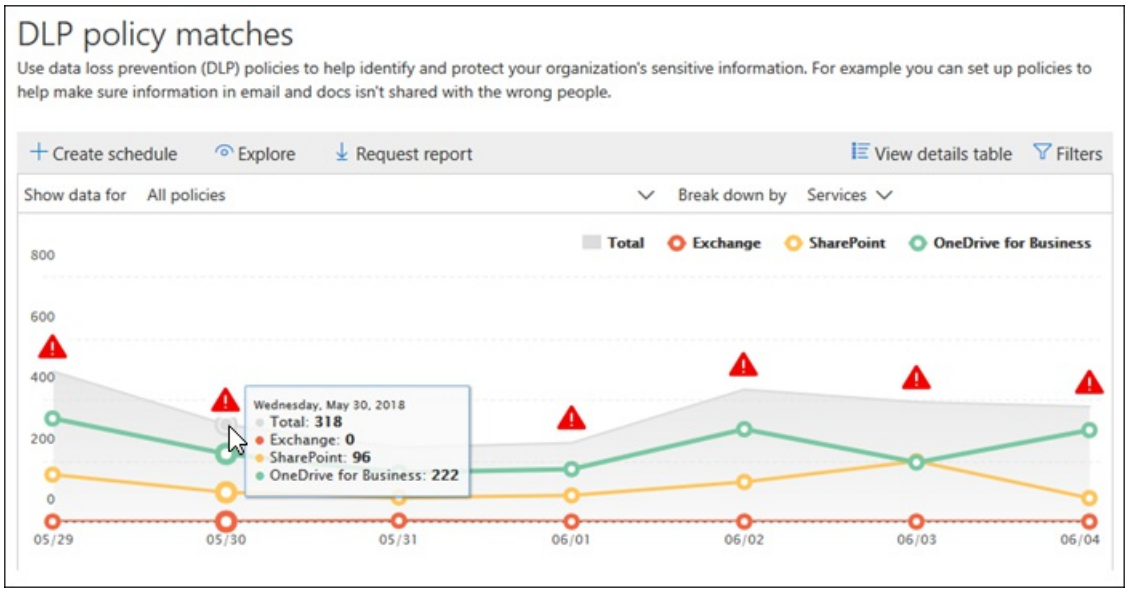

View the justification submitted by a user for an override

If your DLP policy allows users to override it, you can use the false positive and override report to view the text submitted by users in the policy tip.

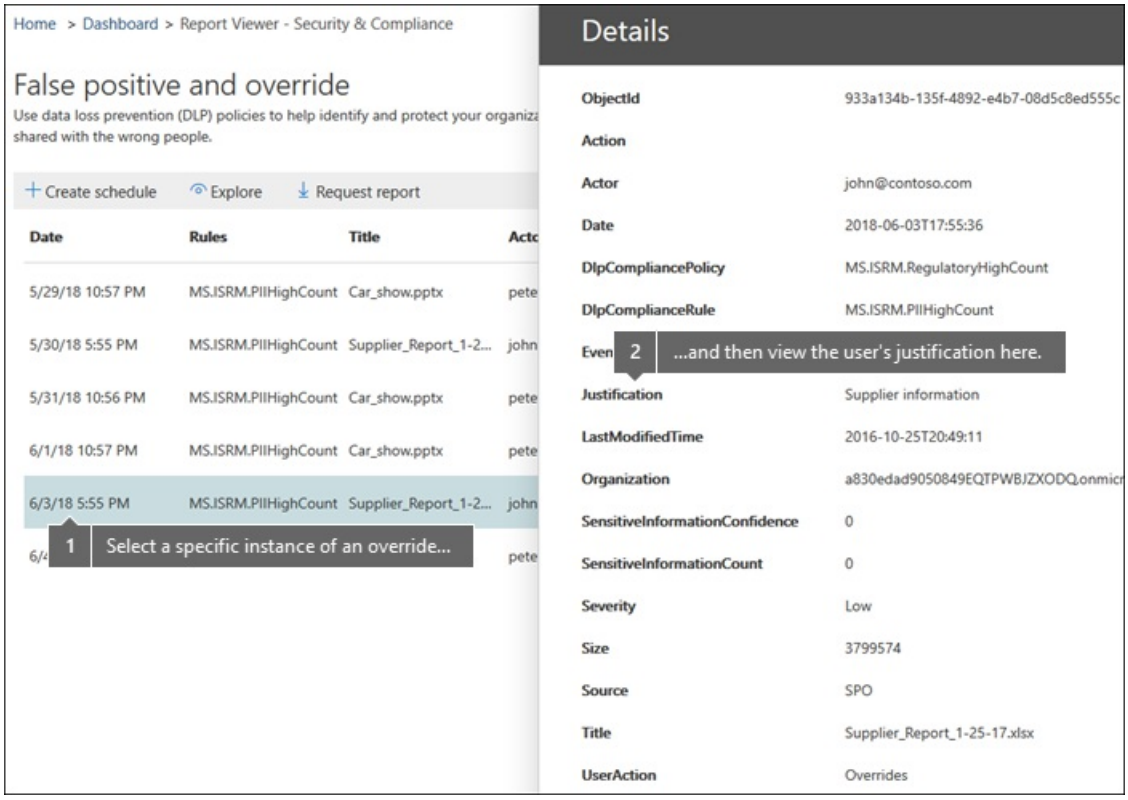

Take action on insights and recommendations

Reports can show insights and recommendations where you can click the red warning icon to see details about potential issues and take possible remedial action.

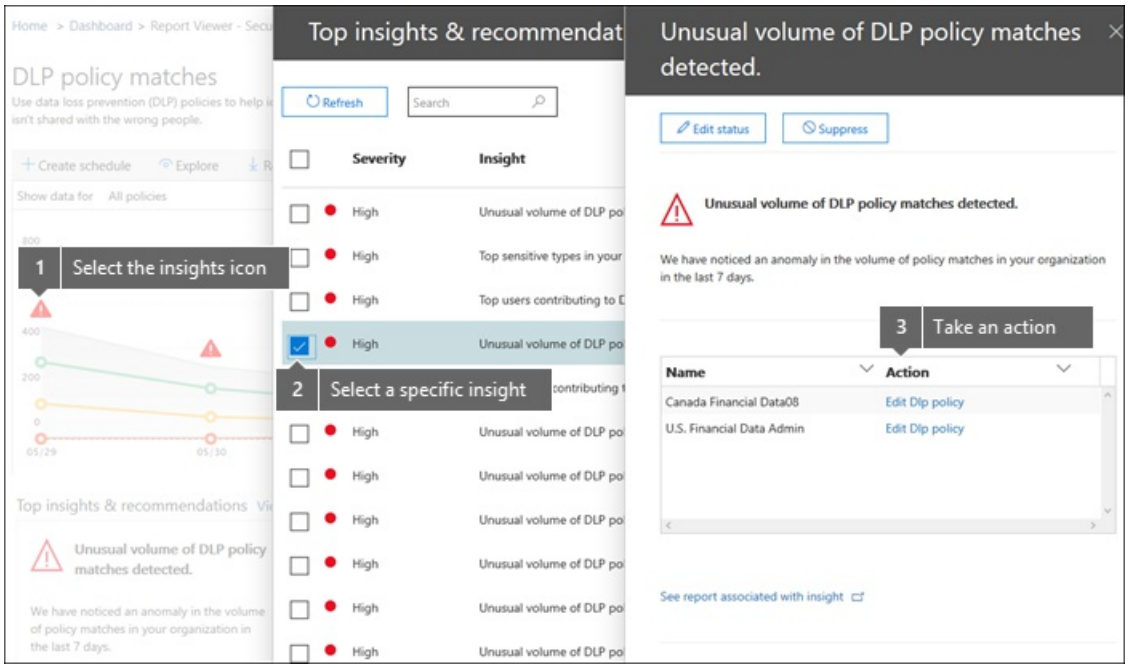

Permissions for DLP reports

To view DLP reports in the Security & Compliance Center, you have to be assigned the:

- **Security Reader** role in the Exchange admin center. By default, this role is assigned to the Organization Management and Security Reader role groups in the Exchange admin center.
- **View-Only DLP Compliance Management** role in the Security & Compliance Center. By default, this role is assigned to the Compliance Administrator, Organization Management, Security Administrator, and Security Reader role groups in the Security & Compliance Center.
- **View-Only Recipients** role in the Exchange admin center. By default, this role is assigned to the Compliance Management, Organization Management, and View-Only Organization Management role groups in the Exchange admin center.

Find the cmdlets for the DLP reports

To use most of the cmdlets for the Security & Compliance Center, you need to:

1. Connect to the Security & Compliance Center using remote PowerShell
2. Use any of these Security & Compliance Center cmdlets

However, DLP reports need pull data from across Office 365, including Exchange Online. For this reason, the cmdlets for the DLP reports are available in Exchange Online Powershell—not in Security & Compliance Center Powershell. Therefore, to use the cmdlets for the DLP reports, you need to:

1. Connect to Exchange Online using remote PowerShell
2. Use any of these cmdlets for the DLP reports:
   - Get-DlpDetectionsReport
   - Get-DlpDetailReport

Form a query to find sensitive data stored on sites

11/2/2020 • 5 minutes to read • Edit Online

Users often store sensitive data, such as credit card numbers, social security numbers, or personal, on their sites, and over time this can expose an organization to significant risk of data loss. Documents stored on sites—including OneDrive for Business sites—could be shared with people outside the organization who shouldn't have access to the information. With data loss prevention (DLP) in SharePoint Online, you can discover documents that contain sensitive data throughout your tenant. After discovering the documents, you can work with the document owners to protect the data. This topic can help you form a query to search for sensitive data.

> **NOTE**
>
> Electronic discovery, or eDiscovery, and DLP are premium features that require SharePoint Online Plan 2.

Forming a basic DLP query

There are three parts that make up a basic DLP query: SensitiveType, count range, and confidence range. As illustrated in the following graphic, **SensitiveType:"<type>"** is required, and both**|<count range>** and**|<confidence range>** are optional.

Sensitive type - required

So what is each part? SharePoint DLP queries typically begin with the property `SensitiveType:"` and an information type name from the sensitive information types inventory, and end with a `"`. You can also use the name of a custom sensitive information type that you created for your organization. For example, you might be looking for documents that contain credit card numbers. In such an instance, you'd use the following format: `SensitiveType:"Credit Card Number"`. Because you didn't include count range or confidence range, the query returns every document in which a credit card number is detected. This is the simplest query that you can run, and it returns the most results. Keep in mind that the spelling and spacing of the sensitive type matters.

Ranges - optional

Both of the next two parts are ranges, so let's quickly examine what a range looks like. In SharePoint DLP queries, a basic range is represented by two numbers separated by two periods, which looks like this: `[number]..[number]`. For instance, if `10..20` is used, that range would capture numbers from 10 through 20. There are many different range combinations and several are covered in this topic.

Let's add a count range to the query. You can use count range to define the number of occurrences of sensitive information a document needs to contain before it's included in the query results. For example, if you want your query to return only documents that contain exactly five credit card numbers, use this: `SensitiveType:"Credit Card Number|5"`. Count range can also help you identify documents that pose high degrees of risk. For example, your organization might consider documents with five or more credit card numbers a high risk. To find documents fitting this criterion, you would use this query: `SensitiveType:"Credit Card Number|5.."`. Alternatively, you can find documents with five or fewer credit card

numbers by using this query: `SensitiveType:"Credit Card Number|..5"`.

Confidence range

Finally, confidence range is the level of confidence that the detected sensitive type is actually a match. The values for confidence range work similarly to count range. You can form a query without including a count range. For example, to search for documents with any number of credit card numbers—as long as the confidence range is 85 percent or higher—you would use this query: `SensitiveType:"Credit Card Number|*|85.."`.

> **IMPORTANT**
> The asterisk (`*`) is a wildcard character that means any value works. You can use the wildcard character (`*`) either in the count range or in the confidence range, but not in a sensitive type.

Additional query properties and search operators available in the eDiscovery Center

DLP in SharePoint also introduces the LastSensitiveContentScan property, which can help you search for files scanned within a specific timeframe. For query examples with the `LastSensitiveContentScan` property, see the Examples of complex queries in the next section.

You can use not only DLP-specific properties to create a query, but also standard SharePoint eDiscovery search properties such as `Author` or `FileExtension`. You can use operators to build complex queries. For the list of available properties and operators, see the Using Search Properties and Operators with eDiscovery blog post.

Examples of complex queries

The following examples use different sensitive types, properties, and operators to illustrate how you can refine your queries to find exactly what you're looking for.

QUERY	EXPLANATION
`SensitiveType:"International Banking Account Number (IBAN)"`	The name might seem strange because it's so long, but it's the correct name for that sensitive type. Make sure to use exact names from the sensitive information types inventory. You can also use the name of a custom sensitive information type that you created for your organization.
`SensitiveType:"Credit Card Number|1..4294967295|1..100"`	This returns documents with at least one match to the sensitive type "Credit Card Number." The values for each range are the respective minimum and maximum values. A simpler way to write this query is `SensitiveType:"Credit Card Number"`, but where's the fun in that?
`SensitiveType:"Credit Card Number| 5..25" AND LastSensitiveContentScan:"8/11/2018..8/13/2018"`	This returns documents with 5-25 credit card numbers that were scanned from August 11, 2018 through August 13, 2018.
`SensitiveType:"Credit Card Number| 5..25" AND LastSensitiveContentScan:"8/11/2018..8/13/2018" NOT FileExtension:XLSX`	This returns documents with 5-25 credit card numbers that were scanned from August 11, 2018 through August 13, 2018. Files with an XLSX extension aren't included in the query results. `FileExtension` is one of many properties that you can include in a query. For more information, see Using Search Properties and Operators with eDiscovery.
`SensitiveType:"Credit Card Number" OR SensitiveType:"U.S. Social Security Number (SSN)"`	This returns documents that contain either a credit card number or a social security number.

Examples of queries to avoid

Not all queries are created equal. The following table gives examples of queries that don't work with DLP in SharePoint and describes why.

UNSUPPORTED QUERY	REASON	
`SensitiveType:"Credit Card Number|.."`	You must add at least one number.	
`SensitiveType:"NotARule"`	"NotARule" isn't a valid sensitive type name. Only names in the sensitive information types inventory work in DLP queries.	
`SensitiveType:"Credit Card Number|0"`	Zero isn't valid as either the minimum value or the maximum value in a range.	
`SensitiveType:"Credit Card Number"`	It's might be difficult to see, but there's extra white space between "Credit" and "Card" that makes the query invalid. Use exact sensitive type names from the sensitive information types inventory.	
`SensitiveType:"Credit Card Number|1. .3"`	The two-period portion shouldn't be separated by a space.	
`SensitiveType:"Credit Card Number| |1..|80.."`	There are too many pipe delimiters (	). Follow this format instead: `SensitiveType: "Credit Card Number|1..|80.."`
`SensitiveType:"Credit Card Number|1..|80..101"`	Because confidence values represent a percentage, they can't exceed 100. Choose a number from 1 through 100 instead.	

For more information

- Sensitive information type entity definitions
- Run a Content Search
- Keyword queries and search conditions for Content Search

How DLP works between the Security & Compliance Center and Exchange admin center

5/5/2020 • 2 minutes to read • Edit Online

In Office 365, you can create a data loss prevention (DLP) policy in two different admin centers:

- In the **Security & Compliance Center**, you can create a single DLP policy to help protect content in SharePoint, OneDrive, Exchange, and now Microsoft Teams. When possible, we recommend that you create a DLP policy here. For more information, see DLP in the Security & Compliance Center.
- In the **Exchange admin center**, you can create a DLP policy to help protect content only in Exchange. This policy can use Exchange mail flow rules (also known as transport rules), so it has more options specific to handling email. For more information, see DLP in the Exchange admin center.

DLP polices created in these admin centers work side by side - this topic explains how.

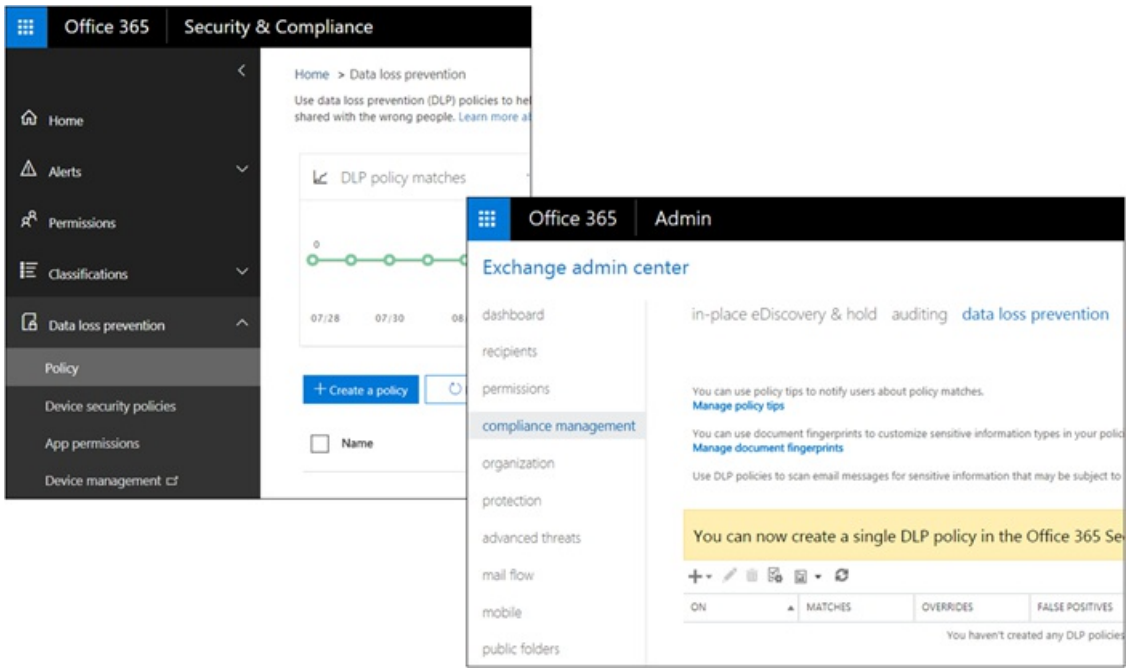

How DLP in the Security & Compliance Center works with DLP and mail flow rules in the Exchange admin center

After you create a DLP policy in the Security & Compliance Center, the policy is deployed to all of the locations included in the policy. If the policy includes Exchange Online, the policy's synced there and enforced in exactly the same way as a DLP policy created in the Exchange admin center.

If you've created DLP policies in the Exchange admin center, those policies will continue to work side by side with any policies for email that you create in the Security & Compliance Center. But note that rules created in the Exchange admin center take precedence. All Exchange mail flow rules are processed first, and then the DLP rules from the Security & Compliance Center are processed.

This means that:

- Messages that are blocked by Exchange mail flow rules won't get scanned by DLP rules created in the Security & Compliance Center.

- If an Exchange mail flow rule modifies a message in a way that causes it to match a DLP policy in the Security & Compliance Center - such as adding external users - then the DLP rules will detect this and enforce the policy as needed.

Also note that Exchange mail flow rules that use the "stop processing" action don't affect the processing of DLP rules in the Security & Compliance Center - they'll still be processed.

Policy tips in the Security & Compliance Center vs. the Exchange admin center

Policy tips can work either with DLP policies and mail flow rules created in the Exchange admin center, or with DLP policies created in the Security & Compliance Center, but not both. This is because these policies are stored in different locations, but policy tips can draw only from a single location.

If you've configured policy tips in the Exchange admin center, any policy tips that you configure in the Security & Compliance Center won't appear to users in Outlook on the web and Outlook 2013 and later until you turn off the tips in the Exchange admin center. This ensures that your current Exchange mail flow rules will continue to work until you choose to switch over to the Security & Compliance Center.

Note that while policy tips can draw only from a single location, email notifications are always sent, even if you're using DLP policies in both the Security & Compliance Center and the Exchange admin center.

Configure and view alerts for DLP polices (preview)

11/2/2020 • 5 minutes to read • Edit Online

This article shows you how to define rich alert policies that are linked to your data loss prevention (DLP) policies. You'll see how to use the new DLP alert management dashboard in the Microsoft 365 compliance center to view alerts, events, and associated metadata for DLP policy violations.

Features

The following features are part of this preview:

- **DLP alert management dashboard**: In the Microsoft 365 compliance center, this dashboard shows alerts for DLP policies that are enforced on the following workloads:
  - Exchange
  - SharePoint
  - OneDrive
  - Teams
  - Devices
- **Advanced alert configuration options**: These options are part of the DLP policy authoring flow. Use them to create rich alert configurations. You can create a single-event alert or an aggregated alert, based on the number of events or the size of the leaked data.

Before you begin

Before you begin, make sure you have the necessary prerequisites:

- Licensing for the DLP alerts management dashboard
- Licensing for alert configuration options
- Roles

Licensing for the DLP alert management dashboard

All eligible tenants for Office 365 DLP can access the new DLP alert management dashboard. To get started, you should be eligible for Office 365 DLP for Exchange Online, SharePoint Online, and OneDrive for Business. For more information about the licensing requirements for Office 365 DLP, see Which licenses provide the rights for a user to benefit from the service?.

Customers who participate in the Endpoint DLP public preview or who are eligible for Teams DLP will see their endpoint DLP policy alerts and Teams DLP policy alerts in the DLP alert management dashboard.

Licensing for alert configuration options

- **Single-event alert configuration**: Organizations that have an E1, F1, or G1 subscription or an E3 or G3 subscription can create alert policies only where an alert is triggered every time an activity occurs.
- **Aggregated alert configuration**: To configure aggregate alert policies based on a threshold, you must have either of the following configurations:
  - An E5 or G5 subscription
  - An E1, F1, or G1 subscription or an E3 or G3 subscription that includes one of the following features:
    - Office 365 Advanced Threat Protection Plan 2
    - Microsoft 365 E5 Compliance

- Microsoft 365 eDiscovery and Audit add-on license

Roles

If you want to view the DLP alert management dashboard or to edit the alert configuration options in a DLP policy, you must be a member of one of these role groups:

- Compliance Administrator
- Compliance Data Administrator
- Security Administrator
- Security Operator
- Security Reader

To access the DLP alert management dashboard, you need the Manage alerts role and either of the following roles:

- DLP Compliance Management
- View-Only DLP Compliance Management

Alert configuration experience

If you're eligible for aggregated alert configuration options, then you see the following options inline in the DLP policy authoring experience.

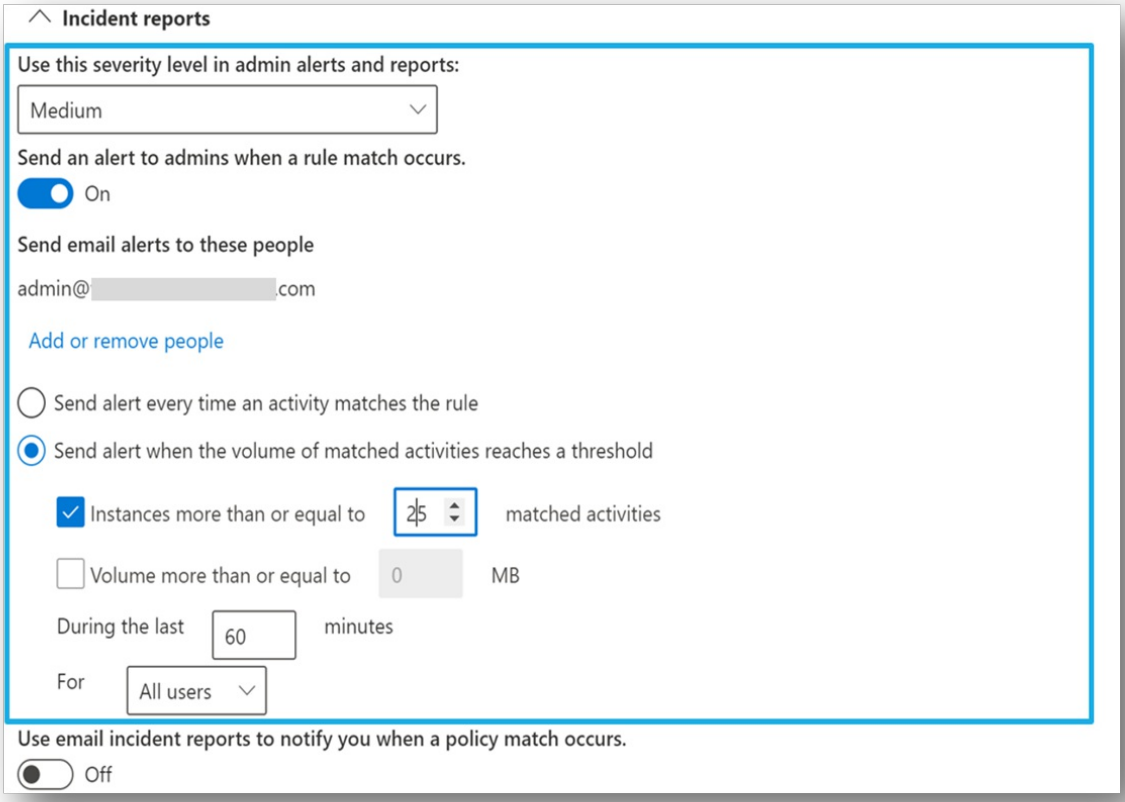

You can use these alert configuration options to configure a setting that defines how often a DLP rule match can occur before an alert is triggered. This configuration allows you to set up a policy to generate an alert every time an activity matches the policy conditions or when a certain threshold is exceeded, based on the number of activities or based on the volume of exfiltrated data.

If you're eligible for single-event alert configuration options, then you see the following alert configuration option in the DLP policy authoring experience. Use this option to create an alert that's raised every time a DLP

rule match happens because of a user activity.

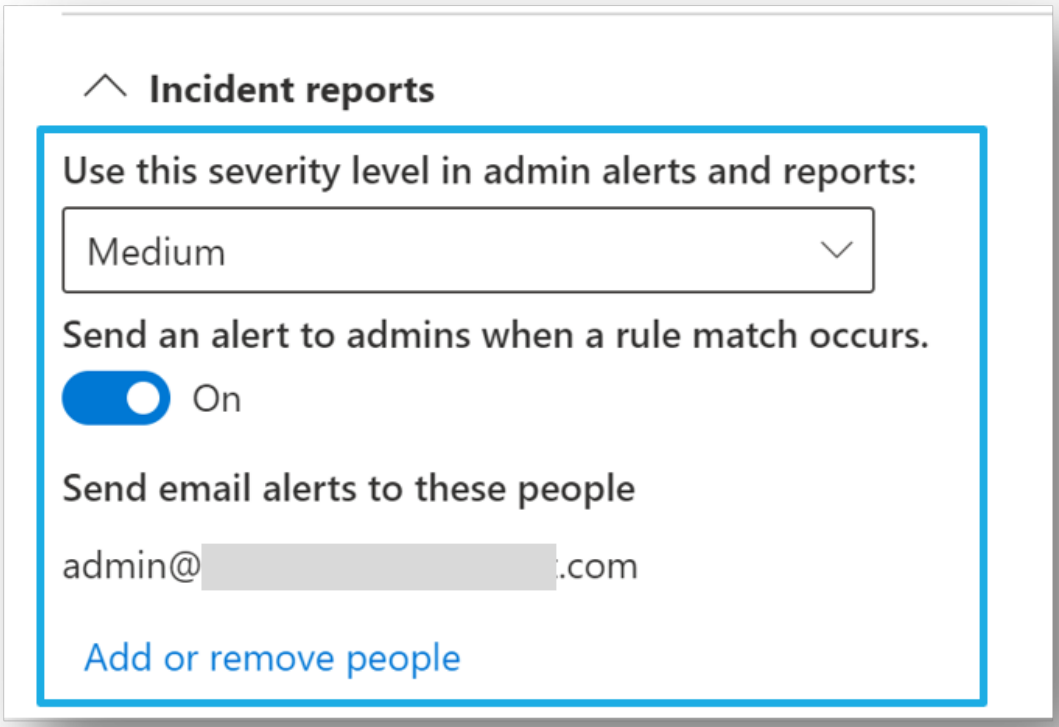

DLP alert management dashboard

To work with the DLP alert management dashboard:

1. In the Microsoft 365 compliance center, go to **Data Loss Prevention**.
2. Select the **Alerts** tab to view the DLP alerts dashboard.
   - Choose filters to refine the list of alerts. Choose **Customize columns** to list the properties you want to see. You can also choose to sort the alerts in ascending or descending order in any column.
   - Select an alert to see details:

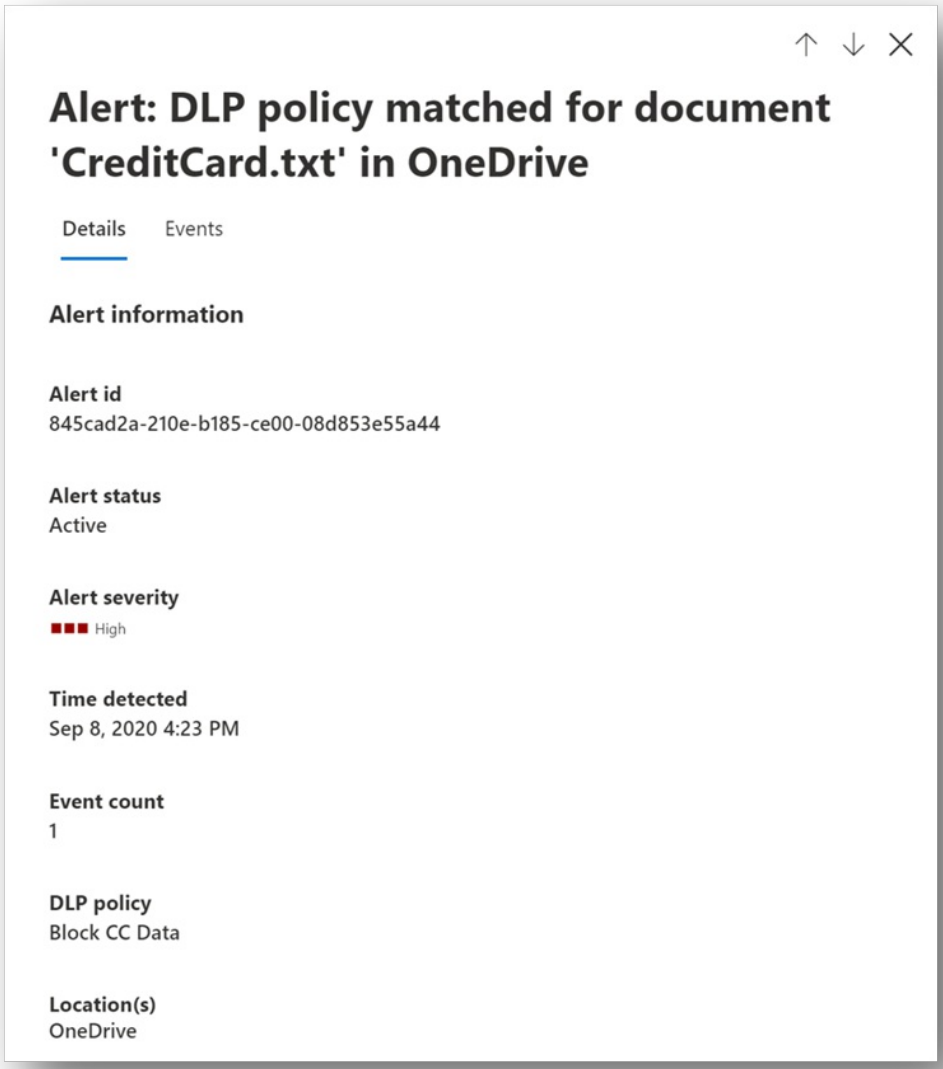

3. Select the **Events** tab to view all of the events associated with the alert. You can choose a particular event to view its details. The following table shows some of the event details.

CATEGORY	PROPERTY NAME	DESCRIPTION	APPLICABLE EVENT TYPES
Event details			
	Id	Unique ID associated with the event	All events
	Location	Workload where the event was detected	All events
	Time of activity	Time of the user activity that caused the DLP violation	All events
Impacted entities			
	User	User who caused the DLP violation	All events

CATEGORY	PROPERTY NAME	DESCRIPTION	APPLICABLE EVENT TYPES
	Hostname	Host name of the machine where the DLP violation was detected	Devices events
	IP address	IP address of the machine	Devices events
	File path	Absolute path of the file involved in the violation	SharePoint, OneDrive, and Devices events
	Email recipients	Recipients of the email that violated the DLP policy	Exchange events
	Email subject	Subject of the email that violated the DLP policy	Exchange events
	Email attachments	Names of the attachments in the email that violated the DLP policy	Exchange events
	Site owner	Name of the site owner	SharePoint and OneDrive events
	Site URL	Full URL of the SharePoint or OneDrive site	SharePoint and OneDrive events
	File created	Time of file creation	SharePoint and OneDrive events
	File last modified	Time of the last modification of the file	SharePoint and OneDrive events
	File size	Size of the file	SharePoint and OneDrive events
	File owner	Owner of the file	SharePoint and OneDrive events
Policy details			
	DLP policy matched	Name of the DLP policy that was matched	All events
	Rule matched	Name of the DLP rule in the DLP policy that was matched	All events
	Sensitive info types detected	Sensitive information types that were detected as a part of the DLP policy	All events

CATEGORY	PROPERTY NAME	DESCRIPTION	APPLICABLE EVENT TYPES
	Actions taken	Actions taken as a part of the matched DLP policy	All events
	User overrode policy	Whether the user overrode the policy through the policy tip	All events
	Override justification text	Justification provided to override the policy tip	All events

4. Select the **Sensitive Info Types** tab to view details about the sensitive information types detected in the content. Details include confidence and count.

5. After you investigate the alert, choose **Manage alert** to change the status (**Active**, **Investigating**, **Dismissed**, or **Resolved**). You can also add comments and assign the alert to someone in your organization.

   - To see the history of workflow management, choose **Management log**.
   - After you take the required action for the alert, set the status of the alert to **Resolved**.

Use sensitivity labels as conditions in DLP policies (preview)

11/2/2020 • 2 minutes to read • Edit Online

You can use sensitivity labels as a condition in DLP policies for these location:

- Exchange Online email messages
- SharePoint Online
- OneDrive for Business sites
- Windows 10 devices

Sensitivity labels appear as an option in the **Content contains** list.

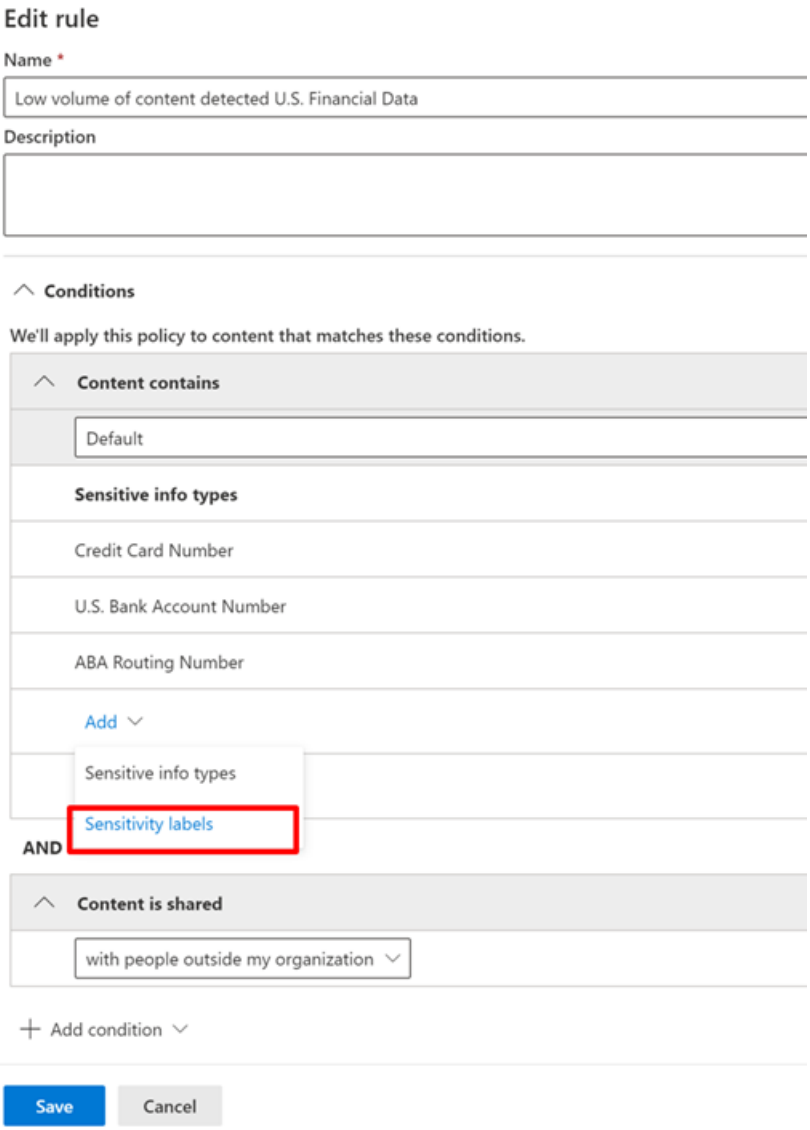

IMPORTANT

Sensitivity Labels as a condition will not be available if you have selected **Teams chat and channel messages** as a location to apply the DLP policy.

Supported items, scenarios, and policy tips

You can use sensitivity labels as conditions on these items and in these scenarios.

Supported items

SERVICE	ITEM TYPE	AVAILABLE TO POLICY TIP	ENFORCEABLE
Exchange	email message	yes	yes
Exchange	email attachment	no *	no *
SharePoint Online	items in SharePoint Online	yes	yes
OneDrive for Business	items	yes	yes
Teams	Teams and channel messages	not applicable	not applicable
Teams	attachments	yes **	yes **
Windows 10 devices (preview)	items	yes	yes
MCAS (preview)	items	yes	yes

* DLP detection of sensitivity labels on emails are supported. DLP detection of sensitivity labeled email attachments are not.

** Attachments sent in Teams over 1:1 chat or channels are automatically uploaded to OneDrive for Business and SharePoint. So if SharePoint Online or OneDrive for Business are included as locations in your DLP policy, then labeled attachments sent in Teams will be automatically included in the scope of this condition. Teams as a location does not need to be selected in the DLP policy.

Supported scenarios

- DLP Admin will be able to see a list of all sensitivity labels in the tenant when they choose to include one or more sensitivity labels as a condition.
- Using sensitivity labels as a condition is supported across all workloads as indicated in the support matrix above.
- DLP policy tips will continue to be shown across workloads (except Outlook Win32) for DLP policies which contain sensitivity label as a condition.
- Sensitivity labels will also appear as a part of the incident report email if a DLP policy with sensitivity label as a condition is matched.
- Sensitivity label details will also be shown in the DLP rule match audit log for a DLP policy match which contains sensitivity label as a condition.

Support policy tips

WORKLOAD	POLICY TIPS SUPPORTED/NOT SUPPORTED
OWA	supported

WORKLOAD	POLICY TIPS SUPPORTED/NOT SUPPORTED
Outlook Win 32	not supported
SharePoint	supported
OneDrive for Business	supported
endpoint devices	not supported

Use data loss prevention policies for non-Microsoft cloud apps (preview)

11/2/2020 • 2 minutes to read • Edit Online

Data loss prevention (DLP) policies to non-Microsoft cloud apps are part of the Microsoft 365 DLP suite of features; using these features, you can discover and protect sensitive items across Microsoft 365 services. For more information about all Microsoft DLP offerings, see Overview of data loss prevention.

You can use DLP policies to non-Microsoft cloud apps to monitor and detect when sensitive items are used and shared via non-Microsoft cloud apps. Using these policies gives you the visibility and control that you need to ensure that they're correctly used and protected, and it helps prevent risky behavior that might compromise them.

Before you begin

SKU/subscriptions licensing

Before you start using DLP policies to non-Microsoft cloud apps, confirm your Microsoft 365 subscription and any add-ons. To access and use this functionality, you must have one of these subscriptions or add-ons:

- Microsoft 365 E5
- Microsoft 365 E5 Compliance
- Microsoft 365 E5 Security

Prepare your Cloud App Security environment

DLP policies to non-Microsoft cloud apps use Cloud App Security DLP capabilities. To use it, you should prepare your Cloud App Security environment. For instructions, see Set instant visibility, protection, and governance actions for your apps.

Connect a non-Microsoft cloud app

To use DLP policy to a specific non-Microsoft cloud app, the app must be connected to Cloud App Security. For information, see:

- Connect Box
- Connect Dropbox
- Connect G-Suite
- Connect Salesforce
- Connect Cisco Webex

After you connect your cloud apps to Cloud App Security, you can create Microsoft 365 DLP policies for them.

> **NOTE**
> It's also possible to use Microsoft Cloud App Security to create DLP policies to Microsoft cloud apps. However, it's recommended to use Microsoft 365 to create and manage DLP policies to Microsoft cloud apps.

Create a DLP policy to a non-Microsoft cloud app

When you select a location for the DLP policy, turn on the **Microsoft Cloud App Security** location.

- To select a specific app or instance, select **Choose instance**.
- If you don't select an instance, the policy uses all connected apps in your Microsoft Cloud App Security tenant.

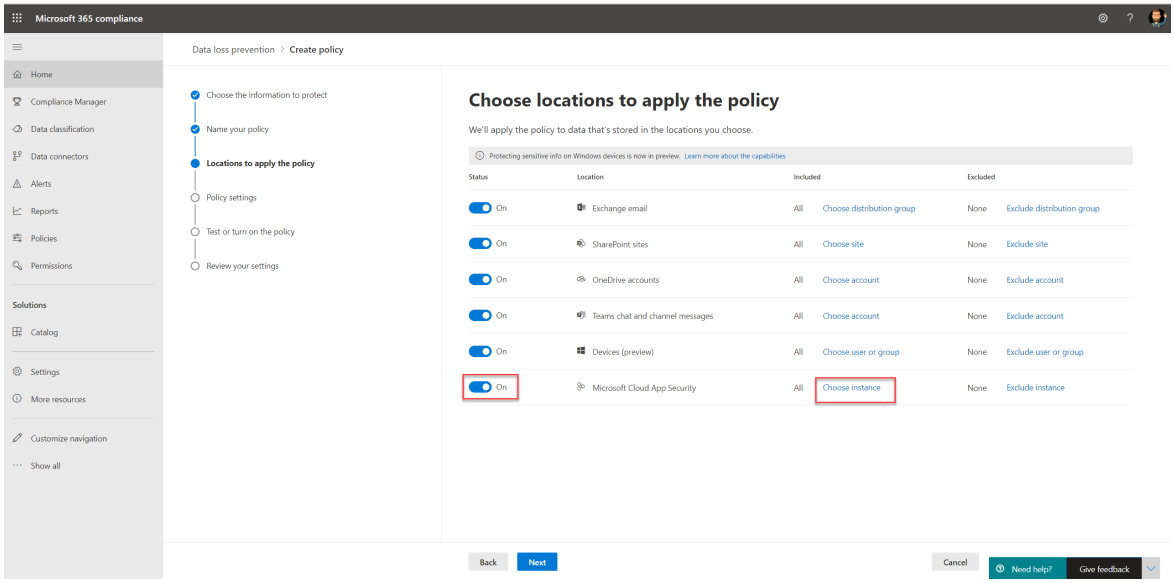

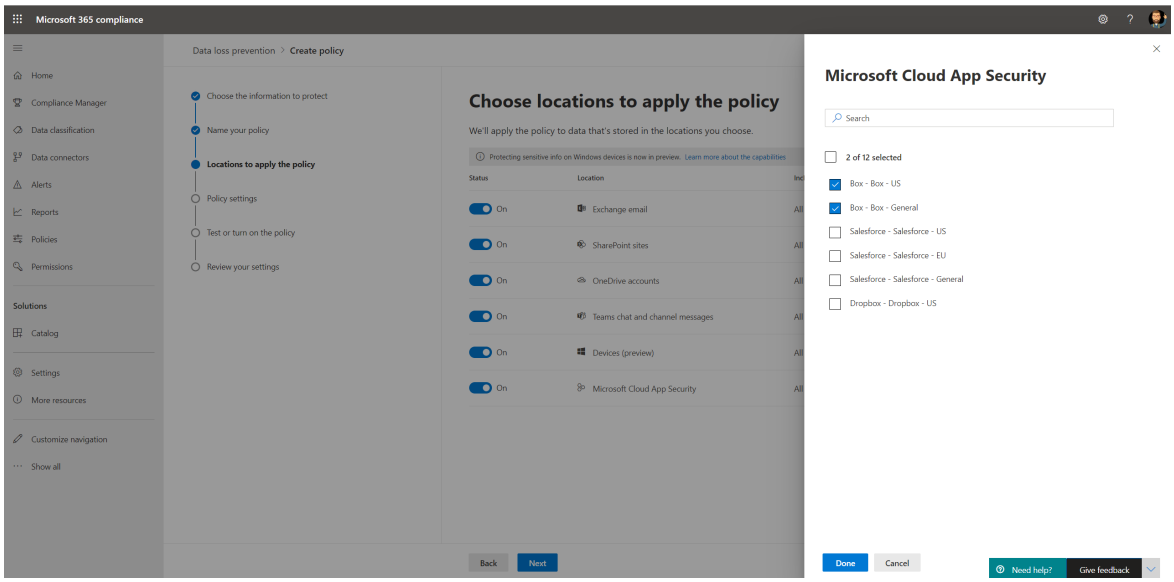

You can choose various actions for every supported non-Microsoft cloud app. For every app, there are different possible actions (depends on the cloud app API).

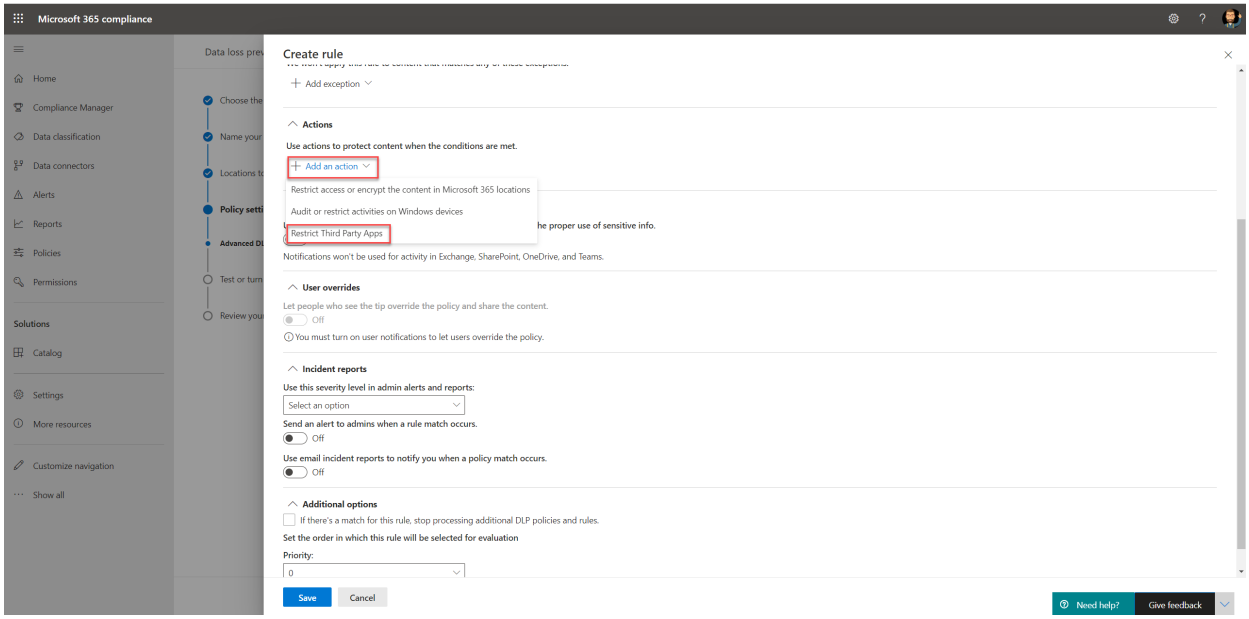

When you create a rule in the DLP policy, you can select an action for non-Microsoft cloud apps. To restrict third-party apps, select **Restrict Third Party Apps**.

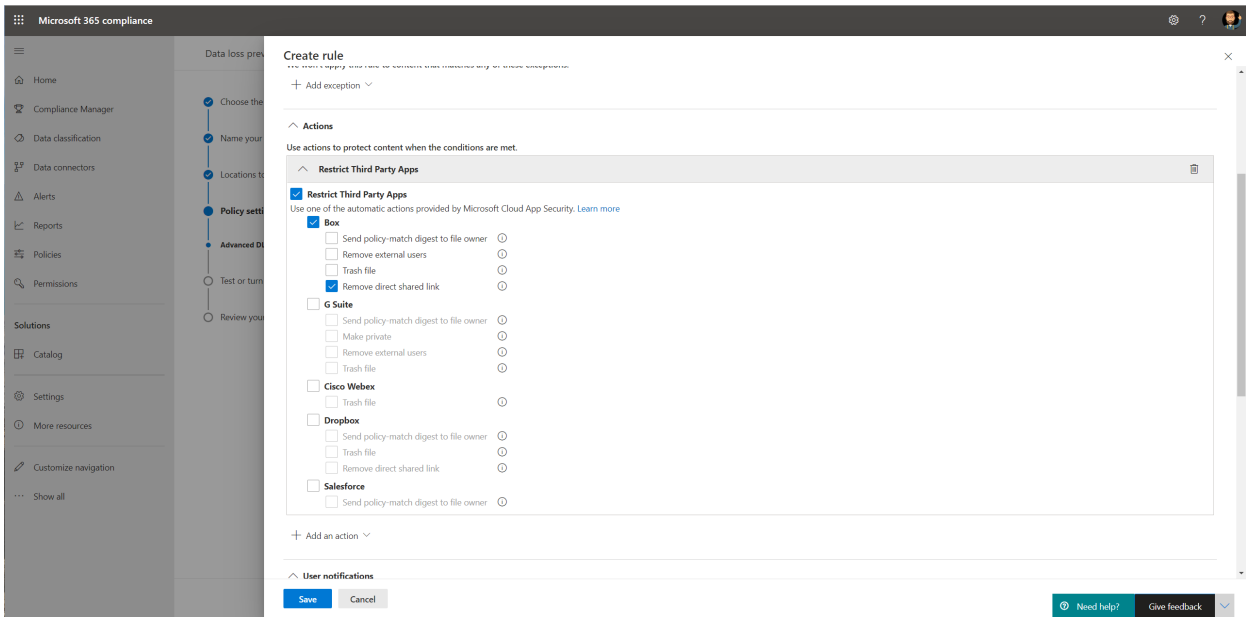

For information about creating and configuring DLP policies, see Create test and tune a DLP policy.

See Also

- Create test and tune a DLP policy
- Get started with the default DLP policy
- Create a DLP policy from a template

DLP policy conditions, exceptions, and actions (preview)

2/18/2021 • 8 minutes to read • Edit Online

Conditions and exceptions in DLP policies identify sensitive items that the policy is applied to. Actions define what happens as a consequence of a condition of exception being met.

- Conditions define what to include
- Exceptions define what to exclude.
- Actions define what happens as a consequence of condition or exception being met

Most conditions and exceptions have one property that supports one or more values. For example, if the DLP policy is being applied to Exchange emails, the **The sender** is condition requires the sender of the message. Some conditions have two properties. For example, the **A message header includes any of these words** condition requires one property to specify the message header field, and a second property to specify the text to look for in the header field. Some conditions or exceptions don't have any properties. For example, the **Attachment is password protected** condition simply looks for attachments in messages that are password protected.

Actions typically require additional properties. For example, when the DLP policy rule redirects a message, you need to specify where the message is redirected to.

Conditions and exceptions for DLP policies

The tables in the following sections describe the conditions and exceptions that are available in DLP.

- Senders
- Recipients
- Message subject or body
- Attachments
- Message headers
- Message properties

Senders

CONDITION OR EXCEPTION IN DLP	CONDITION/EXCEPTION PARAMETERS IN MICROSOFT 365 POWERSHELL	PROPERTY TYPE	DESCRIPTION
Sender is	condition: *From* exception: *ExceptIfFrom*	Addresses	Messages that are sent by the specified mailboxes, mail users, mail contacts, or Microsoft 365 groups in the organization.
Sender IP address is	condition: *SenderIPRanges* exception: *ExceptIfSenderIPRanges*	IPAddressRanges	Messages where the sender's IP address matches the specified IP address, or falls within the specified IP address range.

CONDITION OR EXCEPTION IN DLP	CONDITION/EXCEPTION PARAMETERS IN MICROSOFT 365 POWERSHELL	PROPERTY TYPE	DESCRIPTION
Sender address contains words	condition: *FromAddressContainsWords* exception: *ExceptIfFromAddressContainsWords*	Words	Messages that contain the specified words in the sender's email address.
Sender address matches patterns	condition: *FromAddressMatchesPatterns* exception: *ExceptFromAddressMatchesPatterns*	Patterns	Messages where the sender's email address contains text patterns that match the specified regular expressions.
Sender domain is	condition: *SenderDomainIs* exception: *ExceptIfSenderDomainIs*	DomainName	Messages where the domain of the sender's email address matches the specified value. If you need to find sender domains that *contain* the specified domain (for example, any subdomain of a domain), use **The sender address matches**(*FromAddressMatchesPatterns*) condition and specify the domain by using the syntax: '.domain.com$'.
Sender scope	condition: *FromScope* exception: *ExceptIfFromScope*	UserScopeFrom	Messages that are sent by either internal or external senders.

Recipients

CONDITION OR EXCEPTION IN DLP	CONDITION/EXCEPTION PARAMETERS IN MICROSOFT 365 POWERSHELL	PROPERTY TYPE	DESCRIPTION
Recipient is	condition: *SentTo* exception: *ExceptIfSentTo*	Addresses	Messages where one of the recipients is the specified mailbox, mail user, or mail contact in the organization. The recipients can be in the **To**, **Cc**, or **Bcc** fields of the message.
Recipient domain is	condition: *RecipientDomainIs* exception: *ExceptIfRecipientDomainIs*	DomainName	Messages where the domain of the sender's email address matches the specified value.

CONDITION OR EXCEPTION IN DLP	CONDITION/EXCEPTION PARAMETERS IN MICROSOFT 365 POWERSHELL	PROPERTY TYPE	DESCRIPTION
Recipient address contains words	condition: *RecipientAddressContainsWords* exception: *ExceptIfRecipientAddressContainsWords*	Words	Messages that contain the specified words in the recipient's email address. **Note**: This condition doesn't consider messages that are sent to recipient proxy addresses. It only matches messages that are sent to the recipient's primary email address.
Recipient address matches patterns	condition: *RecipientAddressMatchesPatterns* exception: *ExceptIfRecipientAddressMatchesPatterns*	Patterns	Messages where a recipient's email address contains text patterns that match the specified regular expressions. **Note**: This condition doesn't consider messages that are sent to recipient proxy addresses. It only matches messages that are sent to the recipient's primary email address.
Sent to member of	condition: *SentToMemberOf* exception: *ExceptIfSentToMemberOf*	Addresses	Messages that contain recipients who are members of the specified distribution group, mail-enabled security group, or Microsoft 365 group. The group can be in the **To**, **Cc**, or **Bcc** fields of the message.

Message subject or body

CONDITION OR EXCEPTION IN DLP	CONDITION/EXCEPTION PARAMETERS IN MICROSOFT 365 POWERSHELL	PROPERTY TYPE	DESCRIPTION
Subject contains words or phrases	condition: *SubjectContainsWords* exception: *ExceptIf SubjectContainsWords*	Words	Messages that have the specified words in the Subject field.
Subject matches patterns	condition: *SubjectMatchesPatterns* exception: *ExceptIf SubjectMatchesPatterns*	Patterns	Messages where the Subject field contain text patterns that match the specified regular expressions.
Content contains	condition: *ContentContainsSensitiveInformation* exception *ExceptIfContentContainsSensitiveInformation*	SensitiveInformationTypes	Messages or documents that contain sensitive information as defined by data loss prevention (DLP) policies.

CONDITION OR EXCEPTION IN DLP	CONDITION/EXCEPTION PARAMETERS IN MICROSOFT 365 POWERSHELL	PROPERTY TYPE	DESCRIPTION
Subject or Body matches pattern	condition: *SubjectOrBodyMatchesPatterns* exception: *ExceptIfSubjectOrBodyMatchesPatterns*	Patterns	Messages where the subject field or message body contains text patterns that match the specified regular expressions.
Subject or Body contains words	condition: *SubjectOrBodyContainsWords* exception: *ExceptIfSubjectOrBodyContainsWords*	Words	Messages that have the specified words in the subject field or message body

Attachments

CONDITION OR EXCEPTION IN DLP	CONDITION/EXCEPTION PARAMETERS IN MICROSOFT 365 POWERSHELL	PROPERTY TYPE	DESCRIPTION
Attachment is password protected	condition: *DocumentIsPasswordProtected* exception: *ExceptIfDocumentIsPasswordProtected*	none	Messages where an attachment is password protected (and therefore can't be scanned). Password detection only works for Office documents, .zip files, and .7z files.
Attachment's file extension is	condition: *ContentExtensionMatchesWords* exception: *ExceptIfContentExtensionMatchesWords*	Words	Messages where an attachment's file extension matches any of the specified words.
Any email attachment's content could not be scanned	condition: *DocumentIsUnsupported* exception: *ExceptIfDocumentIsUnsupported*	n/a	Messages where an attachment isn't natively recognized by Exchange Online.
Any email attachment's content didn't complete scanning	condition: *ProcessingLimitExceeded* exception: *ExceptIfProcessingLimitExceeded*	n/a	Messages where the rules engine couldn't complete the scanning of the attachments. You can use this condition to create rules that work together to identify and process messages where the content couldn't be fully scanned.

CONDITION OR EXCEPTION IN DLP	CONDITION/EXCEPTION PARAMETERS IN MICROSOFT 365 POWERSHELL	PROPERTY TYPE	DESCRIPTION
Document name contains words	condition: *DocumentNameMatchesWords* exception: *ExceptIfDocumentNameMatchesWords*	Words	Messages where an attachment's file name matches any of the specified words.
Document name matches patterns	condition: *DocumentNameMatchesPatterns* exception: *ExceptIfDocumentNameMatchesPatterns*	Patterns	Messages where an attachment's file name contains text patterns that match the specified regular expressions.
Document property is	condition: *ContentPropertyContainsWords* exception: *ExceptIfContentPropertyContainsWords*	Words	Messages or documents where an attachment's file extension matches any of the specified words.
Document size equals or is greater than	condition: *DocumentSizeOver* exception: *ExceptIfDocumentSizeOver*	Size	Messages where any attachment is greater than or equal to the specified value.

Message Headers

CONDITION OR EXCEPTION IN DLP	CONDITION/EXCEPTION PARAMETERS IN MICROSOFT 365 POWERSHELL	PROPERTY TYPE	DESCRIPTION
Header contains words or phrases	condition: *HeaderContainsWords* exception: *ExceptIfHeaderContainsWords*	Hash Table	Messages that contain the specified header field, and the value of that header field contains the specified words.
Header matches patterns	condition: *HeaderMatchesPatterns* exception: *ExceptIfHeaderMatchesPatterns*	Hash Table	Messages that contain the specified header field, and the value of that header field contains the specified regular expressions.

Message properties

CONDITION OR EXCEPTION IN DLP	CONDITION/EXCEPTION PARAMETERS IN MICROSOFT 365 POWERSHELL	PROPERTY TYPE	DESCRIPTION

CONDITION OR EXCEPTION IN DLP	CONDITION/EXCEPTION PARAMETERS IN MICROSOFT 365 POWERSHELL	PROPERTY TYPE	DESCRIPTION
Message size over	condition: *MessageSizeOver* exception: *ExceptIfMessageSizeOver*	Size	Messages where the total size (message plus attachments) is greater than or equal to the specified value. **Note**: Message size limits on mailboxes are evaluated before mail flow rules. A message that's too large for a mailbox will be rejected before a rule with this condition is able to act on the message.
With importance	condition: *WithImportance* exception: *ExceptIfWithImportance*	Importance	Messages that are marked with the specified importance level.
Content character set contains words	condition: *ContentCharacterSetContainsWords* *ExceptIfContentCharacterSetContainsWords*	CharacterSets	Messages that have any of the specified character set names.
Has sender override	condition: *HasSenderOverride* exception: *ExceptIfHasSenderOverride*	n/a	Messages where the sender has chosen to override a data loss prevention (DLP) policy. For more information about DLP policies see Data loss prevention.
Message type matches	condition: *MessageTypeMatches* exception: *ExceptIfMessageTypeMatches*	MessageType	Messages of the specified type.

Actions for DLP policies

This table describes the actions that are available in DLP.

ACTION IN DLP	ACTION PARAMETERS IN MICROSOFT 365 POWERSHELL	PROPERTY TYPE	DESCRIPTION

ACTION IN DLP	ACTION PARAMETERS IN MICROSOFT 365 POWERSHELL	PROPERTY TYPE	DESCRIPTION
Set header	SetHeader	First property: *Header Name* Second property: *Header Value*	The SetHeader parameter specifies an action for the DLP rule that adds or modifies a header field and value in the message header. This parameter uses the syntax "HeaderName:HeaderValue". You can specify multiple header name and value pairs separated by commas
Remove header	RemoveHeader	First property: *MessageHeaderField* Second property: *String*	The RemoveHeader parameter specifies an action for the DLP rule that removes a header field from the message header. This parameter uses the syntax “HeaderName” or "HeaderName:HeaderValue". You can specify multiple header names or header name and value pairs separated by commas
Redirect the message to specific users	*RedirectMessageTo*	Addresses	Redirects the message to the specified recipients. The message isn't delivered to the original recipients, and no notification is sent to the sender or the original recipients.
Forward the message for approval to sender’s manager	Moderate	First property: *ModerateMessageByManager* Second property: *Boolean*	The Moderate parameter specifies an action for the DLP rule that sends the email message to a moderator. This parameter uses the syntax: @{ModerateMessageByManager = <$true | $false>;
Forward the message for approval to specific approvers	Moderate	First property: *ModerateMessageByUser* Second property: *Addresses*	The Moderate parameter specifies an action for the DLP rule that sends the email message to a moderator. This parameter uses the syntax: @{ ModerateMessageByUser = @("emailaddress1","emailaddress2",..."emailaddressN")}

ACTION IN DLP	ACTION PARAMETERS IN MICROSOFT 365 POWERSHELL	PROPERTY TYPE	DESCRIPTION
Add recipient	AddRecipients	First property: *Field* Second property: *Addresses*	Adds one or more recipients to the To/Cc/Bcc field of the message. This parameter uses the syntax: @{<AddToRecipients | CopyTo | BlindCopyTo> = "emailaddress"}
Add the sender's manager as recipient	AddRecipients	First property: *AddedManagerAction* Second property: *Field*	Adds the sender's manager to the message as the specified recipient type (To, Cc, Bcc), or redirects the message to the sender's manager without notifying the sender or the recipient. This action only works if the sender's Manager attribute is defined in Active Directory. This parameter uses the syntax: @{AddManagerAsRecipientType = "<To | Cc | Bcc>"}
Prepend subject	PrependSubject	String	Adds the specified text to the beginning of the Subject field of the message. Consider using a space or a colon (:) as the last character of the specified text to differentiate it from the original subject text. To prevent the same string from being added to messages that already contain the text in the subject (for example, replies), add the "The subject contains words" (ExceptIfSubjectContainsWords) exception to the rule.
Apply HTML disclaimer	ApplyHtmlDisclaimer	First property: *Text* Second property: *Location* Third property: *Fallback action*	Applies the specified HTML disclaimer to the required location of the message. This parameter uses the syntax: @{ Text = " " ; Location = <Append | Prepend>; FallbackAction = <Wrap | Ignore | Reject> }

Microsoft Information Governance in Microsoft 365

2/18/2021 • 2 minutes to read • Edit Online

Microsoft 365 licensing guidance for security & compliance.

Use Microsoft Information Governance (sometimes abbreviated to MIG) capabilities to govern your data for compliance or regulatory requirements.

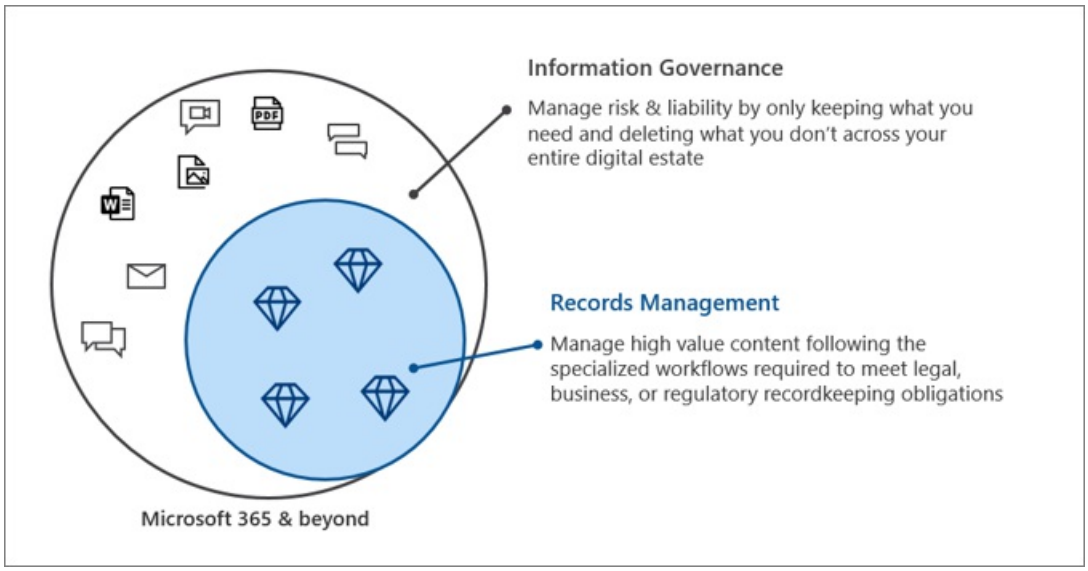

Looking to protect your data? See Microsoft Information Protection in Microsoft 365.

Information governance

To keep what you need and delete what you don't:

CAPABILITY	WHAT PROBLEMS DOES IT SOLVE?	GET STARTED
Retention policies and retention labels	Retain or delete content with policy management and a deletion workflow for email, documents, instant messages, and more Example scenario: Apply a retention label to content automatically	Get started with retention policies and retention labels
Import service	Bulk-import PST files to Exchange Online mailboxes to retain and search email messages for compliance or regulatory requirements	Use network upload to import your organization's PST files to Microsoft 365
Archive third-party data	Import, archive, and apply compliance solutions to third-party data from social media platforms, instant messaging platforms, and document collaboration platforms	Third-party connectors
Inactive mailboxes	Retain mailbox content after employees leave the organization	Create and manage inactive mailboxes

CAPABILITY	WHAT PROBLEMS DOES IT SOLVE?	GET STARTED

Records management

To manage high-value content for legal, business, or regulatory obligations:

CAPABILITY	WHAT PROBLEMS DOES IT SOLVE?	GET STARTED
Records management	A single solution for email and documents that incorporates retention schedules and requirements into a file plan that supports the full lifecycle of your content with records declaration, retention, and disposition Example scenario: Disposition of records	Get started with records management

Learn about retention policies and retention labels

2/18/2021 • 24 minutes to read • Edit Online

Microsoft 365 licensing guidance for security & compliance.

For most organizations, the volume and complexity of their data is increasing daily—email, documents, instant messages, and more. Effectively managing or governing this information is important because you need to:

- **Comply proactively with industry regulations and internal policies** that require you to retain content for a minimum period of time—for example, the Sarbanes-Oxley Act might require you to retain certain types of content for seven years.
- **Reduce your risk in the event of litigation or a security breach** by permanently deleting old content that you're no longer required to keep.
- **Help your organization to share knowledge effectively and be more agile** by ensuring that your users work only with content that's current and relevant to them.

Retention settings that you configure can help you achieve all these goals. Managing content commonly requires two actions:

- **Retaining** content so that it can't be permanently deleted before the end of the retention period.
- **Deleting** content permanently at the end of the retention period.

With these two retention actions, you can configure retention settings for the following outcomes:

- Retain-only: Retain content forever or for a specified period of time.
- Delete-only: Delete content after a specified period of time.
- Retain and then delete: Retain content for a specified period of time and then delete it.

These retention settings work with content in place that saves you the additional overheads of creating and configuring additional storage when you need to retain content for compliance reasons. In addition, you don't need to implement customized processes to copy and synchronize this data.

How retention settings work with content in place

When content has retention settings assigned to it, that content remains in its original location. People can continue to work with their documents or mail as if nothing's changed. But if they edit or delete content that's included in the retention policy, a copy of the content is automatically retained.

- For SharePoint and OneDrive sites: The copy is retained in the **Preservation Hold** library.
- For Exchange mailboxes: The copy is retained in the **Recoverable Items** folder.
- For Teams and Yammer messages: The copy is retained in a hidden folder named **SubstrateHolds** as a subfolder in the Exchange **Recoverable Items** folder.

> **NOTE**
> The Preservation Hold library consumes storage that isn't exempt from a site's storage quota. You might need to increase your storage when you use retention settings for SharePoint and Microsoft 365 groups.

These secure locations and the retained content are not visible to most people. In most cases, people do not even need to know that their content is subject to retention settings.

For more detailed information about how retention settings work for different workloads, see the following articles:

- Learn about retention for SharePoint and OneDrive
- Learn about retention for Microsoft Teams
- Learn about retention for Yammer
- Learn about retention for Exchange

Retention policies and retention labels

To assign your retention settings to content, use **retention policies** and **retention labels with label policies**. You can use just one of these methods, or combine them.

Use a retention policy to assign the same retention settings for content at a site or mailbox level, and use a retention label to assign retention settings at an item level (folder, document, email).

For example, if all documents in a SharePoint site should be retained for 5 years, it's more efficient to do this with a retention policy than apply the same retention label to all documents in that site. However, if some documents in that site should be retained for 5 years and others retained for 10 years, a retention policy wouldn't be able to do this. When you need to specify retention settings at the item level, use retention labels.

Unlike retention policies, retention settings from retention labels travel with the content if it's moved to a different location within your Microsoft 365 tenant. In addition, retention labels have the following capabilities that retention policies don't support:

- Options to start the retention period from when the content was labeled or based on an event, in addition to the age of the content or when it was last modified.
- Use trainable classifiers to identify content to label.
- Apply a default label for SharePoint documents.
- Support disposition reviewto review the content before it's permanently deleted.
- Mark the content as a record as part of the label settings, and always haveproof of dispositionwhen content is deleted at the end of its retention period.

Retention policies

Retention policies can be applied to the following locations:

- Exchange email
- SharePoint site
- OneDrive accounts
- Microsoft 365 Groups
- Skype for Business
- Exchange public folders
- Teams channel messages
- Teams chats
- Yammer community messages
- Yammer private messages

You can very efficiently apply a single policy to multiple locations, or to specific locations or users.

For the start of the retention period, you can choose when the content was created or, supported only for files

and the SharePoint, OneDrive, and Microsoft 365 Groups locations, when the content was last modified.

Items inherit the retention settings from their container specified in the retention policy. If they are then moved outside that container when the policy is configured to retain content, a copy of that item is retained in the workload's secured location. However, the retention settings don't travel with the content in its new location. If that's required, use retention labels instead of retention policies.

Retention labels

Use retention labels for different types of content that require different retention settings. For example:

- Tax forms that need to be retained for a minimum period of time.
- Press materials that need to be permanently deleted when they reach a specific age.
- Competitive research that needs to be retained for a specific period and then permanently deleted.
- Work visas that must be marked as a record so that they can't be edited or deleted.

In all these cases, retention labels let you apply retention settings for governance control at the item level (document or email).

With retention labels, you can:

- **Enable people in your organization to apply a retention label manually** to content in Outlook and Outlook on the web, OneDrive, SharePoint, and Microsoft 365 groups. Users often know best what type of content they're working with, so they can classify it and have the appropriate retention settings applied.
- **Apply retention labels to content automatically** if it matches specific conditions, such as when the content contains:
  - Specific types of sensitive information.
  - Specific keywords that match a query you create.
  - Pattern matches for a trainable classifier.
- **Start the retention period from when the content was labeled** for documents in SharePoint sites and OneDrive accounts, and to email items with the exception of calendar items. If you apply a retention label with this configuration to a calendar item, the retention period starts from when it is sent.
- **Start the retention period when an event occurs**, such as employees leave the organization, or contracts expire.
- **Apply a default retention label to a document library, folder, or document set** in SharePoint, so that all documents that are stored in that location inherit the default retention label.

Additionally, retention labels support records management for email and documents across Microsoft 365 apps and services. You can use a retention label to mark items as a record. When this happens and the content remains in Microsoft 365, the label places further restrictions on the content that might be needed for regulatory reasons. For more information, see Compare restrictions for what actions are allowed or blocked.

Retention labels, unlike sensitivity labels, do not persist if the content is moved outside Microsoft 365.

There is no limit to the number of retention labels that are supported for a tenant. However, 10,000 is the maximum number of policies that are supported for a tenant and these include the policies that apply the labels (retention label policies and auto-apply retention policies), as well as retention policies.

Classifying content without applying any actions

Although the main purpose of retention labels is to retain or delete content, you can also use retention labels without turning on any retention or other actions. In this case, you can use a retention label simply as a text label, without enforcing any actions.

For example, you can create and apply a retention label named "Review later" with no actions, and then use that label to find that content later.

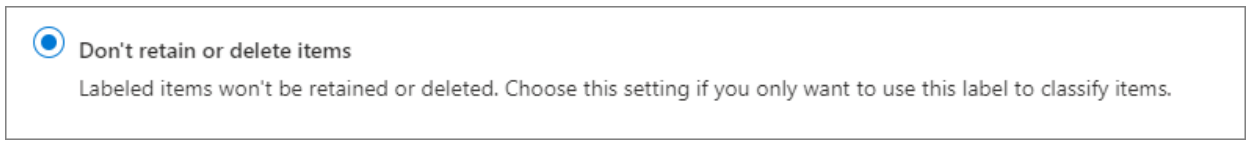

Using a retention label as a condition in a DLP policy

You can specify a retention label as a condition in a data loss prevention (DLP) policy for documents in SharePoint. For example, configure a DLP policy to prevent documents from being shared outside the organization if they have a specified retention label applied to it.

For more information, see Using a retention label as a condition in a DLP policy.

Retention labels and policies that apply them

When you publish retention labels, they're included in a **retention label policy** that makes them available for admins and users to apply to content. As the following diagram shows:

1. A single retention label can be included in multiple retention label policies.
2. Retention label policies specify the locations to publish the retention labels. The same location can be included in multiple retention label policies.

You can also create one or more **auto-apply retention label policies**, each with a single retention label. With this policy, a retention label is automatically applied when conditions that you specify in the policy are met.

Retention label policies and locations

Different types of retention labels can be published to different locations, depending on what the retention label does.

IF THE RETENTION LABEL IS...	THEN THE LABEL POLICY CAN BE APPLIED TO...
Published to admins and end users	Exchange, SharePoint, OneDrive, Microsoft 365 Groups
Auto-applied based on sensitive information types or trainable classifiers	Exchange (all mailboxes only), SharePoint, OneDrive
Auto-applied based on a query	Exchange, SharePoint, OneDrive, Microsoft 365 Groups

In Exchange, retention labels that you auto-apply are applied only to messages newly sent (data in transit), not

to all items currently in the mailbox (data at rest). Also, auto-apply retention labels for sensitive information types and trainable classifiers apply to all mailboxes; you can't select specific mailboxes.

Exchange public folders, Skype, Teams and Yammer messages do not support retention labels. To retain and delete contain from these locations, use retention policies instead.

Only one retention label at a time

An email or document can have only a single retention label applied to it at a time. A retention label can be applied manually by an end user or admin, or automatically by using any of the following methods:

- Auto-apply label policy
- Document understanding model for SharePoint Syntex
- Default label for SharePoint or Outlook
- Outlook rules

For standard retention labels (they don't mark items as a record or regulatory record):

- Admins and end users can manually change or remove an existing retention label that's applied on content.
- When content already has a retention label applied, the existing label won't be automatically removed or replaced by another retention label with one possible exception: The existing label was applied as a default label.

  For more information about the label behavior when it's applied by using a default label:

  - Default label for SharePoint: Label behavior when you use a default label for SharePoint
  - Default label for Outlook: Applying a default retention label to an Outlook folder
- If there are multiple auto-apply label policies that could apply a retention label, and content meets the conditions of multiple policies, the retention label for the oldest auto-apply label policy (by date created) is applied.

When retention labels mark items as a record or a regulatory record, these labels are never automatically changed. Only admins for the container can manually change or remove retention labels that mark items as a record, but not regulatory records. For more information, see Compare restrictions for what actions are allowed or blocked.

Monitoring retention labels

From the Microsoft 365 compliance center, use **Data classification** > **Overview** to monitor how your retention labels are being used in your tenant, and identify where your labeled items are located. For more information, including important prerequisites, see Know your data - data classification overview.

You can then drill down into details by using content explorer and activity explorer.

> **TIP**
> Consider using some of the other data classification insights, such as trainable classifiers and sensitive info types, to help you identify content that you might need to retain or delete, or manage as records.

The Office 365 Security & Compliance Center has the equivalent overview information for retention labels from **Information governance** > **Dashboard**, and more detailed information from **Information governance** > **Label activity explorer**. For more information about monitoring retention labels from this older admin center, see the following documentation:

- View the data governance reports
- Get started with data classification.

- View label activity for documents

Using Content Search to find all content with a specific retention label

After retention labels are applied to content, either by users or auto-applied, you can use content search to find all items that have a specific retention label applied.

When you create a content search, choose the **Retention label** condition, and then enter the complete retention label name or part of the label name and use a wildcard. For more information, see Keyword queries and search conditions for Content Search.

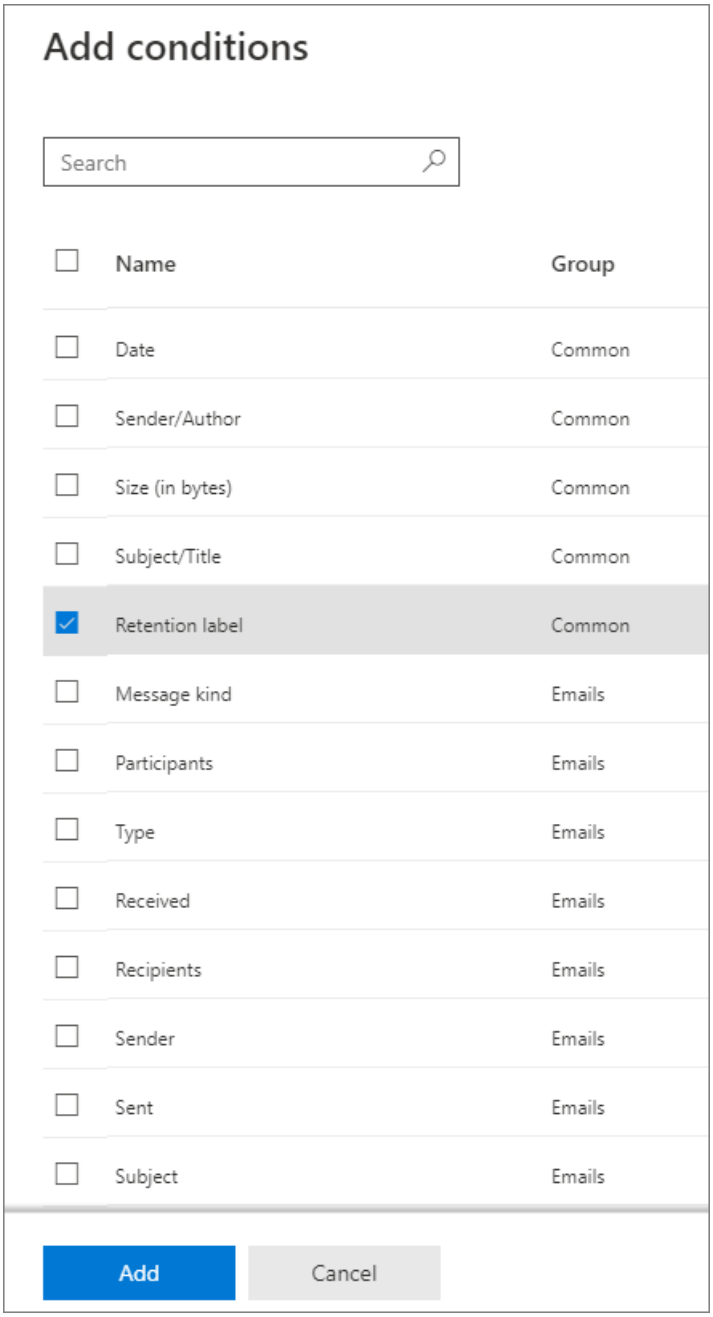

Compare capabilities for retention policies and retention labels

Use the following table to help you identify whether to use a retention policy or retention label, based on capabilities.

CAPABILITY	RETENTION POLICY	RETENTION LABEL
Retention settings that can retain and then delete, retain-only, or delete-only	Yes	Yes

CAPABILITY	RETENTION POLICY	RETENTION LABEL
Workloads supported: - Exchange - SharePoint - OneDrive - Microsoft 365 groups - Skype for Business - Teams - Yammer	 Yes Yes Yes Yes Yes Yes Yes	 Yes, except public folders Yes Yes Yes No No No
Retention applied automatically	Yes	Yes
Retention applied based on conditions - sensitive info types, KQL queries and keywords, trainable classifiers	No	Yes
Retention applied manually	No	Yes
UI presence for end users	No	Yes
Persists if the content is moved	No	Yes, within your Microsoft 365 tenant
Declare item as a record	No	Yes
Start the retention period when labeled or based on an event	No	Yes
Disposition review	No	Yes
Proof of disposition for up to 7 years	No	Yes, when item is declared a record
Audit admin activities	Yes	Yes
Identify items subject to retention: - Content Search - Data classification page, content explorer, activity explorer	 No No	 Yes Yes

Note that you can use both retention policies and retention labels as complementary retention methods. For example:

1. You create and configure a retention policy that automatically deletes content five years after it's last modified, and apply the policy to all OneDrive accounts.
2. You create and configure a retention label that keeps content forever and add this to a label policy that you publish to all OneDrive accounts. You explain to users how to manually apply this label to specific documents that should be excluded from automatic deletion if not modified after five years.

For more information about how retention policies and retention labels work together and how to determine their combined outcome, see the next section that explains the principles of retention and what takes precedence.

The principles of retention, or what takes precedence?

Unlike retention labels, you can apply more than one retention policy to the same content. Each retention policy

can result in a retain action and a delete action. Additionally, that item could also be subject to these actions from a retention label.

In this scenario, when items can be subject to multiple retention settings that could conflict with one another, what takes precedence to determine the outcome?

The outcome isn't which single retention policy or single retention label wins, but how long an item is retained (if applicable) and when an item is deleted (if applicable). These two actions are calculated independently from each other, from all the retention settings applied to an item.

For example, an item might be subject to one retention policy that is configured for a delete-only action, and another retention policy that is configured to retain and then delete. Consequently, this item has just one retain action but two delete actions. The retention and deletion actions could be in conflict with one another and the two deletion actions might have a conflicting date. To work out the outcome, you must apply the principles of retention.

At a high level, you can be assured that retention always takes precedence over deletion, and the longest retention period wins. These two simple rules always decide how long an item will be retained.

There are a few more factors that determine when an item will be deleted, which include the delete action from a retention label always takes precedence over the delete action from a retention policy.

Use the following flow to understand the retention and deletion outcomes for a single item, where each level acts as a tie-breaker for conflicts, from top to bottom. If the outcome is determined by the first level because there are no further conflicts, there's no need to progress to the next level, and so on.

> **IMPORTANT**
> If you are using retention labels: Before using this flow to determine the outcome of multiple retention settings on the same item, make sure you know which retention label is applied.

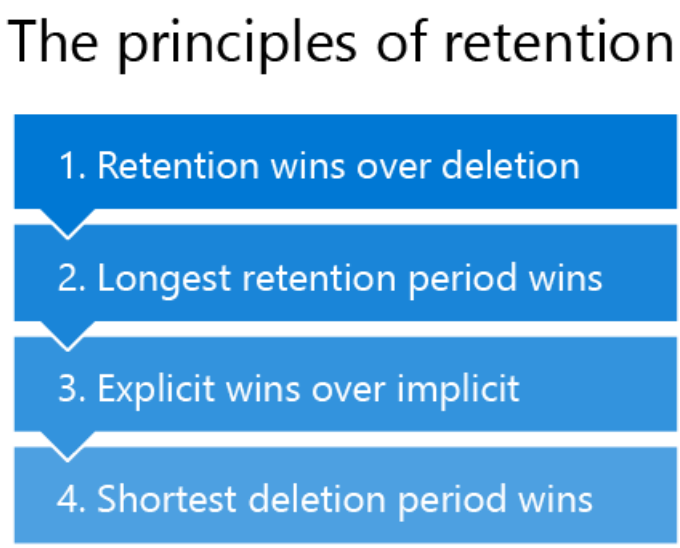

Explanation for the four different levels:

1. **Retention wins over deletion.** Content won't be permanently deleted when it also has retention settings to retain it.

   Example: An email message is subject to a retention policy for Exchange that is configured to delete items after three years and it also has a retention label applied that is configured to retain items for five years.

   The email message is retained for five years because this retention action takes precedence over deletion.

The email message is deleted at the end of the five years because of the deferred delete action.

2. **The longest retention period wins.** If content is subject to multiple retention settings that retain content for different periods of time, the content will be retained until the end of the longest retention period.

   Example: Documents in the Marketing SharePoint site are subject to two retention policies. The first retention policy is configured for all SharePoint sites to retain items for five years. The second retention policy is configured for specific SharePoint sites to retain items for ten years.

   Documents in this Marketing SharePoint site are retained for ten years because that's the longest retention period.

3. **Explicit wins over implicit.** Applicable to determine when items will be deleted:

   a. A retention label (however it was applied) provides explicit retention in comparison with retention policies, because the retention settings are applied to an individual item rather than implicitly assigned from a container. This means that a delete action from a retention label always takes precedence over a delete action from any retention policy.

      Example: A document is subject to two retention policies that have a delete action of five years and ten years respectively, and also a retention label that has a delete action of seven years.

      The document is deleted after seven years because the delete action from the retention label takes precedence.

   b. When you have retention policies only: If a retention policy for a location is scoped to use an include configuration (such as specific users for Exchange email) that retention policy takes precedence over unscoped retention policies for the same location.

      An unscoped retention policy is where a location is selected without specifying specific instances. For example, **Exchange email** and the default setting of **All recipients** is an unscoped retention policy. Or, **SharePoint sites** and the default setting of **All sites**. When retention policies are scoped, they have equal precedence at this level.

      Example 1: An email message is subject to two retention policies. The first retention policy is unscoped and deletes items after ten years. The second retention policy is scoped to specific mailboxes and deletes items after five years.

      The email message is deleted after five years because the deletion action from the scoped retention policy takes precedence over the unscoped retention policy.

      Example 2: A document in a user's OneDrive account is subject to two retention policies. The first retention policy is scoped to include this user's OneDrive account and has a delete action after 10 years. The second retention policy is scoped to include this user's OneDrive account and has a delete action after seven years.

      When this document will be deleted can't be determined at this level because both retention policies are scoped.

4. **The shortest deletion period wins.** Applicable to determine when items will be deleted from retention policies and the outcome couldn't be resolved from the previous level: Content is deleted at the end of the shortest retention period.

   Example: A document in a user's OneDrive account is subject to two retention policies. The first retention policy is scoped to include this user's OneDrive account and has a delete action after 10 years. The second retention policy is scoped to include this user's OneDrive account and has a delete action after seven years.

This document will be deleted after seven years because that's the shortest retention period for these two scoped retention policies.

Note that items subject to eDiscovery hold also fall under the first principle of retention; they cannot be deleted by any retention policy or retention label. When that hold is released, the principles of retention continue to apply to them. For example, they could then be subject to an unexpired retention period or a deferred delete action.

More complex examples that combine retain and delete actions:

1. An item has the following retention settings applied to it:

   - A retention policy for delete-only after five years
   - A retention policy that retains for three years and then deletes
   - A retention label that retains-only for seven years

   Outcome: The item is retained for seven years because retention takes precedence over deletion and seven years is the longest retention period. At the end of this retention period, the item is deleted because of the delete action from the retention policies that was deferred while the item was retained.

   Although the two retention policies have different dates for the delete actions, the earliest the item can be deleted is at the end of the longest retention period, which is longer than both deletion dates. In this example, there is no conflict to resolve for the deletion dates so all conflicts are resolved by the second level.

2. An item has the following retention settings applied to it:

   - An unscoped retention policy that deletes-only after ten years
   - A scoped retention policy that retains for five years and then deletes
   - A retention label that retains for three years and then deletes

   Outcome: The item is retained for five years because that's the longest retention period. At the end of that retention period, the item is deleted because of the delete action of three years from the retention label that was deferred while the item was retained. Deletion from retention labels takes precedence over deletion from all retention policies. In this example, all conflicts are resolved by the third level.

Use Preservation Lock to restrict changes to policies

Some organizations might need to comply with rules defined by regulatory bodies such as the Securities and Exchange Commission (SEC) Rule 17a-4, which requires that after a policy for retention is turned on, it cannot be turned off or made less restrictive.

Preservation Lock ensures your organization can meet such regulatory requirements because it locks a retention policy or retention label policy so that no one—including an administrator—can turn off the policy, delete the policy, or make it less restrictive.

You apply Preservation Lock after the retention policy or retention label policy is created. For more information and instructions, see Use Preservation Lock to restrict changes to retention policies and retention label policies.

Releasing a policy for retention

Providing your policies for retention don't have a Preservation Lock, you can delete your policies at any time, which effectively turns off the previously applied retention settings. You can also keep the policy but change the location status to off.

When you do either of these actions, any SharePoint or OneDrive content that's being retained in the Preservation Hold library is not immediately and permanently deleted. Instead, to help prevent inadvertent data loss, there is a 30-day grace period, during which content expiration for that policy does not happen in the

Preservation Hold library, so that you can restore any content from there, if needed. Additionally, you can't manually delete this content during the grace period.

You can change the location status back to on during the grace period, and no content will be deleted for that policy.

This 30-day grace period in SharePoint and OneDrive corresponds to the 30-day delay hold in Exchange. For more information, see Managing mailboxes on delay hold.

Auditing retention configuration

Administrator actions for retention policies and retention labels are saved to the audit log when auditing is enabled. For example, an audit event is created when a retention policy or label is created, configured, or deleted. For the full list, see Retention policy and retention label activities.

PowerShell cmdlets for retention policies and retention labels

To use the retention cmdlets, you must first connect to the Office 365 Security & Compliance Center PowerShell. Then, use any of the following cmdlets:

- Get-ComplianceTag
- New-ComplianceTag
- Remove-ComplianceTag
- Set-ComplianceTag
- Enable-ComplianceTagStorage
- Get-ComplianceTagStorage
- Get-RetentionCompliancePolicy
- New-RetentionCompliancePolicy
- Remove-RetentionCompliancePolicy
- Set-RetentionCompliancePolicy
- Get-RetentionComplianceRule
- New-RetentionComplianceRule
- Remove-RetentionComplianceRule
- Set-RetentionComplianceRule

When to use retention policies and retention labels or eDiscovery holds

Although retention settings and holds that you create with an eDiscovery case can both prevent data from being permanently deleted, they are designed for different scenarios. To help you understand the differences and decide which to use, use the following guidance:

- Retention settings that you specify in retention policies and retention labels are designed for a long-term information governance strategy to retain or delete data for compliance requirements. The scope is usually broad with the main focus being the location and content rather than individual users. The start and end of the retention period is configurable, with the option to automatically delete content without additional administrator intervention.

- Holds for eDiscovery (either Core eDiscovery or Advanced eDiscovery cases) are designed for a limited duration to preserve data for a legal investigation. The scope is specific with the focus being content owned by identified users. The start and end of the preservation period isn't configurable but dependent on individual administrator actions, without an option to automatically delete content when the hold is released.

Summary to compare retention with holds:

CONSIDERATION	RETENTION	EDISCOVERY HOLDS
Business need:	Compliance	Legal
Time scope:	Long-term	Short-term
Focus:	Broad, content-based	Specific, user-based
Start and end date configurable:	Yes	No
Content deletion:	Yes (optional)	No
Administrative overheads:	Low	High

If content is subject to both retention settings and an eDiscovery hold, preserving content for the eDiscovery hold always takes precedence. In this way, the principles of retention expand to eDiscovery holds because they preserve data until an administrator manually releases the hold. However, despite this precedence, don't use eDiscovery holds for long-term information governance. If you are concerned about automatic deletion of data, you can configure retention settings to retain items forever, or use disposition review with retention labels.

If you are using older eDiscovery tools to preserve data, see the following resources:

- Exchange:
  - In-Place Hold and Litigation Hold
  - How to identify the type of hold placed on an Exchange Online mailbox
- SharePoint and OneDrive:
  - Add content to a case and place sources on hold in the eDiscovery Center
- Retirement of legacy eDiscovery tools

Use retention policies and retention labels instead of older features

If you need to proactively retain or delete content in Microsoft 365 for information governance, we recommend that you use retention policies and retention labels instead of the following older features.

If you currently use these older features, they will continue to work side-by-side with retention policies and retention labels. However, we recommend that going forward, you use retention policies and retention labels instead. They provide you with a single mechanism to centrally manage both retention and deletion of content across Microsoft 365.

Older features from Exchange Online:

- Retention tags and retention policies, also known as messaging records management (MRM) (deletion only)

Older features from SharePoint and OneDrive:

- Document deletion policies (deletion only)

- Configuring in place records management (retention only)
- Use policies for site closure and deletion (deletion only)
- Information management policies (deletion only)

If you have configured SharePoint sites for content type policies or information management policies to retain content for a list or library, those policies are ignored while a retention policy is in effect.

Related information

- SharePoint Online Limits
- Limits and specifications for Microsoft Teams
- Resources to help you meet regulatory requirements for information governance and records management

Configuration guidance

If you are ready to create retention policies, see Create and configure retention policies.

To create and apply retention labels:

- Create retention labels and apply them in apps
- Apply a retention label to content automatically

Learn about retention for SharePoint and OneDrive

2/18/2021 • 9 minutes to read • Edit Online

> *Microsoft 365 licensing guidance for security & compliance.*

The information in this article supplements Learn about retention because it has information that's specific to SharePoint and OneDrive.

For other workloads, see:

- Learn about retention for Microsoft Teams
- Learn about retention for Yammer
- Learn about retention for Exchange

What's included for retention and deletion

All files stored in SharePoint or OneDrive sites can be retained by applying a retention policy or retention label.

The following files can be deleted:

- When you use a retention policy: All files in document libraries, which include any automatically created SharePoint document libraries, such as **Site Assets**.
- When you use retention labels: All files in all document libraries, and all files at the root level that aren't in a folder.

> **TIP**
> When you use a query with an auto-apply policy for a retention label, you can exclude specific document libraries by using the following entry: `NOT(DocumentLink:"<URL to document library>")`

List items are not supported by retention policies but are supported by retention labels with the exception of items in system lists. These are hidden lists used by SharePoint to manage the system and include the master page catalog, solution catalog, and data sources.

Retention settings from both retention policies and retention labels do not apply to organizing structures that include libraries, lists, and folders.

For retention policies and auto-apply label policies: SharePoint sites must be indexed for the retention settings to be applied. However, if items in SharePoint document libraries are configured to not appear in search results, this configuration doesn't exclude files from the retention settings.

How retention works for SharePoint and OneDrive

To store content that needs to be retained, SharePoint and OneDrive create a Preservation Hold library if one doesn't exist. You can view this library on the **Site contents** page in the top-level site of the site collection. Most users can't view the Preservation Hold library because it's visible only to site collection administrators.

Items in SharePoint that have a standard retention label (doesn't declare the item to be a record) don't need the Preservation Hold library because these items remain in their original location. SharePoint prevents users from deleting items when the applied retention label is configured to retain the content, and SharePoint versioning preserves older versions when items are edited. But for other scenarios, the Preservation Hold library is used

when items must be retained:

- Items in OneDrive that have standard retention labels
- Items in SharePoint or OneDrive that have retention labels that declares them a record, and the item is unlocked for editing
- Items that are subject to retention policies

To retain this content when a user attempts to change or delete it, a check is made whether the content's been changed since the retention settings were applied. If this is the first change since the retention settings were applied, the content is copied to the Preservation Hold library, which allows the person to change or delete the original content. Any content in a site collection can be copied to the Preservation Hold library, independently from retention settings.

A timer job periodically cleans up the Preservation Hold library. This job compares all content in the Preservation Hold library to all queries used by the retention settings for that content. Content that is older than their configured retention period is deleted from the Preservation Hold library, and the original location if it is still there. This timer job runs every seven days, which means that it can take up to seven days for content to be deleted.

This behavior applies to content that exists when the retention settings were applied. In addition, for retention policies, any new content that's created or added to the site collection after it was included in the policy will be retained after deletion. However, new content isn't copied to the Preservation Hold library the first time it's edited, only when it's deleted. To retain all versions of a file, you must turn on versioning.

Users see an error message if they try to delete a library, list, folder, or site that's subject to retention. They can delete a folder if they first move or delete any files in the folder that are subject to retention.

> **NOTE**
> Because the Preservation Hold library is created only when it's needed, and not when you apply a retention policy or retention label, to see this working, you must first edit or delete an item that's subject to retention. Then browse to the Preservation Hold library to view the retained copy.

After retention settings are assigned to content in a OneDrive account or SharePoint site, the paths the content takes depend on whether the retention settings are to retain and delete, to retain only, or delete only.

When the retention settings are to retain and delete:

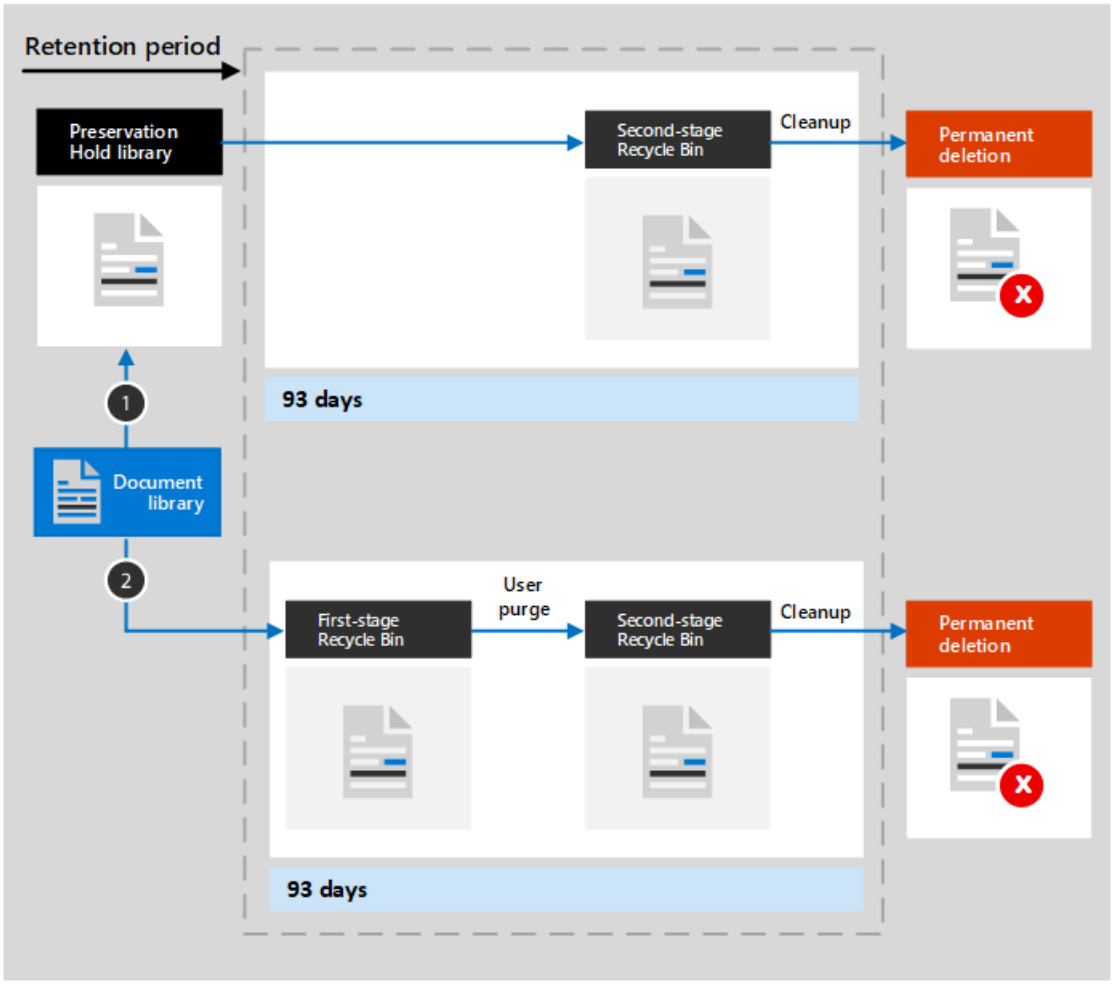

1. **If the content is modified or deleted** during the retention period, a copy of the original content as it existed when the retention settings were assigned is created in the Preservation Hold library. There, the timer job identifies items whose retention period has expired. Those items are moved to the second-stage Recycle Bin, where they're permanently deleted at the end of 93 days. The second-stage Recycle Bin is not visible to end users (only the first-stage Recycle Bin is), but site collection admins can view and restore content from there.

   > **NOTE**
   > To help prevent inadvertent data loss, we no longer permanently delete content from the Preservation Hold library. Instead, we permanently delete content only from the Recycle Bin, so all content from the Preservation Hold library now goes through the second-stage Recycle Bin.

2. **If the content is not modified or deleted** during the retention period, the timer job moves this content to the first-stage Recycle Bin at the end of the retention period. If a user deletes the content from there or empties this Recycle Bin (also known as purging), the document is moved to the second-stage Recycle Bin. A 93-day retention period spans both the first- and second-stage recycle bins. At the end of 93 days, the document is permanently deleted from wherever it resides, in either the first-stage or second-stage Recycle Bin. The Recycle Bin is not indexed and therefore unavailable for searching. As a result, an eDiscovery search can't find any Recycle Bin content on which to place a hold.

When the retention settings are retain-only, or delete-only, the contents paths are variations of retain and delete:

Content paths for retain-only retention settings

1. **If the content is modified or deleted** during the retention period: A copy of the original document is created in the Preservation Hold library and retained until the end of the retention period, when the copy in the Preservation Hold library is moved to the second-stage Recycle Bin and is permanently deleted

after 93 days.

2. **If the content is not modified or deleted** during the retention period: Nothing happens before and after the retention period; the document remains in its original location.

Content paths for delete-only retention settings

1. **If the content is deleted** during the configured period: The document is moved to first-stage Recycle Bin. If a user deletes the document from there or empties this Recycle Bin, the document is moved to the second-stage Recycle Bin. A 93-day retention period spans both the first-stage and second-stage recycle bins. At the end of 93 days, the document is permanently deleted from wherever it resides, in either the first-stage or second-stage Recycle Bin. If the content is modified during the configured period, it follows the same deletion path after the configured period.

2. **If the content is not deleted** during the configured period: At the end of the configured period in the retention policy, the document is moved to the first-stage Recycle Bin. If a user deletes the document from there or empties this Recycle Bin (also known as purging), the document is moved to the second-stage Recycle Bin. A 93-day retention period spans both the first-stage and second-stage recycle bins. At the end of 93 days, the document is permanently deleted from wherever it resides, in either the first-stage or second-stage Recycle Bin. The Recycle Bin is not indexed and therefore unavailable for searching. As a result, an eDiscovery search can't find any Recycle Bin content on which to place a hold.

How retention works for OneNote content

When you apply a retention policy to a location that includes OneNote content, behind the scenes, the different OneNote sections are individual files. This means that each section will be individually retained and deleted, according to the retention settings you specify.

How retention works with document versions

Versioning is a feature of all document lists and libraries in SharePoint and OneDrive. By default, versioning retains a minimum of 500 major versions, although you can increase this limit. For more information, see Enable and configure versioning for a list or library and How versioning works in lists and libraries.

When a document with versions is subject to retention settings to retain that content, versions that get copied to the Preservation Hold library exist as a separate item. If the retention settings are configured to delete at the end of the retention period:

- If the retention period is based on when the content was created, each version has the same expiration date as the original document. The original document and its versions all expire at the same time.
- If the retention period is based on when the content was last modified, each version has its own expiration date based on when the original document was modified to create that version. The original document and its versions expire independently of each other.

> **NOTE**
> The retained versions of these SharePoint and OneDrive documents are not searchable by eDiscovery tools.

When the retention action is to delete the document, all versions not in the Preservation Hold library are deleted at the same time according to the current version.

For items that are subject to a retention policy (or an eDiscovery hold), the versioning limits for the document library are ignored until the retention period of the document is reached (or the eDiscovery hold is released). In this scenario, old versions are not automatically purged and users are prevented from deleting versions.

That's not the case for retention labels when the content isn't subject to a retention policy (or an eDiscovery hold). Instead, the versioning limits are honored so that older versions are automatically deleted to accommodate new versions, but users are still prevented from deleting versions.

When a user leaves the organization

SharePoint:

When a user leaves your organization, any content created by that user is not affected because SharePoint is considered a collaborative environment, unlike a user's mailbox or OneDrive account.

OneDrive:

If a user leaves your organization, any files that are subject to a retention policy or has a retention label will remain for the duration of the policy or label. During that time period, all sharing access continues to work. When the retention period expires, content moves into the Site Collection Recycle Bin and is not accessible to anyone except the admin. If a document is marked by a retention label as a record, the document will not be deleted until the retention period is over, after which time the content is permanently deleted.

Configuration guidance

If you're new to configuring retention in Microsoft 365, see Get started with retention policies and retention labels.

If you're ready to configure a retention policy or retention label for Exchange, see the following instructions:

- Create and configure retention policies
- Create retention labels and apply them in apps
- Apply a retention label to content automatically

Learn about retention for Microsoft Teams

2/18/2021 • 7 minutes to read • Edit Online

Microsoft 365 licensing guidance for security & compliance.

The information in this article supplements Learn about retention because it has information that's specific to Microsoft Teams messages.

For other workloads, see:

- Learn about retention for SharePoint and OneDrive
- Learn about retention for Yammer
- Learn about retention for Exchange

What's included for retention and deletion

The following Teams items can be retained and deleted by using retention policies for Teams: Chat messages and channel messages, including embedded images, tables, hypertext links and links to other Teams messages and files, and card content. Chat messages include all the names of the people in the chat, and channel messages include the team name and the message title (if supplied).

> **NOTE**
> Including card content is a recent addition and currently rolling out to tenants. For more information, see Microsoft 365 compliance capabilities for Adaptive Card content through apps in Teams now available.

Teams messages in private channels are currently not supported for retention policies. Code snippets, recorded voice memos from the Teams mobile client, and reactions from others in the form of emoticons are not included when you use retention policies for Teams.

Emails and files that you use with Teams aren't included in retention policies for Teams. These items have their own retention policies.

How retention works with Microsoft Teams

You can use a retention policy to retain and delete data from chats and channel messages in Teams. Behind the scenes, Exchange mailboxes are used to store these messages. Data from Teams chats is stored in a hidden folder in the mailbox of each user included in the chat, and a similar hidden folder in a group mailbox is used for Teams channel messages.

These mailboxes are, listed by their RecipientTypeDetails attribute:

- **UserMailbox**: These mailboxes store messages for Teams users who have an Exchange Online mailbox.
- **MailUser**: These mailboxes store messages for Teams users who have a mailbox for an on-premises Exchange server and not Exchange Online.
- **GroupMailbox**: These mailboxes store messages for Teams channels.

Other mailbox types, such as RoomMailbox that is used for Teams conference rooms, are not supported for Teams retention policies.

It's important to understand that Teams uses an Azure-powered chat service that also stores this data, and by

default this service stores the data indefinitely. For this reason, if you need to delete Teams messages for compliance reasons, we recommend that you use retention policies for Teams that can permanently delete this data from both the Exchange mailboxes and the underlying Azure-powered chat service. For more information about the underlying architecture, see Security and compliance in Microsoft Teams and specifically, the Information Protection Architecture section.

Although Teams chats and channel messages are stored in mailboxes, this Teams data is included only by a retention policy that's configured for the **Teams channel messages** and **Teams chats** locations. Teams chats and channel messages are not affected by retention policies that are configured for Exchange user or group mailboxes.

> **NOTE**
>
> If a user is included in an active retention policy that retains Teams data and you a delete a mailbox of a user who is included in this policy, to retain the Teams data, the mailbox is converted into an inactive mailbox. If you don't need to retain this Teams data for the user, exclude the user account from the retention policy before you delete their mailbox.

After a retention policy is configured for chat and channel messages, a timer job from the Exchange service periodically evaluates items in the hidden folder where these Teams messages are stored. The timer job takes up to seven days to run. When these items have expired their retention period, they are moved to the SubstrateHolds folder—another hidden folder that's in every user or group mailbox to store "soft-deleted" items before they are permanently deleted.

After a retention policy is configured for chat and channel messages, the paths the content takes depend on whether the retention policy is to retain and then delete, to retain only, or delete only.

When the retention policy is to retain and then delete:

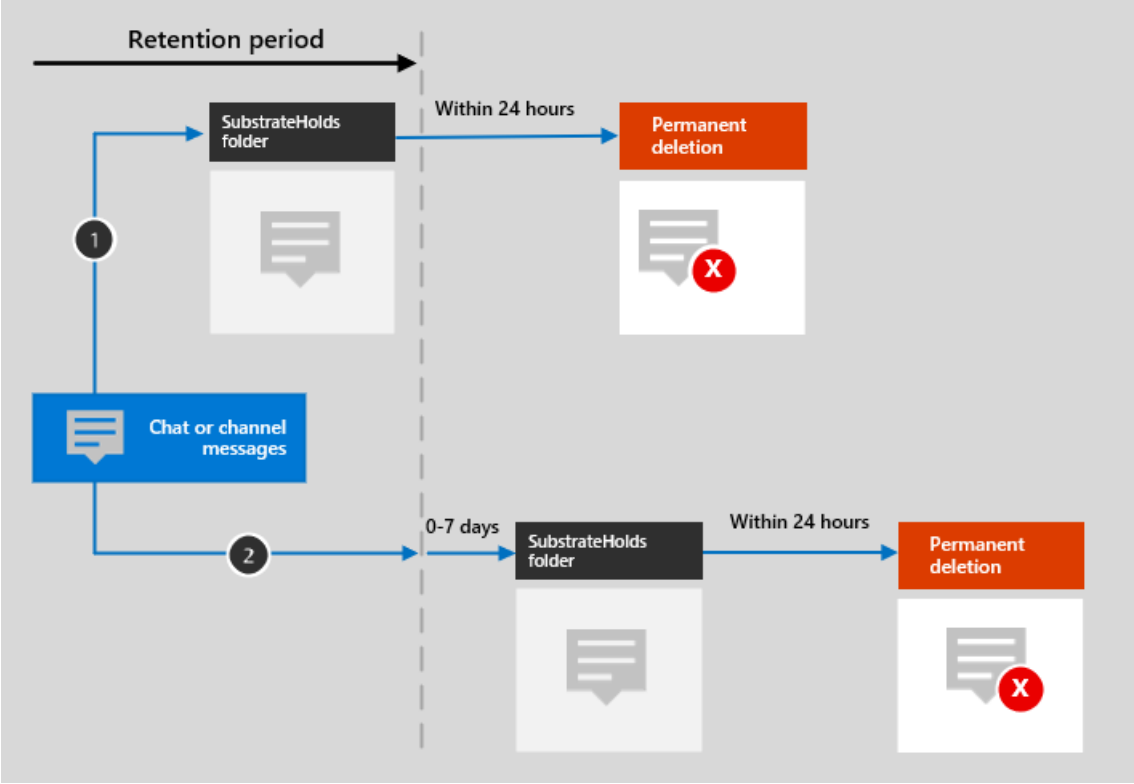

For the two paths in the diagram:

1. **If a chat or channel message is edited or deleted** by the user during the retention period, the original message is copied (if edited) or moved (if deleted) to the SubstrateHolds folder within 21 days. The message is stored there until the retention period expires and then the message is permanently deleted within 24 hours.

2. **If a chat or channel message is not deleted** and for current messages after editing, the message is moved to the SubstrateHolds folder after the retention period expires. This action takes up to 7 days from the expiry date. When the message is in the SubstrateHolds folder, it is then permanently deleted within 24 hours.

> **NOTE**
> Messages in the SubstrateHolds folder are searchable by eDiscovery tools. Until messages are permanently deleted from this SubstrateHolds folder, they remain searchable by eDiscovery tools.

When the retention policy is retain-only, or delete-only, the content's paths are variations of retain and delete.

Content paths for retain-only retention policy

1. **If a chat or channel message is edited or deleted**: A copy of the original message is created in the SubstrateHolds folder within 21 days, and retained there until the retention period expires. Then the message is permanently deleted from the SubstrateHolds folder within 24 hours.
2. **If the item is not modified or deleted** and for current messages after editing during the retention period: Nothing happens before and after the retention period; the message remains in its original location.

Content paths for delete-only retention policy

1. **If the message is not deleted** during the retention period: At the end of the retention period, the message is moved to the SubstrateHolds folder. This action takes up to seven days from the expiry date. Then the message is permanently deleted from the SubstrateHolds folder within 24 hours.
2. **If the item is deleted by the user** during the period, the item is moved to the SubstrateHolds folder within 21 days where it is then permanently deleted within 24 hours.

Skype for Business and Teams interop chats

When a Skype for Business chat comes into Teams, it becomes a message in a Teams chat thread and is ingested into the appropriate mailbox. Teams retention policies will apply to these messages from the Teams thread.

However, if conversation history is turned on for Skype for Business and from the Skype for Business client side that history is being saved into a mailbox, that chat data isn't handled by a Teams retention policy. For this content, use a retention policy that's configured for Skype for Business.

Meetings and external users

Channel meeting messages are stored the same way as channel messages, so for this data, select the **Teams channel messages** location when you configure your retention policy.

Impromptu and scheduled meeting messages are stored in the same way as group chat messages, so for this data, select the **Teams chats** location when you configure your retention policy.

When external users are included in a meeting that your organization hosts:

- If an external user joins by using a guest account in your tenant, this user has a shadow mailbox that can be subject to your organization's retention policy for Teams. Any messages from the meeting are stored in both your users' mailbox and the shadow mailbox.
- If an external user joins by using an account from another Microsoft 365 organization, your retention policies can't delete messages for this user because they are stored in that user's mailbox in another tenant. For the same meeting however, your retention policies can delete messages for your users.

When a user leaves the organization

If a user who has a mailbox in Exchange Online leaves your organization and their Microsoft 365 account is deleted, their chat messages that are subject to retention are stored in an inactive mailbox. The chat messages remain subject to any retention policy that was placed on the user before their mailbox was made inactive, and the contents are available to an eDiscovery search. For more information, see Inactive mailboxes in Exchange Online.

If the user stored any files in Teams, see the equivalent section for SharePoint and OneDrive.

Limitations

We're continuously working on optimizing retention functionality in Teams. In the meantime, here are a few limitations to be aware of when you use retention for Teams channel messages and chats:

- **Incorrect display issue in Outlook**. If you create retention policies for Skype or Teams locations, one of those policies is shown as the default folder policy when a user views the properties of a mailbox folder in the Outlook desktop client. This is an incorrect display issue in Outlook and a known issue. What should be displayed as the default folder policy is the mailbox retention policy that's applied to the folder. The Skype or Teams retention policy is not applied to the user's mailbox.
- **Configuration issues**:
  - When you select **Choose teams** for the **Teams channel messages** location, you might see Microsoft 365 groups that aren't also teams. Don't select these groups.
  - When you select **Choose users** for the **Teams chats** location, you might see guests and non-mailbox users. Retention policies aren't designed for these users, so don't select them.

Configuration guidance

If you're new to configuring retention in Microsoft 365, see Get started with retention policies and retention labels.

If you're ready to configure a retention policy for Teams, see Create and configure retention policies.

Learn about retention for Yammer

2/18/2021 • 4 minutes to read • Edit Online

Microsoft 365 licensing guidance for security & compliance.

NOTE

This feature is in preview and not yet available for all customers.

The information in this article supplements Learn about retention because it has information that's specific to Yammer.

For other workloads, see:

- Learn about retention for SharePoint and OneDrive
- Learn about retention for Microsoft Teams
- Learn about retention for Exchange

What's included for retention and deletion

The following Yammer items can be retained and deleted by using retention policies for Yammer: Community messages and private messages.

Reactions from others in the form of emoticons are not included in these messages.

How retention works with Yammer

You can use a retention policy to retain and delete community messages and private messages in Yammer. Private messages are stored in a hidden folder in the mailbox of each user included in the message, and community messages are stored in a similar hidden folder in the group mailbox for the community.

Yammer messages are not affected by retention policies that are configured for user or group mailboxes. Even though Yammer messages are stored in Exchange, this Yammer data is included only by a retention policy that's configured for the **Yammer community messages** and **Yammer private messages** locations.

NOTE

If a user is included in an active retention policy that retains Yammer data and you a delete a mailbox of a user who is included in this policy, to retain the Yammer data, the mailbox is converted into an inactive mailbox. If you don't need to retain this Yammer data for the user, exclude the user account from the retention policy before you delete their mailbox.

After a retention policy is configured for Yammer messages, a timer job from the Exchange service periodically evaluates items in the hidden folder where these Yammer messages are stored. The timer job takes up to seven days to run. When these items have expired their retention period, they are moved to the SubstrateHolds folder —a hidden folder that's in every user or group mailbox to store "soft-deleted" items before they are permanently deleted.

After a retention policy is configured for Yammer messages, the paths the content takes depend on whether the retention policy is to retain and then delete, to retain only, or delete only.

When the retention policy is to retain and then delete:

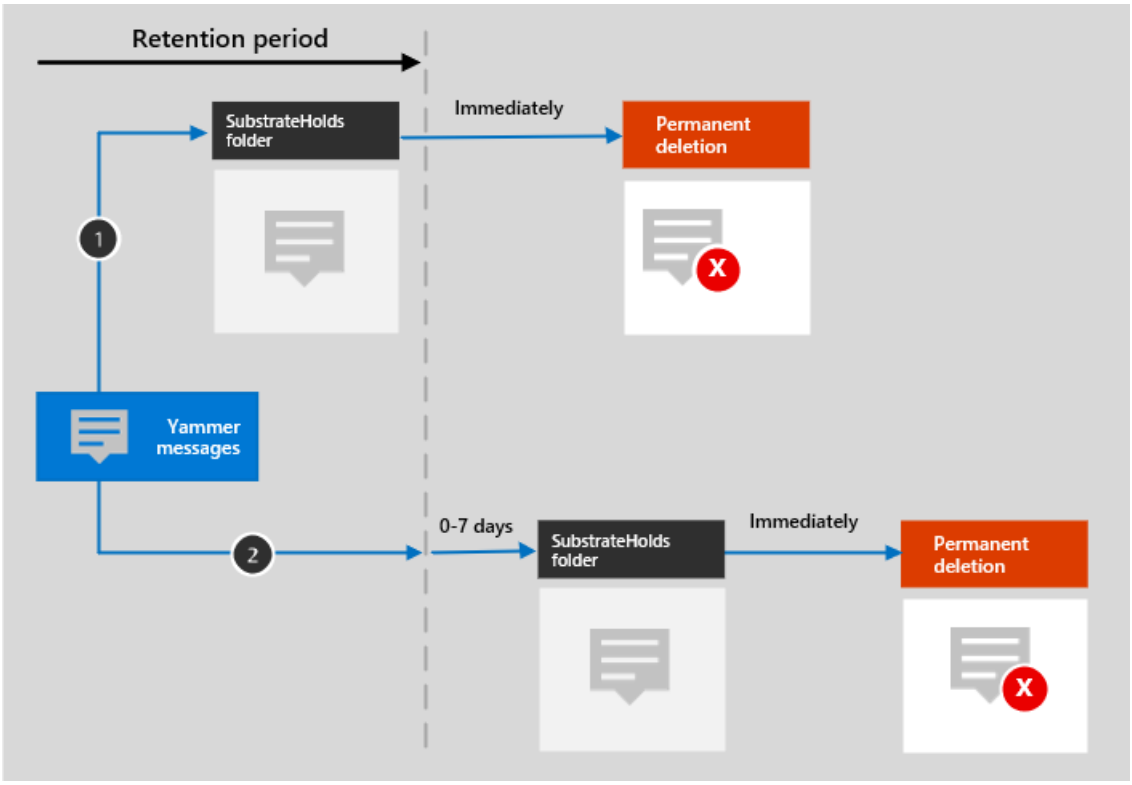

For the two paths in the diagram:

1. **If a Yammer message is edited or deleted** by the user during the retention period, the original message is immediately copied (if edited) or moved (if deleted) to the SubstrateHolds folder. The message is stored there until the retention period expires and then the message is immediately permanently deleted.
2. **If a Yammer message is not deleted** and for current messages after editing, the message is moved to the SubstrateHolds folder after the retention period expires. This action takes up to seven days from the expiry date. When the message is in the SubstrateHolds folder, it is then immediately permanently deleted.

> **NOTE**
>
> Messages in the SubstrateHolds folder are searchable by eDiscovery tools. Until messages are permanently deleted (in the SubstrateHolds folder), they remain searchable by eDiscovery tools.

When the retention policy is retain-only, or delete-only, the content's paths are variations of retain and delete.

Content paths for retain-only retention policy

1. **If a Yammer message is edited or deleted**: A copy of the original message is immediately created in the SubstrateHolds folder and retained there until the retention period expires. Then the message is immediately permanently deleted from the SubstrateHolds folder.
2. **If the Yammer message is not modified or deleted** and for current messages after editing during the retention period: Nothing happens before and after the retention period; the message remains in its original location.

Content paths for delete-only retention policy

1. **If the Yammer message is not deleted** during the retention period: At the end of the retention period, the message is moved to the SubstrateHolds folder. This action takes up to seven days from the expiry date. Then the message is immediately permanently deleted from the SubstrateHolds folder.

2. **If the Yammer message is deleted by the user** during the period, the item is immediately moved to the SubstrateHolds folder where it is immediately permanently deleted.

Messages and external users

By default, a retention policy for Yammer private messages applies to all users in your organization, but not external users. You can apply a retention policy to external users if you use the **Choose user** and specify their account.

At this time, Azure B2B guest users are not supported.

When a user leaves the organization

If a user leaves your organization and their Microsoft 365 account is deleted, their Yammer private messages that are subject to retention are stored in an inactive mailbox. These messages remain subject to any retention policy that was placed on the user before their mailbox was made inactive, and the contents are available to an eDiscovery search. For more information, see Inactive mailboxes in Exchange Online.

If the user stored any files in Yammer, see the equivalent section for SharePoint and OneDrive.

Limitations

Yammer retention policies are currently in preview and we're continuously working on optimizing retention functionality. In the meantime, be aware of the following limitation when you use retention for Yammer community messages and private messages:

- When you select **Choose users** for the **Yammer private messages** location, you might see guests and non-mailbox users. Retention policies aren't designed for these users, so don't select them.

Configuration guidance

If you're new to configuring retention in Microsoft 365, see Get started with retention policies and retention labels.

If you're ready to configure a retention policy for Yammer, see Create and configure retention policies.

Learn about retention for Exchange

11/2/2020 • 4 minutes to read • Edit Online

The information in this article supplements Learn about retention because it has information that's specific to Exchange. For other workloads, see:

- Learn about retention for SharePoint and OneDrive
- Learn about retention for Microsoft Teams
- Learn about retention for Yammer

What's included for retention and deletion

The following Exchange items can be retained and deleted by using retention policies and retention labels: Mail messages (includes drafts) with any attachments, tasks when they have an end date, and notes.

Calendar items that have an end date are supported for retention policies but aren't supported for retention labels.

Contacts, and any tasks and calendar items that don't have an end date are not supported.

Other items stored in a mailbox, such as Skype and Teams messages, aren't included in retention policies or labels for Exchange. These items have their own retention policies.

How retention works for Exchange

Both a mailbox and a public folder use the Recoverable Items folder to retain items. Only people who have been assigned eDiscovery permissions can view items in another user's Recoverable Items folder.

When a person deletes a message in a folder other than the Deleted Items folder, by default, the message moves to the Deleted Items folder. When a person deletes an item in the Deleted Items folder, the message is moved to the Recoverable Items folder. However, a user can soft delete an item (Shift+Delete) in any folder, which bypasses the Deleted Items folder and moves the item directly to the Recoverable Items folder.

When you apply retention settings to Exchange data, a timer job periodically evaluates items in the Recoverable Items folder. If an item doesn't match the rules of at least one retention policy or retention label, the item is permanently deleted (also called hard deleted) from the Recoverable Items folder.

The timer job can take up to seven days to run and the Exchange location must contain at least 10 MB.

When a user attempts to change properties of a mailbox item—such as the subject, body, attachments, senders and recipients, or date sent or received for a message—a copy of the original item is saved to the Recoverable Items folder before the change is committed. This action happens for each subsequent change. At the end of the retention period, copies in the Recoverable Items folder are permanently deleted.

After retention settings are applied to Exchange content, the paths the content takes depend on whether the retention settings are to retain and delete, to retain only, or delete only.

When the retention settings are to retain and delete:

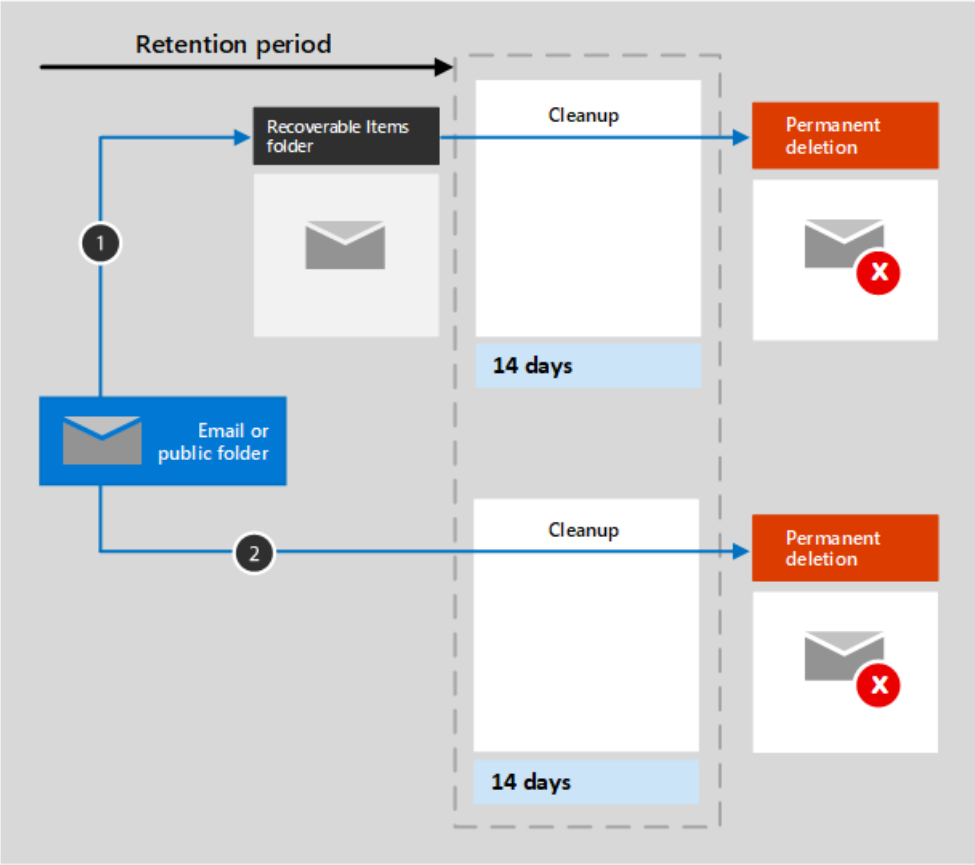

1. **If the item is modified or permanently deleted** by the user (either SHIFT+DELETE or deleted from Deleted Items) during the retention period: The item is moved (or copied, in the case of edit) to the Recoverable Items folder. There, a timer job runs periodically and identifies items whose retention period has expired, and these items are permanently deleted within 14 days of the end of the retention period. Note that 14 days is the default setting, but it can be configured up to 30 days.
2. **If the item is not modified or deleted** during the retention period: The same process runs periodically on all folders in the mailbox and identifies items whose retention period has expired, and these items are permanently deleted within 14 days of the end of the retention period. Note that 14 days is the default setting, but it can be configured up to 30 days.

When the retention settings are retain-only, or delete-only, the contents paths are variations of retain and delete:

Content paths for retain-only retention settings

1. **If the item is modified or deleted** during the retention period: A copy of the original item is created in the Recoverable Items folder and retained until the end of the retention period, when the copy in the Recoverable Items folder is permanently deleted within 14 days after the item expires.
2. **If the item is not modified or deleted** during the retention period: Nothing happens before and after the retention period; the item remains in its original location.

Content paths for delete-only retention settings

1. **If the item is not deleted** during the configured period: At the end of the configured period in the retention policy, the item is moved to the Recoverable Items folder.
2. **If the item is deleted** during the configured period: The item is immediately moved to the Recoverable Items folder. If a user deletes the item from there or empties the Recoverable Items folder, the item is permanently deleted. Otherwise, the item is permanently deleted after being in the Recoverable Items folder for 14 days.

When a user leaves the organization

If a user leaves your organization and the user's mailbox is included in a retention policy, the mailbox becomes an inactive mailbox when the user's Microsoft 365 account is deleted. The contents of an inactive mailbox are still subject to any retention policy that was placed on the mailbox before it was made inactive, and the contents are available to an eDiscovery search. For more information, see Inactive mailboxes in Exchange Online.

Configuration guidance

If you're new to configuring retention in Microsoft 365, see Get started with retention policies and retention labels.

If you're ready to configure a retention policy or retention label for Exchange, see the following instructions:

- Create and configure retention policies
- Create retention labels and apply them in apps
- Apply a retention label to content automatically

Limits for retention policies and retention label policies

2/18/2021 • 3 minutes to read • Edit Online

Microsoft 365 licensing guidance for security & compliance.

When you use retention policies and retention label policies to automatically retain or delete data for your organization, there are some maximum numbers to be aware of.

Maximum number of policies per tenant

A single tenant can have a maximum of 10,000 policies (any configuration). This maximum number includes the different policies for retention and other policies for compliance, such as DLP policies.

Maximum number of policies for retention per workload:

- Exchange Online (any configuration): 1,800
- SharePoint or OneDrive: (all sites automatically included): 13
- SharePoint or OneDrive (specific locations included or excluded): 2,600

Maximum number of items per policy

If you use the optional configuration to scope your retention settings to specific users, specific Microsoft 365 groups, or specific sites, there are some limits per policy to be aware of:

Maximum numbers of items per policy for retention:

- 1,000 mailboxes (user mailboxes or group mailboxes)
- 1,000 Microsoft 365 groups
- 1,000 users for Teams private chats
- 100 sites (OneDrive or SharePoint)

Because these limitations are per policy, if you need to use specific inclusions or exclusions that result in going over these numbers, you can create additional policies that have the same retention settings. See the next section for some example scenarios and solutions that use multiple retention policies for this reason.

However, multiple policies result in higher administrative overheads, so always confirm whether you really need inclusions and exclusions. Remember that the default configuration that applies to the entire location doesn't have any limitations, and this configuration choice might be a better solution than creating and maintaining multiple policies.

> **TIP**
> If do you need to create and maintain multiple policies for this scenario, consider using PowerShell for more efficient configuration.

Examples of using multiple policies to avoid exceeding maximum numbers

The following examples provide some design solutions for when you can't specify just the location for a retention policy, and must take into account the maximum number of items documented in the previous section.

Exchange example:

- **Requirement**: In an organization that has over 40,000 user mailboxes, most users must have their email retained for 7 years but a subset of identified users (425) must have their email retained for only 5 years.
- **Solution**: Create one retention policy for Exchange email with a retention period of 7 years and exclude the subset of users. Then create a second retention policy for Exchange email with a retention period of 5 years and include the subset of users.

  In both cases, the number included and excluded is below the maximum number of specified mailboxes for a single policy, and the subset of users must be explicitly excluded from the first policy because it has a longer retention period than the second policy. If the subset of users required a longer retention policy, you wouldn't need to exclude them from the first policy.

  With this solution, if anybody new joins the organization, their mailbox is automatically included in the first policy for 7 years and there is no impact to the maximum numbers supported. However, new users that require the 5-year retention period add to the include and exclude numbers, and this limit would be reached at 1,000.

SharePoint example:

- **Requirement**: An organization has several thousand SharePoint sites but only 2,000 sites require a retention period of 10 years, and 8,000 sites require a retention period of 4 years.
- **Solution**: Create 20 retention policies for SharePoint with a retention period of 10 years that includes 100 specific sites, and create 80 retention policies for SharePoint with a retention period of 4 years that includes 100 specific sites.

  Because you don't need to retain all SharePoint sites, you must create retention policies that specify the specific sites. Because a retention policy doesn't support more than 100 specified sites, you must create multiple policies for the two retention periods. These retention policies have the maximum number of included sites, so the next new site that needs retaining would require a new retention policy, irrespective of the retention period.

Get started with retention policies and retention labels

2/18/2021 • 4 minutes to read • Edit Online

Microsoft 365 licensing guidance for security & compliance.

Ready to start governing your organization's data by retaining the content that you need to keep, and deleting the content that you don't? Use the following high-level guidance to get started:

1. **Understand how retention works** in Microsoft 365, and then identify whether you need to use retention policies or retention labels, or a combination: Learn about retention
2. **Identify the retention settings and actions** that are required by your organization policies or industry regulations.

   As part of this assessment, determine whether you will use records management.
3. **Create retention policies and retention labels**, based on the retention settings and actions that you identified.

   For retention labels, you might find it useful to use file plan to define and refine your retention labels in a spreadsheet. Then, import that spreadsheet to create your labels.
4. **Publish and apply your retention labels**. While retention policies are designed for "set it and forget it" configuration, retention labels are reusable building blocks that can be used in multiple policies and can be incorporated into user workflows. See the list of common scenarios to help you identify how retention labels can be used.

Subscription and licensing requirements for retention policies and retention labels

A number of different subscriptions support retention policies and retention labels and the licensing requirements for users depend on the features you use.

To see the options for licensing your users to benefit from Microsoft 365 compliance features, see the Microsoft 365 licensing guidance for security & compliance. For retention, see the Information Governance section and related PDF or Excel download for feature-level licensing requirements.

Permissions required to create and manage retention policies and retention labels

Members of your compliance team who will create and manage retention policies and retention labels need permissions to the Microsoft 365 compliance center. By default, the tenant admin (global administrator) has access to this location and can give compliance officers and other people access without giving them all the permissions of a tenant admin. To grant permissions for this limited administration, we recommend that you add users to the **Compliance Administrator** admin role group.

Alternatively to using this default role, you can create a new role group and add the **Retention Management** role to this group. For a read-only role, use **View-Only Retention Management**.

For more information about role groups and roles, see Permissions in the Security & Compliance Center.

For instructions to add users to role groups and assign roles, see Give users access to the Security & Compliance Center.

These permissions are required only to create, configure, and apply retention policies and retention labels. The person configuring these policies and labels doesn't require access to the content.

Common scenarios for retention policies and retention labels

Use the following table to help you map your business requirements to retention scenarios supported by retention policies and retention labels.

I WANT TO ...	DOCUMENTATION
Efficiently set retain and delete actions by Microsoft 365 service: - Exchange - SharePoint - OneDrive - Microsoft 365 Groups - Skype for Business - Microsoft Teams - Yammer network	Create and configure retention policies
Let admins and users manually apply retain and delete actions for documents and emails: - SharePoint - OneDrive - Outlook and Outlook on the web	Create retention labels and apply them in apps
Let site admins set default retain and delete actions for all content in a SharePoint library, folder, or document set	Create retention labels and apply them in apps
Let users automatically apply retain and delete actions to emails by using Outlook rules	Create retention labels and apply them in apps
Let admins apply retain and delete actions to a document understanding model, so that these are automatically applied to identified documents in a SharePoint library	Create retention labels and apply them in apps
Automatically apply retain and delete actions to documents and emails	Apply a retention label to content automatically
Start the retention period when an event occurs, such as: - Employees leave the organization - Contracts expire - End of product lifetime	Start retention when an event occurs
Restrict changes to policies to help meet regulatory requirements or safeguard against rogue administrators	Use Preservation Lock to restrict changes to retention policies and retention label policies
Make sure somebody reviews and approves before content is deleted at the end of its retention period	Disposition reviews
Monitor how and where retain and delete settings are applied to items	Monitoring retention labels

I WANT TO ...	DOCUMENTATION
Use a single records management solution for documents and emails	Learn about records management

If you use retention labels for records management, there are additional scenarios that are unique to retention labels that mark content as a record. See Common scenarios for records management.

End-user documentation for retention labels

Retention labels, unlike retention policies, have a UI presence in Microsoft 365 apps. Make sure you provide guidance for end users and your help desk before you deploy retention labels to your production network.

The most effective end-user documentation will be customized guidance and instructions you provide for the retention label names and configurations you choose. See the following blog post for a download package that you can use to train users and drive adoption: End User Training for Retention Labels in M365 – How to Accelerate Your Adoption.

You will also find basic user instructions in the follow section: Manually apply retention labels.

Create and configure retention policies

2/18/2021 • 15 minutes to read • Edit Online

> *Microsoft 365 licensing guidance for security & compliance.*

Use a retention policy to decide proactively whether to retain content, delete content, or both - retain and then delete the content.

A retention policy lets you do this very efficiently by assigning the same retention settings for content by location, at a site or mailbox level. If you're not sure whether to use a retention policy or a retention label, see Retention policies and retention labels.

For more information about retention policies and how retention works, see Learn about retention policies and retention labels.

Before you begin

The global admin for your organization has full permissions to create and edit retention policies. If you aren't signing in as a global admin, see Permissions required to create and manage retention policies and retention labels.

Create and configure a retention policy

Although a retention policy can support multiple locations, you can't create a single retention policy that includes all the supported locations:

- Exchange email
- SharePoint site
- OneDrive accounts
- Microsoft 365 groups
- Skype for Business
- Exchange public folders
- Teams channel messages
- Teams chats
- Yammer community messages
- Yammer private messages

If you select the Teams or Yammer locations when you create a retention policy, the other locations are automatically excluded. Therefore, which instructions to follow depend on whether you need to include the Teams or Yammer locations:

- Instructions for a retention policy for Teams locations
- Instructions for a retention policy for Yammer locations
- Instructions for a retention policy for locations other than Teams and Yammer

When you have more than one retention policy, and when you also use retention labels, see The principles of retention, or what takes precedence? to understand the outcome when multiple retention settings apply to the same content.

Retention policy for Teams locations

1. From the Microsoft 365 compliance center, select **Policies** > **Retention**.
2. Select **New retention policy** to start the Create retention policy wizard, and name your new retention policy.
3. For the **Choose locations to apply the policy** page, select one or both of the locations for Teams: **Teams channel message** and **Teams chats**.

   For **Teams channel messages**, message from standard channels but not private channels are included. Currently, private channels aren't supported by retention policies.

   By default, all teams and all users are selected, but you can refine this by selecting the **Choose** and **Exclude** options.
4. For **Decide if you want to retain content, delete it, or both** page of the wizard, specify the configuration options for retaining and deleting content.

   You can create a retention policy that just retains content without deleting, retains and then deletes after a specified period of time, or just deletes content after a specified period of time. For more information, see Settings for retaining and deleting content on this page.
5. Complete the wizard to save your settings.

For more information about retention policies for Teams, see Retention policies in Microsoft Teams from the Teams documentation.

Additional retention policy needed to support Teams

Teams is more than just chats and channel messages. If you have teams that were created from a Microsoft 365 group (formerly Office 365 group), you should additionally configure a retention policy that includes that Microsoft 365 group by using the **Microsoft 365 Groups** location. This retention policy applies to content in the group's mailbox, site, and files.

If you have team sites that aren't connected to a Microsoft 365 group, you need a retention policy that includes the **SharePoint sites** or **OneDrive accounts** locations to retain and delete files in Teams:

- Files that are shared in chat are stored in the OneDrive account of the user who shared the file.
- Files that are uploaded to channels are stored in the SharePoint site for the team.

> **TIP**
>
> You can apply a retention policy to the files of just a specific team when it's not connected to a Microsoft 365 group by selecting the SharePoint site for the team, and the OneDrive accounts of users in the Team.

It's possible that a retention policy that's applied to Microsoft 365 groups, SharePoint sites, or OneDrive accounts could delete a file that's referenced in a Teams chat or channel message before those messages get deleted. In this scenario, the file still displays in the Teams message, but when users select the file, they get a "File not found" error. This behavior isn't specific to retention policies and could also happen if a user manually deletes a file from SharePoint or OneDrive.

Retention policy for Yammer locations

> **NOTE**
>
> Retention policies for Yammer are rolling out in preview. If you don't yet see the new locations for Yammer, try again in a few weeks.
>
> To use this feature, your Yammer network must be Native Mode, not Hybrid Mode.

1. From the Microsoft 365 compliance center, select **Policies** > **Retention**.
2. Select **New retention policy** to create a new retention policy.
3. For **Decide if you want to retain content, delete it, or both** page of the wizard, specify the configuration options for retaining and deleting content.

   You can create a retention policy that just retains content without deleting, retains and then deletes after a specified period of time, or just deletes content after a specified period of time. For more information, see Settings for retaining and deleting content on this page.

   Do not select **Use advanced retention settings** because this option isn't supported for Yammer locations.
4. For the **Choose locations** page, select **Let me choose specific locations**. Then toggle on one or both of the locations for Yammer: **Yammer community message** and **Yammer private messages**.

   By default, all communities and users are selected, but you can refine this by specifying communities and users to be included or excluded.

   For Yammer private messages:

   - If you leave the default at **All**, Azure B2B guest users are not included.
   - If you select **Choose user**, you can apply a retention policy to external users if you know their account.
5. Complete the wizard to save your settings.

For more information about how retention policies work for Yammer, see Learn about retention for Yammer.

Additional retention policies needed to support Yammer

Yammer is more than just community messages and private messages. To retain and delete email messages for your Yammer network, configure an additional retention policy that includes any Microsoft 365 groups that are used for Yammer, by using the **Microsoft 365 Groups** location.

To retain and delete files that are stored in Yammer, you need a retention policy that includes the **SharePoint sites** or **OneDrive accounts** locations:

- Files that are shared in private messages are stored in the OneDrive account of the user who shared the file.
- Files that are uploaded to communities are stored in the SharePoint site for the Yammer community.

It's possible that a retention policy that's applied to SharePoint sites or OneDrive accounts could delete a file that's referenced in a Yammer message before those messages get deleted. In this scenario, the file still displays in the Yammer message, but when users select the file, they get a "File not found" error. This behavior isn't specific to retention policies and could also happen if a user manually deletes a file from SharePoint or OneDrive.

Retention policy for locations other than Teams and Yammer

Use the following instructions for retention policies that apply to any of these services:

- Exchange: Email and public folders
- SharePoint: Sites
- OneDrive: Accounts
- Microsoft 365 groups
- Skype for Business

1. From the Microsoft 365 compliance center, select **Policies** > **Retention**.

2. Select **New retention policy** to start the Create retention policy wizard, and name your new retention policy.

3. For the **Choose locations** page, toggle on or off any of the locations except the locations for Teams. For each location, you can leave it at the default to apply the policy to the entire location, or specify includes and excludes.

   Information specific to locations:

   - Exchange email and Exchange public folders
   - SharePoint sites and OneDrive accounts
   - Microsoft 365 Groups
   - Skype for Business

4. For **Decide if you want to retain content, delete it, or both** page of the wizard, specify the configuration options for retaining and deleting content.

   You can create a retention policy that just retains content without deleting, retains and then deletes after a specified period of time, or just deletes content after a specified period of time. For more information, see Settings for retaining and deleting content on this page.

5. Complete the wizard to save your settings.

Configuration information for Exchange email and Exchange public folders

The **Exchange email** location supports retention for users' email, calendar, and other mailbox items, by applying retention settings at the level of a mailbox.

For detailed information about which items are included and excluded when you configure retention settings for Exchange, see What's included for retention and deletion

Note that even though a Microsoft 365 group has an Exchange mailbox, a retention policy that includes the entire **Exchange email** location won't include content in Microsoft 365 group mailboxes. To retain content in these mailboxes, select the **Microsoft 365 Groups** location.

The **Exchange public folders** location applies retention settings to all public folders and can't be applied at the folder or mailbox level.

Configuration information for SharePoint sites and OneDrive accounts

When you choose the **SharePoint sites** location, the retention policy can retain and delete documents in SharePoint communication sites, team sites that aren't connected by Microsoft 365 groups, and classic sites. Team sites connected by Microsoft 365 groups aren't supported with this option and instead, use the **Microsoft 365 Groups** location that applies to content in the group's mailbox, site, and files.

Although the retention policy is applied at the site level, only documents have retention settings applied to them. For detailed information about what's included and excluded when you configure retention settings for SharePoint and OneDrive, see What's included for retention and deletion.

When you specify your locations for SharePoint sites or OneDrive accounts, you don't need permissions to access the sites and no validation is done at the time you specify the URL on the **Edit locations** page. However, the SharePoint sites that you specify are checked that they exist at the end of the wizard. If this check fails, you see a message that validation failed for the URL you entered, and the wizard won't create the retention policy until the validation check passes. If you see this message, go back in the wizard to change the URL or remove the site from the retention policy.

To specify individual OneDrive accounts to include or exclude, the URL has the following format:

```
https://<tenant name>-my.sharepoint.com/personal/<user_name>_<tenant name>_com
```

For example, for a user in the contoso tenant that has a user name of "rsimone":

```
https://contoso-my.sharepoint.com/personal/rsimone_contoso_onmicrosoft_com
```

To verify the syntax for your tenant and identify URLs for users, see Get a list of all user OneDrive URLs in your organization.

Configuration information for Microsoft 365 Groups

To retain or delete content for a Microsoft 365 group (formerly Office 365 group), use the **Microsoft 365 Groups** location. Even though a Microsoft 365 group has an Exchange mailbox, a retention policy that includes the entire **Exchange email** location won't include content in Microsoft 365 group mailboxes. In addition, although the **Exchange email** location initially allows you to specify a group mailbox to be included or excluded, when you try to save the retention policy, you receive an error that "RemoteGroupMailbox" is not a valid selection for the Exchange location.

A retention policy applied to a Microsoft 365 group includes the group mailbox and SharePoint teams site. Files stored in the SharePoint teams site are covered with this location, but not Teams chats or Teams channel messages that have their own retention policy locations.

Configuration information for Skype for Business

Unlike Exchange email, you can't toggle the status of the Skype location on to automatically include all users, but when you turn on that location, you must then manually choose the users whose conversations you want to retain:

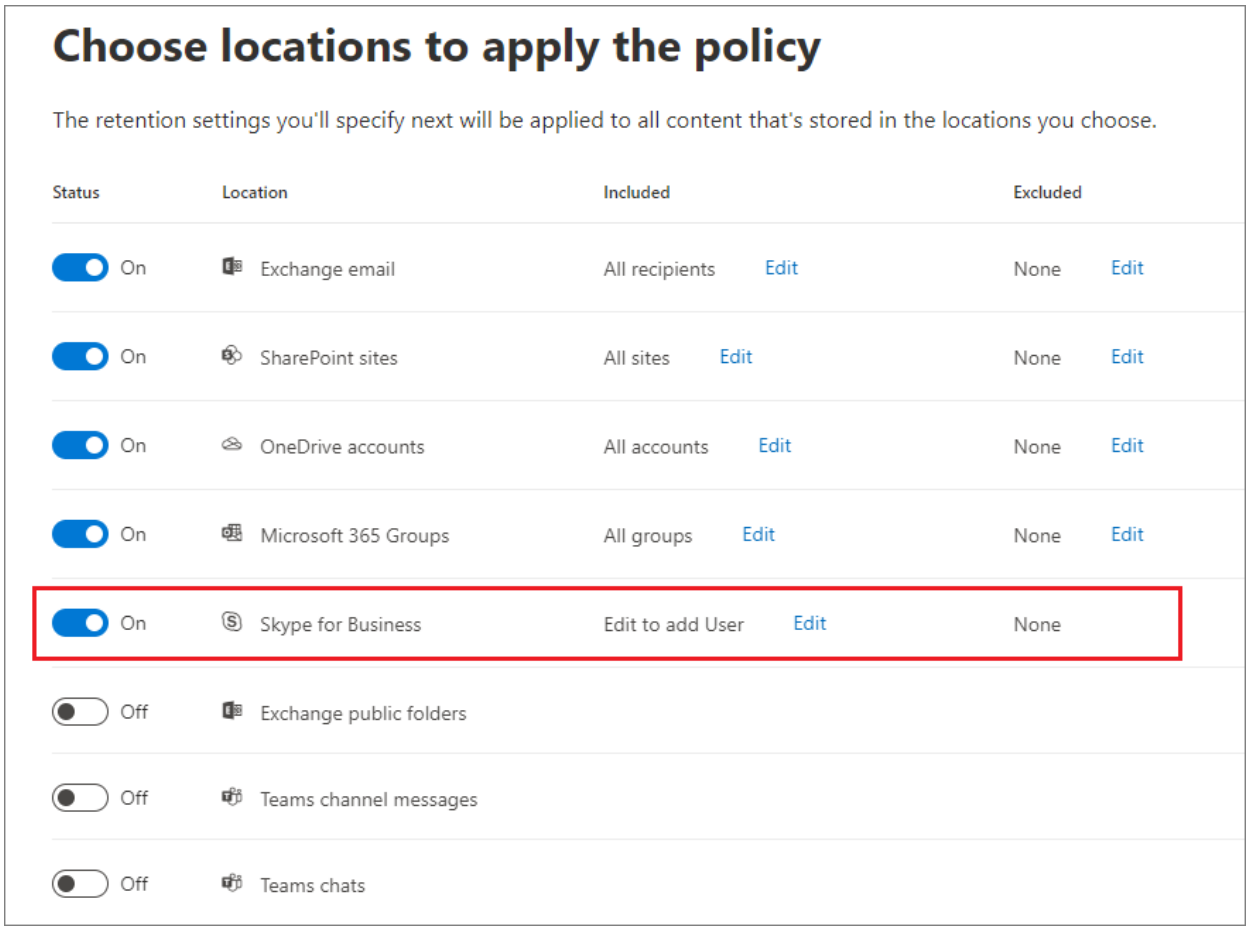

When you select **Choose user**, you can quickly include all users by selecting the **Select all** box. However, it's important to understand that each user counts as a specific inclusion in the policy. So if you include 1,000 users by selecting the **Select all** box, it's the same as if you manually selected 1,000 users to include, which is the maximum supported for Skype for Business.

Be aware that **Conversation History**, a folder in Outlook, is a feature that has nothing to do with Skype archiving. **Conversation History** can be turned off by the end user, but archiving for Skype is done by storing a copy of Skype conversations in a hidden folder that is inaccessible to the user but available to eDiscovery.

Settings for retaining and deleting content

By choosing the settings for retaining and deleting content in your retention policy, your retention policy will have one of the following configurations for a specified period of time:

- Retain-only

  For this configuration, choose **Retain items for a specific period** and **At end of the retention period: Do nothing**. Or, select **Retain items forever**.

- Retain and then delete

  For this configuration, choose **Retain items for a specific period** and **At end of the retention period: Delete items automatically**.

- Delete-only

  For this configuration, choose **Only delete items when they reach a certain age**.

Retaining content for a specific period of time

When you configure a retention policy, you choose to retain items for a specific number of days, months, or years. Or alternatively, retain the items forever.

When you configure a retention policy, you can choose to retain content indefinitely or for a specific number of days, months, or years. The retention period is calculated from the age of the content, not from when the retention policy is applied.

For the start of the retention period, you can also choose when the content was created or, supported only for files and the SharePoint, OneDrive, and Microsoft 365 Groups, when the content was last modified.

Examples:

- SharePoint: If you want to retain items in a site collection for seven years after this content is last modified, and a document in that site collection hasn't been modified in six years, the document will be retained for only another year if it's not modified. If the document is edited again, the age of the document is calculated from the new last modified date, and it will be retained for another seven years.
- Exchange: If you want to retain items in a mailbox for seven years, and a message was sent six years ago, the message will be retained for only one year. For Exchange items, the age is based on the date received for incoming email, or the date sent for outgoing email. Retaining items based on when it was last modified applies only to site content in OneDrive and SharePoint.

At the end of the retention period, you choose whether you want the content to be permanently deleted:

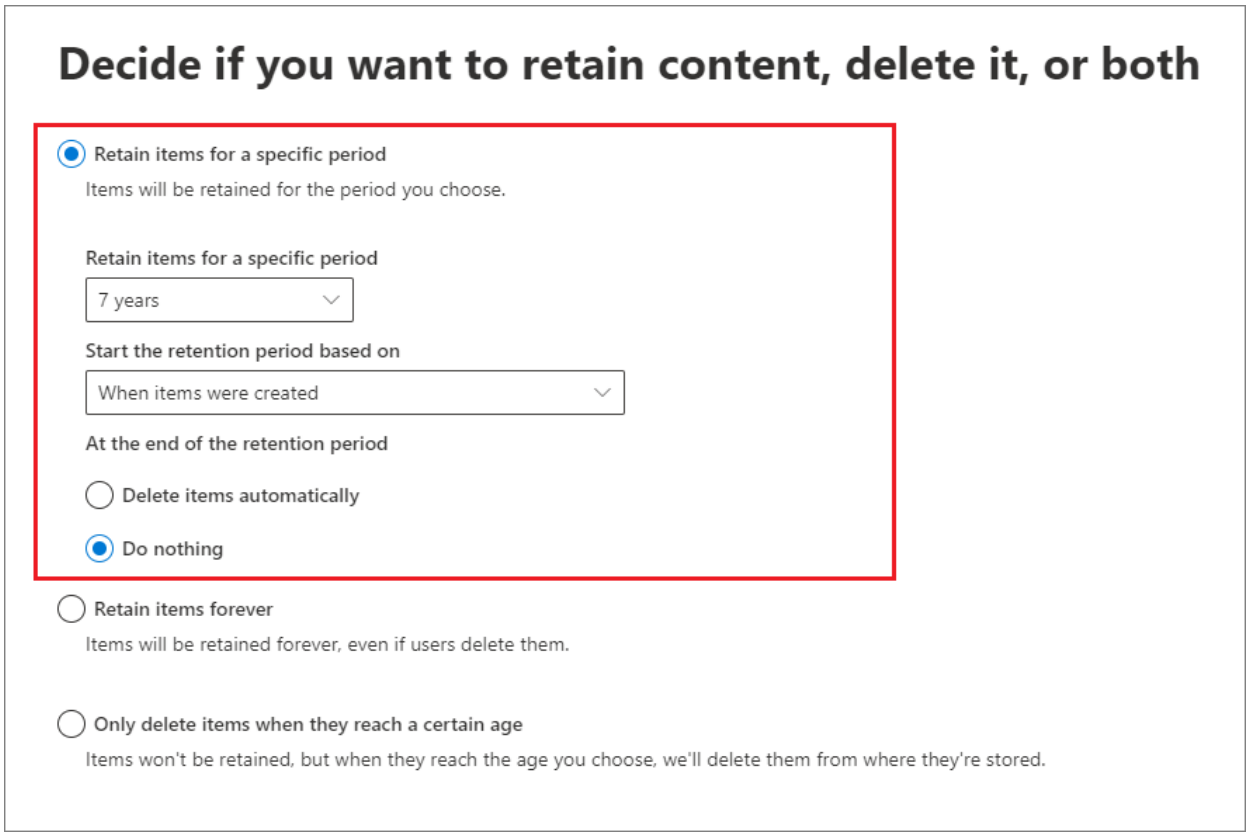

Deleting content that's older than a specific age

A retention policy can both retain and then delete items, or delete old items without retaining them.

In both cases, if your retention policy deletes items, it's important to understand that the time period specified for a retention policy is calculated from the time when the item was created or modified, and not the time since the policy was assigned.

So before you assign a retention policy for the first time, and especially when that policy deletes items, first consider the age of the existing content and how the policy may impact that content. You might also want to communicate the new policy to your users before assigning it, to give them time to assess the possible impact.

A policy that applies to entire locations

When you choose locations, with the exception of Skype for Business, the default setting is **All** when the status of the location is **On**.

When a retention policy applies to any combination of entire locations, there is no limit to the number of recipients, sites, accounts, groups, etc., that the policy can include.

For example, if a policy includes all Exchange email and all SharePoint sites, all sites and recipients will be included, no matter how many. And for Exchange, any new mailbox created after the policy is applied will automatically inherit the policy.

A policy with specific inclusions or exclusions

Be aware that if you use the optional configuration to scope your retention settings to specific users, specific Microsoft 365 groups, or specific sites, there are some limits per policy to be aware of. For more information, see Limits for retention policies and retention label policies.

To use the optional configuration to scope your retention settings, make sure the **Status** of that location is **On**, and then use the links to include or exclude specific users, Microsoft 365 groups, or sites.

WARNING

If you configure includes and then remove the last one, the configuration reverts to **All** for the location. Make sure this is the configuration that you intend before you save the policy.

For example, if you specify one SharePoint site to include in your retention policy that's configured to delete data, and then remove the single site, by default all SharePoint sites will then be subject to the retention policy that permanently deletes data. The same applies to includes for Exchange recipients, OneDrive accounts, Teams chat users etc.

In this scenario, toggle the location off if you don't want the **All** setting for the location to be subject to the retention policy. Alternatively, specify excludes to be exempt from the policy.

Updating retention policies

Some settings can't be changed after a retention policy is created and saved, which include:

- The retention policy name and the retention settings except the retention period and when to start the retention period.

If you edit a retention policy and items are already subject to the original settings in your retention policy, your updated settings will be automatically applied to these items in addition to items that are newly identified.

Usually this update is fairly quick but can take several days. When the policy replication across your Microsoft 365 locations is complete, you'll see the status of the retention policy in the Microsoft 365 compliance center change from **On (Pending)** to **On (Success)**.

Locking the policy to prevent changes

If you need to ensure that no one can turn off the policy, delete the policy, or make it less restrictive, see Use Preservation Lock to restrict changes to retention policies and retention label policies.

Use file plan to manage retention labels

11/2/2020 • 7 minutes to read • Edit Online

Microsoft 365 licensing guidance for security & compliance.

Although you can create and manage retention labels from **Information governance** in the Microsoft 365 compliance center, file plan from **Records management** has additional management capabilities:

- You can bulk-create retention labels by importing the relevant information from a spreadsheet.
- You can export the information from existing retention labels for analysis and offline collaboration, or for bulk-editing.
- More information about the retention labels is displayed to make it easier to see into and across the settings of all your retention labels from one view.
- File plan descriptors support additional and optional information for each label.

File plan can be used for all retention labels, even if they don't mark content as a record.

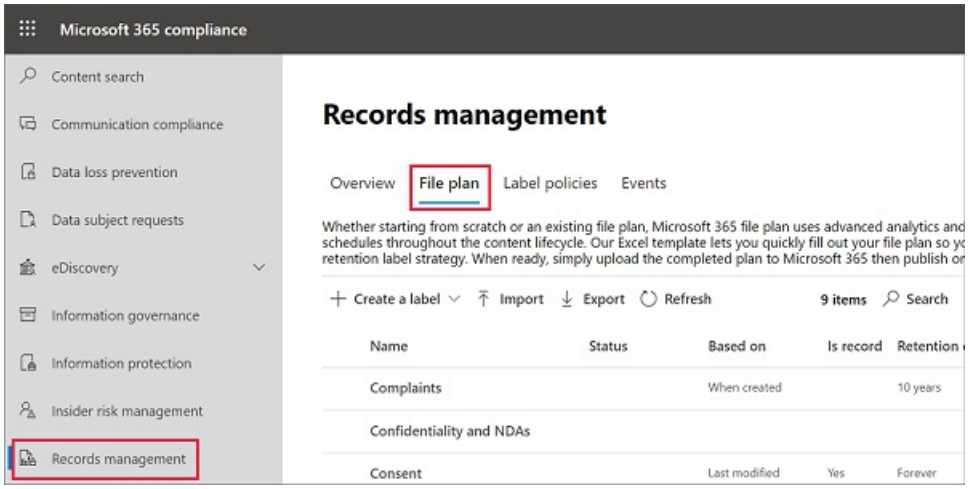

For information about what retention labels are and how to use them, see Learn about retention policies and retention labels.

Accessing file plan

To access file plan, you must have one of the following admin roles:

- Retention Manager
- View-only Retention Manager

In the Microsoft 365 compliance center, go to **Solutions** > **Records management** > **File plan**.

If **Records management** doesn't display in the navigation pane, first scroll down, and select **Show all**.

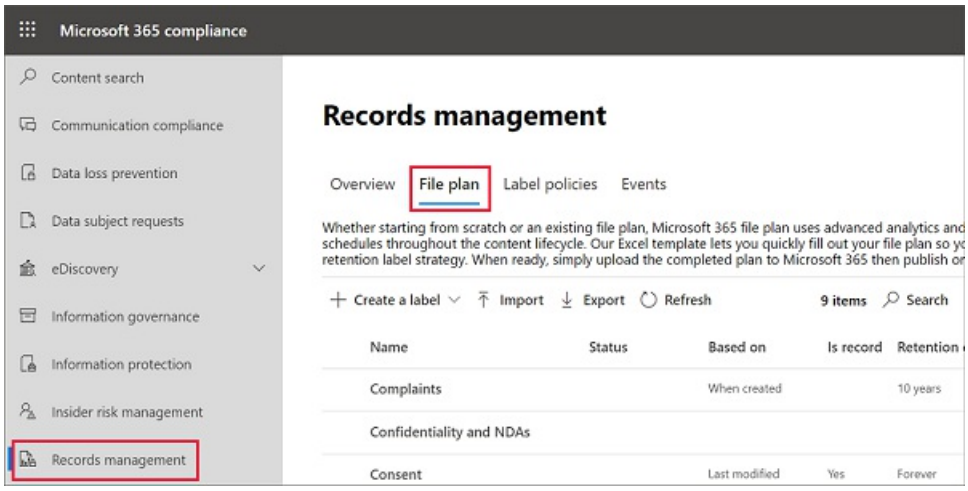

Navigating your file plan

If you've already created retention labels from **Information governance** in the Microsoft 365 compliance center, these labels automatically display in your file plan.

Similarly, if you now create retention labels in file plan, they are also available from **Information governance** if the labels aren't configured to mark content as a record.

On the **File plan** page, you see all your labels with their status and settings, optional file plan descriptors, an export option to analyze or enable offline reviews of your labels, and an import option to create retention labels.

Label settings columns

All columns except the label **Name** can be displayed or hidden by selecting the **Customize columns** option. But by default, the first few columns display information about the label status and its settings:

- **Status** identifies whether the label is included in a label policy or auto-apply policy (**Active**) or not (**Inactive**).
- **Based on** identifies how or when the retention period begins. Valid values:
  - Event
  - When created
  - Last modified
  - When labeled
- **Is record** identifies if the item is marked as a record when the label is applied. Valid values:
  - No
  - Yes
  - Yes(Regulatory)
- **Retention duration** identifies the retention period. Valid values:
  - Days
  - Months
  - Years
  - Forever
  - None
- **Disposition type** identifies what happens to the content at the end of the retention period. Valid values:
  - No action
  - Auto-delete
  - Review required

File plan descriptors columns

File plan lets you include more information as part of your retention labels. These file plan descriptors provide more options to improve the manageability and organization of the content you need to label.

By default, starting with **Reference ID**, the next few columns display these file plan descriptors that you can specify when you create a retention label, or edit an existing label.

To get you started, there are some out-of-box values for the following file plan descriptors:

- Business function/department
- Category
- Authority type
- Provision/citation

Example of file plan descriptors when you create or edit a retention label:

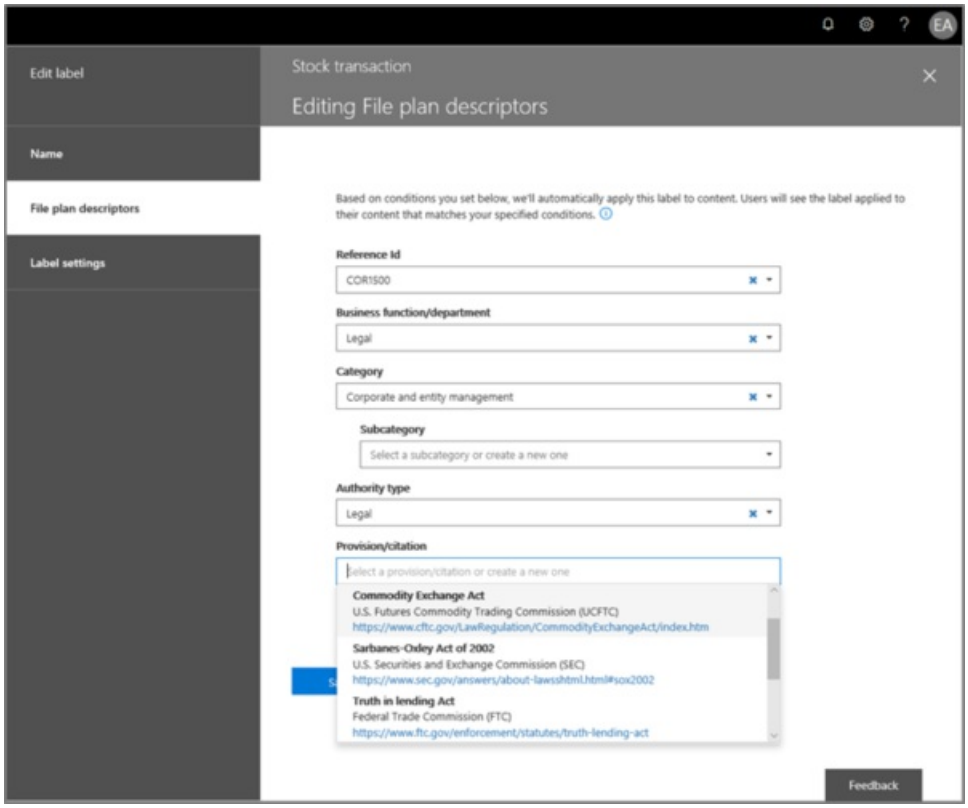

Example view of the file plan descriptors columns:

Reference Id	Function/department	Category	Subcategory	Authority type	Provision/citation
SAL1300	Sales and Marketing	Sales and Marketing		Business	
SAL1300	Sales and Marketing	Sales and Marketing		Business	
	Sales and Marketing	Strategy development r...		Legal	
SAL1100	Sales and Marketing	Sales and Marketing		Business	
LEG1100	Legal	Commercial transactions		Legal	
LEGAL-2360	Legal	Corporate and entity m...		Legal	
1234	Finance	Tax		Business	

Export all retention labels to analyze or enable offline reviews

From your file plan, you can export the details of all retention labels into a .csv file to help you facilitate periodic compliance reviews with data governance stakeholders in your organization.

To export all retention labels: On the **File plan** page, click **Export**:

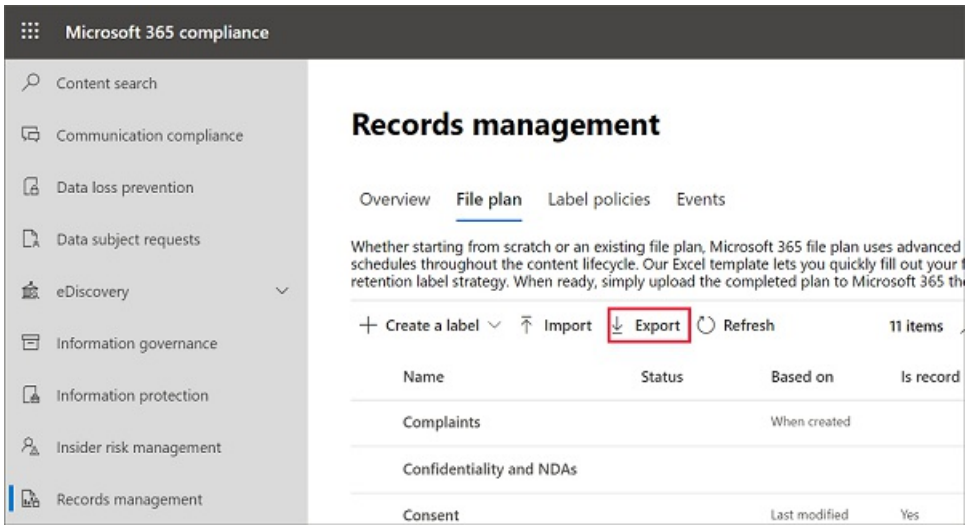

A *.csv file that contains all existing retention labels opens. For example:

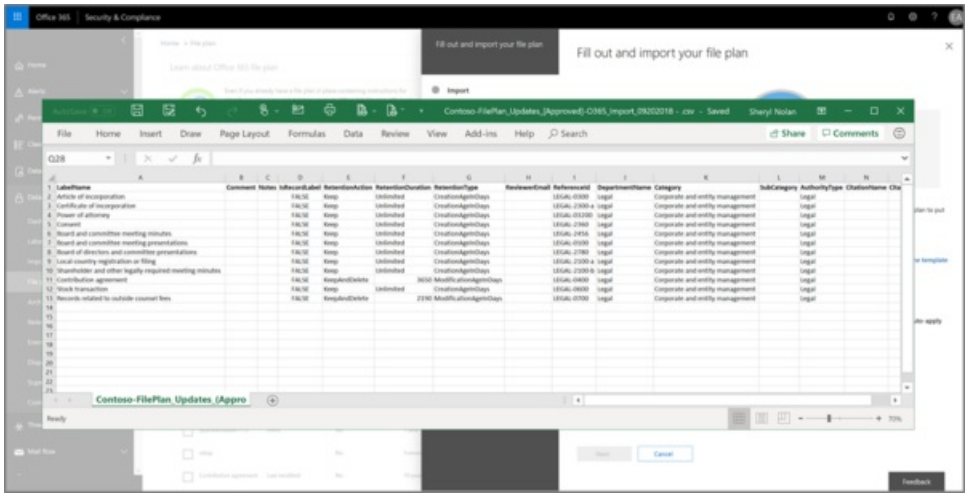

Import retention labels into your file plan

In file plan, you can bulk-import new retention labels, and use the same method to bulk-modify existing retention labels.

To import new retention labels and modify existing retention labels:

1. On the **File plan** page, click **Import** to use the **Fill out and import your file plan** page:

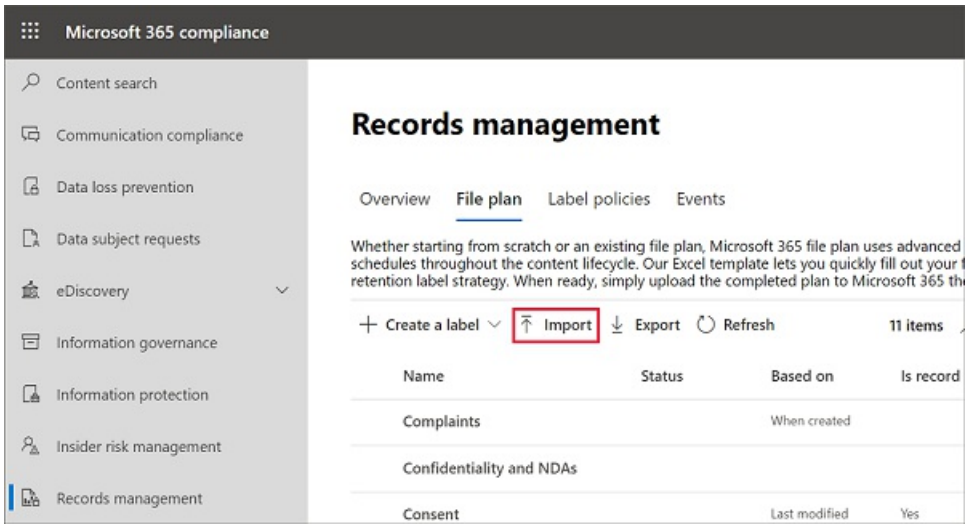

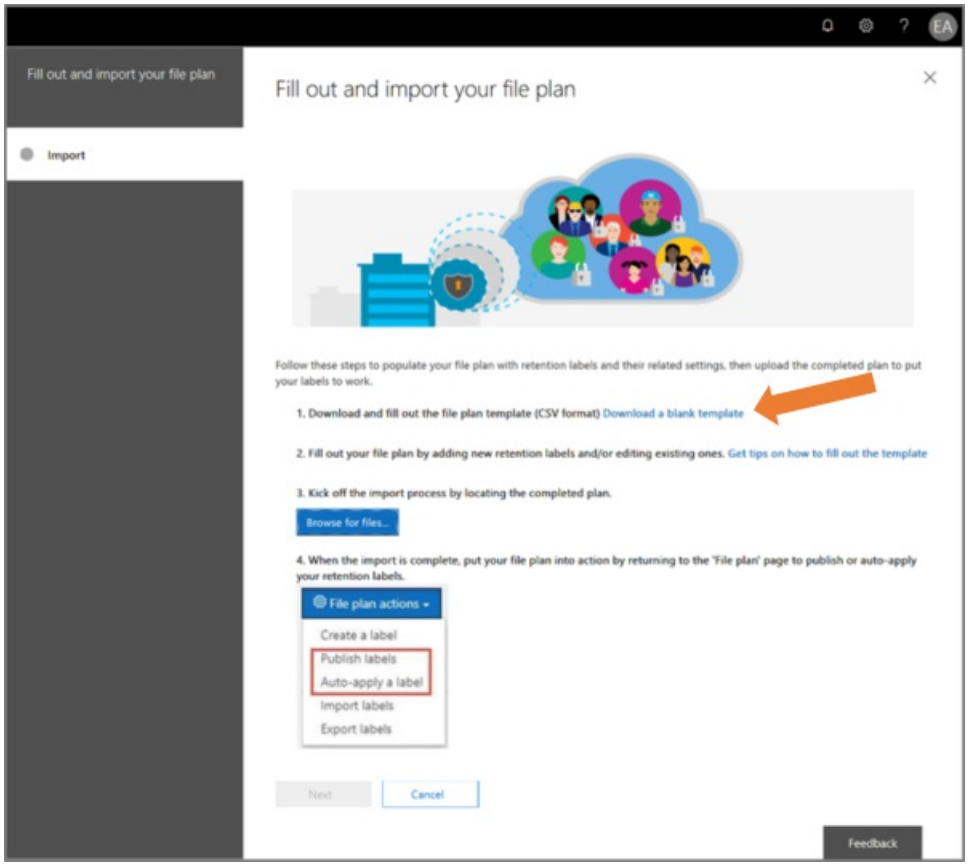

2. Download a blank template to import new retention labels. Alternatively, you can start with the .csv file that is exported when you export the existing retention labels in your organization.

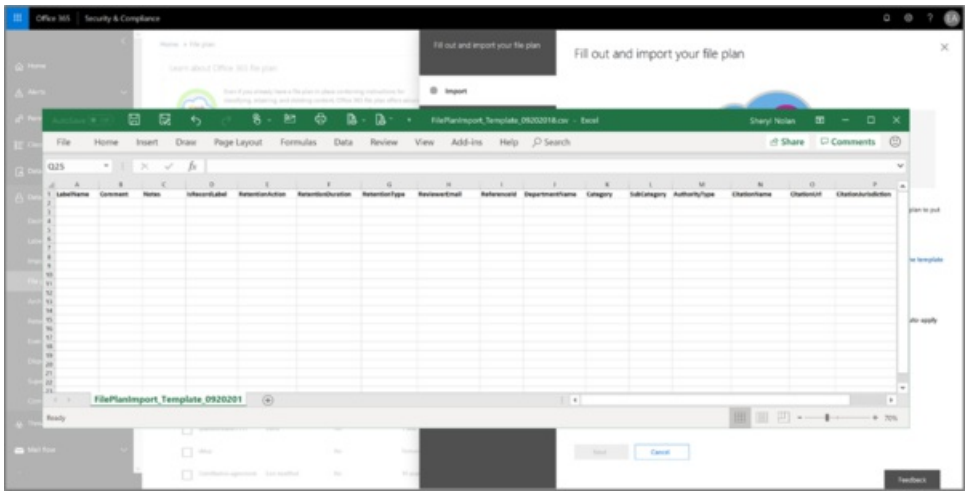

3. Fill out the template, using the following information that describes the properties and valid values for each property. For import, each value has a maximum length of 64 characters.

PROPERTY	TYPE	VALID VALUES
LabelName	String	This property specifies the name of the retention label.
Comment	String	Use this property to add a description about the retention label for admins. This description appears only to admins who manage the retention label in the compliance center.
Notes	String	Use this property to add a description about the retention label for users. This description appears when users hover over the label in apps like Outlook, SharePoint, and OneDrive. If you leave this property blank, a default description is displayed, which explains the label's retention settings.
IsRecordLabel	String	This property specifies whether the label marks the content as a record. Valid values are: **TRUE**: The label marks the item as a record and as a result, the item can't be deleted. **FALSE**: The label doesn't mark the content as a record. This is the default value.
RetentionAction	String	This property specifies what action to take after the value specified by the RetentionDuration property expires. Valid values are: **Delete**: Items older than the value specified by the RetentionDuration property are deleted. **Keep**: Retain items for the duration specified by the RetentionDuration property and then do nothing when the duration period expires. **KeepAndDelete**: Retain items for the duration specified by the RetentionDuration property and then delete them when the duration period expires.
RetentionDuration	String	This property specifies the number of days to retain the content. Valid values are: **Unlimited**: Items will be retained indefinitely. ***n***: A positive integer; for example, **365**.

PROPERTY	TYPE	VALID VALUES
RetentionType	String	This property specifies whether the retention duration is calculated from the content creation date, event date, when labeled date, or last modified date. Valid values are: **CreationAgeInDays** **EventAgeInDays** **TaggedAgeInDays** **ModificationAgeInDays**
ReviewerEmail	SmtpAddress	When this property is populated, a disposition review will be triggered when the retention duration expires. This property specifies the email address of a reviewer for the **KeepAndDelete** retention action. You can include the email address of individual users, distribution groups, or security groups. You can specify multiple email addresses separated by semicolons.
ReferenceId	String	This property specifies the value that's displayed in the **Reference Id** file plan descriptor, which you can use as a unique value to your organization.
DepartmentName	String	This property specifies the value that's displayed in the **Function/department** file plan descriptor.
Category	String	This property specifies the value that's displayed in the **Category** file plan descriptor.
SubCategory	String	This property specifies the value that's displayed in the **Sub category** file plan descriptor.
AuthorityType	String	This property specifies the value that's displayed in the **Authority type** file plan descriptor.
CitationName	String	This property specifies the name of the citation displayed in the **Provision/citation** file plan descriptor. For example, "Sarbanes-Oxley Act of 2002".
CitationUrl	String	This property specifies the URL that's displayed in the **Provision/citation** file plan descriptor.

PROPERTY	TYPE	VALID VALUES
CitationJurisdiction	String	This property specifies the jurisdiction or agency that's displayed in the **Provision/citation** file plan descriptor. For example, "U.S. Securities and Exchange Commission (SEC)".
Regulatory	String	Leave blank. This property isn't used at this time.
EventType	String	This property specifies the retention rule that's associated with the label. You can use any value that uniquely identifies the rule. For example: **Name** **Distinguished name (DN)** **GUID** You can use the Get-RetentionComplianceRule cmdlet to view the available retention rules. Note that because the EventType values are unique to an organization, if you export labels from one organization, you can't use the values for the EventType property from that organization to import labels into a different organization.

Here's an example of the template containing the information about retention labels.

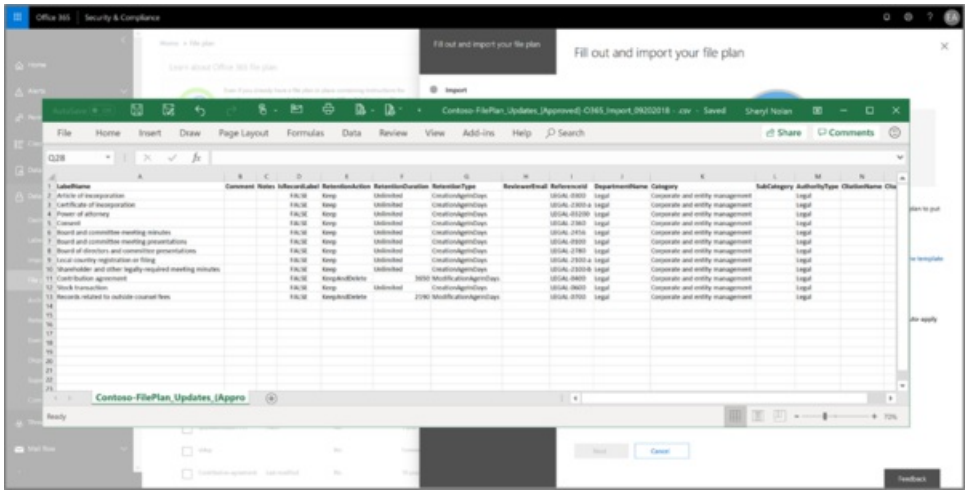

4. Under step 3 on the **Fill out and import your file plan** page, click **Browse for files** to upload the filled-out template.

   File plan validates the entries and displays the import statistics.

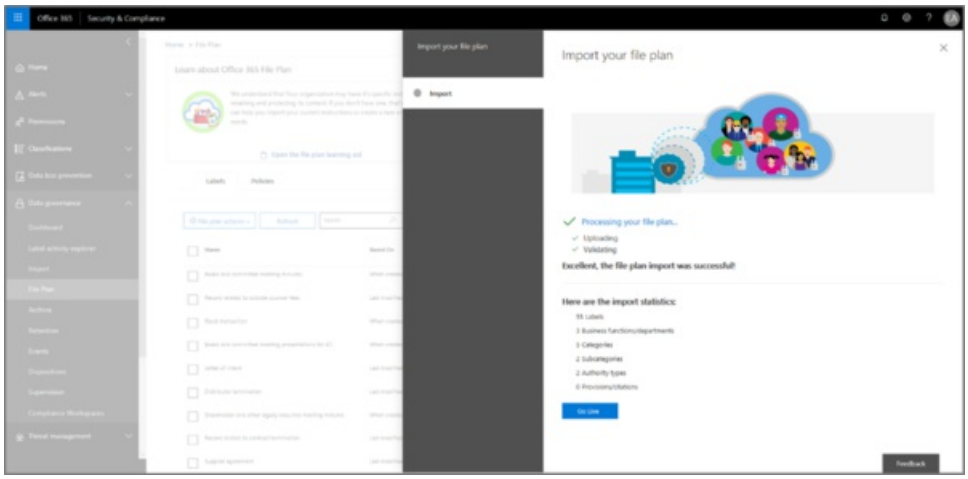

If there's a validation error, file plan import continues to validate every entry in the import file and displays all errors referencing the line and row numbers in the import file. Copy the displayed error results so you can correct them when you return to the import file.

When the import is complete, you can now add the retention labels to a new retention label policy, or auto-apply them. You can do this right from the **File plan** page by selecting the dropdown from **+ Create a label** and then **Policy to publish labels**, or **Policy to auto-apply a label**.

Next steps

For more information about creating and editing retention labels and their policies, see the following guidance:

- Create retention labels and apply them in apps
- Apply a retention label to content automatically

Create retention labels and apply them in apps

2/18/2021 • 12 minutes to read • Edit Online

Microsoft 365 licensing guidance for security & compliance.

NOTE

This scenario is supported for all retention label configurations, including regulatory records.

Use the following information to help you create and publish retention labels, and then apply them to documents and emails.

Retention labels help you retain what you need and delete what you don't at the item level (document or email). They are also used to declare an item as a record as part of a records management solution for your Microsoft 365 data.

Making retention labels available to people in your organization so that they can classify content is a two-step process:

1. Create the retention labels.
2. Publish the retention labels by using a retention label policy.

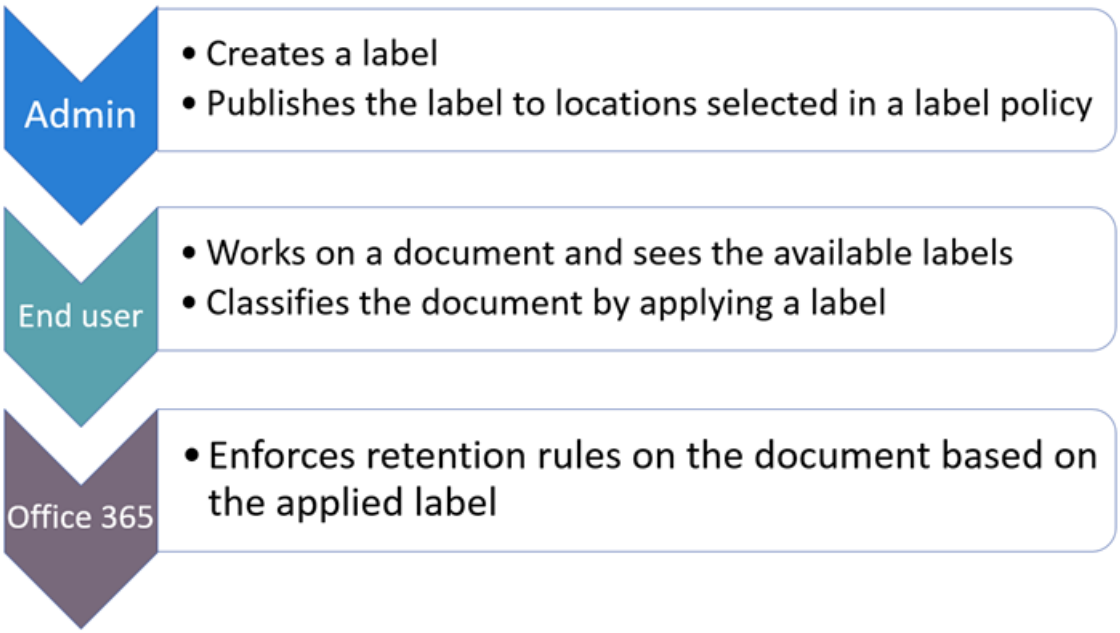

Use the following instructions for the two admin steps.

Before you begin

The global admin for your organization has full permissions to create and edit retention labels and their policies. If you aren't signing in as a global admin, see Permissions required to create and manage retention policies and retention labels.

How to create and publish retention labels

First, create your retention labels. Then create a label policy to make the labels available to apply in apps.

Where you create and configure your retention labels depend on whether you're using records management or not. Instructions are provided for both scenarios.

Step 1: Create retention labels

1. In the Microsoft 365 compliance center, navigate to one of the following locations:

   - If you are using records management:
     - **Solutions** > **Records management** > **File plan** tab > **+ Create a label** > **Retention label**
   - If you are not using records management:
     - **Solutions** > **Information governance** > **Labels** tab > + **Create a label**

   Don't immediately see your option? First select **Show all**.

2. Follow the prompts in the wizard. If you are using records management:

   - For information about the file plan descriptors, see Use file plan to manage retention labels.
   - To use the retention label to declare records, select **Mark items as records**, or **Mark items as regulatory records**. For more information, see Configuring retention labels to declare records.

3. After you have created the label and you see the options to publish the label, auto-apply the label, or just save the label: Select **Just save the label for now**, and then select **Done**.

4. Repeat these steps to create more labels.

To edit an existing label, select it, and then select the **Edit label** option to start the Edit retention wizard that lets you change the label descriptions and any eligible settings from step 2.

Step 2: Publish retention labels

Publish retention labels so that they can be applied by users in apps, such as SharePoint and Outlook.

1. In the Microsoft 365 compliance center, navigate to one of the following locations:

   - If you are using records management:
     - **Solutions** > **Records management** > > **Label policies** tab > **Publish labels**
   - If you are not using records management:
     - **Solutions** > **Information governance** > **Label policies** tab > **Publish labels**

   Don't immediately see your option? First select **Show all**.

2. Follow the prompts in the wizard.

   For information about the locations supported by retention labels, see Retention labels and locations.

To edit an existing retention label policy (the policy type is **Publish**), select it, and then select the **Edit** option to start the Edit retention policy. This wizard lets you change the policy description and any eligible settings from step 2.

When retention labels become available to apply

If you publish retention labels to SharePoint or OneDrive, those labels typically appear for end users to select within one day. However, allow up to seven days.

If you publish retention labels to Exchange, it can take up to seven days for those retention labels to appear for end users, and the mailbox must contain at least 10 MB of data.

For example:

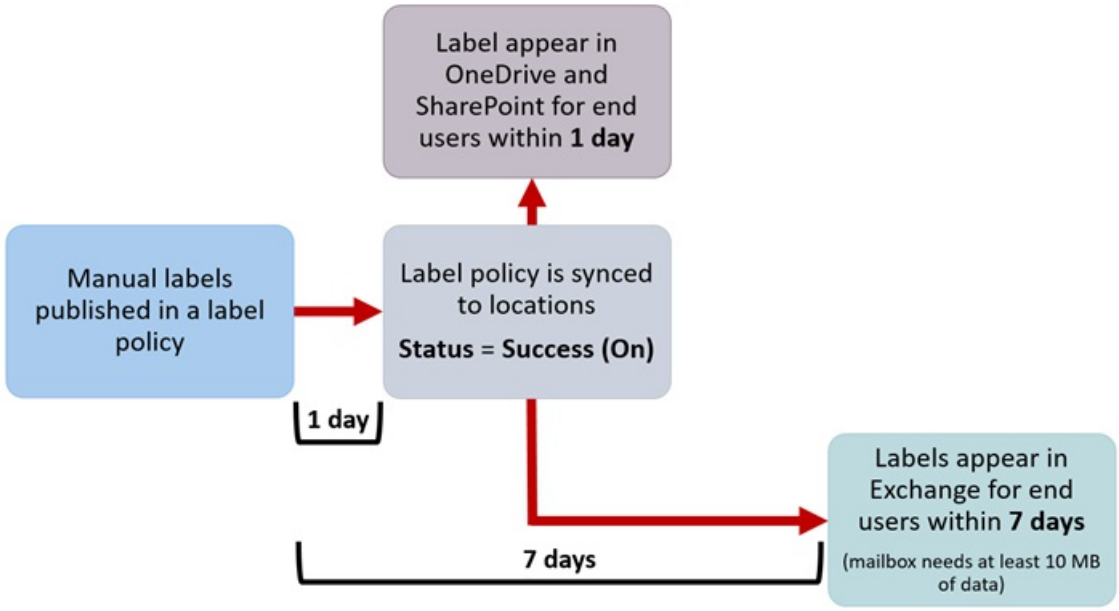

If the labels don't appear after seven days, check the **Status** of the label policy by selecting it from the **Label policies** page in the compliance center. If you see the status of **Off (Error)** and in the details for the locations see a message that it's taking longer than expected to deploy the policy (for SharePoint) or to try redeploying the policy (for OneDrive), try running Set-RetentionCompliancePolicy, a PowerShell command, to retry the policy distribution:

1. Connect to Security & Compliance Center PowerShell
2. Run the following command:

```
Set-RetentionCompliancePolicy -Identity <policy name> -RetryDistribution
```

How to check on the status of retention labels published to Exchange

In Exchange Online, retention labels are made available to end users by a process that runs every seven days. By using PowerShell, you can see when this process last ran and therefore identify when it will run again.

1. Connect to Exchange Online PowerShell.
2. Run these commands.

```
$logProps = Export-MailboxDiagnosticLogs <user> -ExtendedProperties
```

```
$xmlprops = [xml]($logProps.MailboxLog)
```

```
$xmlprops.Properties.MailboxTable.Property | ? {$_.Name -like "ELC*"}
```

In the results, the `ELCLastSuccessTimeStamp` (UTC) property shows when the system last processed your mailbox. If it has not happened since the time you created the policy, the labels are not going to appear. To force processing, run `Start-ManagedFolderAssistant -Identity <user>`.

If labels aren't appearing in Outlook on the web and you think they should be, make sure to clear the cache in your browser (CTRL+F5).

How to apply published retention labels

Use the following sections to learn how published retention labels can be applied in apps:

- Manually apply retention labels
- Applying a default retention label to all content in a SharePoint library, folder, or document set
- Automatically applying a retention label to email by using rules

In addition, when you use SharePoint Syntex and publish retention labels to SharePoint locations, you can apply a retention label to a document understanding model so that identified documents are automatically labeled.

After content is labeled, see the following information to understand when the applied label can be removed or changed: Only one retention label at a time.

Manually apply retention labels

End users, as well as administrators, can manually apply retention labels from the following locations:

- Outlook and Outlook on the web
- OneDrive
- SharePoint
- Microsoft 365 groups (both the group site and group mailbox in Outlook on the web)

Use the following sections to understand how to apply retention labels.

Applying retention labels in Outlook

To label an item in the Outlook desktop client, select the item. On the **Home** tab on the ribbon, click **Assign Policy**, and then choose the retention label.

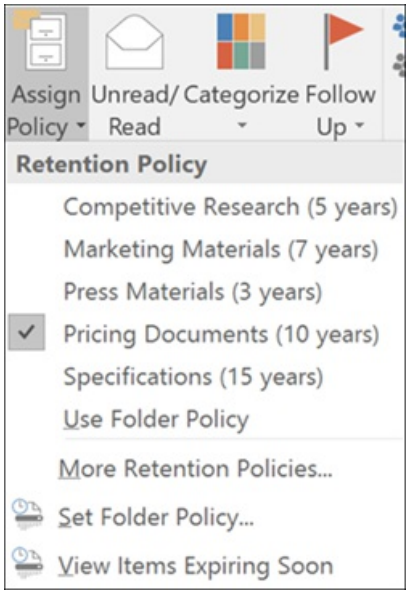

You can also right-click an item, click **Assign Policy** in the context menu, and then choose the retention label.

After the retention label is applied, you can view that retention label and what action it takes at the top of the item. If an email has a retention label applied that has an associated retention period, you can see at a glance when the email expires.

Applying a default retention label to an Outlook folder

You can apply retention labels to Outlook folders as a default label that can be inherited by messages in that folder. Right-click the folder, select **Properties**, the **Policy** tab, and select the retention label you want to use as that folder's default retention label.

When you use a a standard retention label as your default label for an Outlook folder:

- All unlabeled items in the folder have this retention label applied.
- The inheritance flows to any child folders and items inherit the label from their nearest folder.
- Items that are already labeled retain their retention label, unless it was applied by a different default label.
- If you change or remove the default retention label for the folder: Existing retention labels applied to items in that folder are also changed or removed only if those labels were applied by a default label.
- If you move an item with a default retention label from one folder to another folder with a different default retention label: The item gets the new default retention label.
- If you move an item with a default retention label from one folder to another folder with no default retention label: The old default retention label is removed.

When labels are applied that aren't standard retention labels but mark items as records (or regulatory records), these labels can only be manually changed or removed.

Applying retention labels in Outlook on the web

To label an item in Outlook on the web, right-click the item > **Assign policy** > choose the retention label.

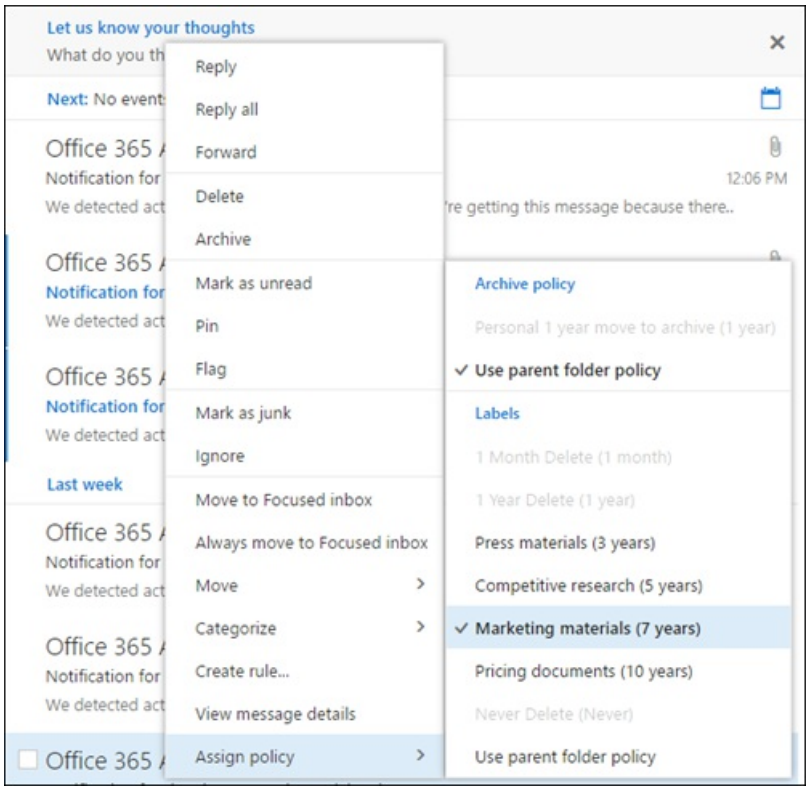

After the retention label is applied, you can view that retention label and what action it takes at the top of the item. If an email is classified and has an associated retention period, you can know at a glance when the email will expire.

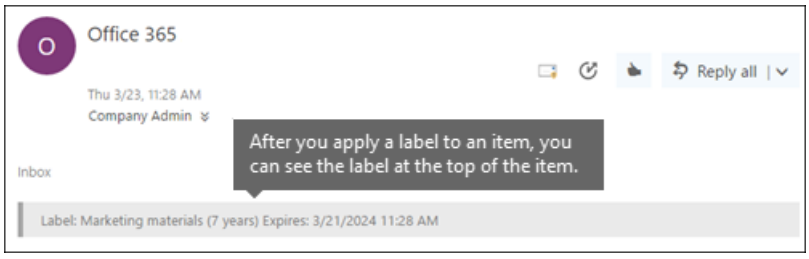

As with the desktop version of Outlook on the web, you can also apply retention labels to folders. Right-click the folder, select **Assign policy**, and change **Use parent folder policy** to the retention label you want to use as

that folder's default retention label.

Applying retention labels in OneDrive and SharePoint

To label a document (including OneNote files) in OneDrive or SharePoint, select the item > in the upper-right corner, choose **Open the details pane** > **Apply retention label** > choose the retention label.

You can also apply a retention label to a folder or document set, and you can set a default retention label for a document library.

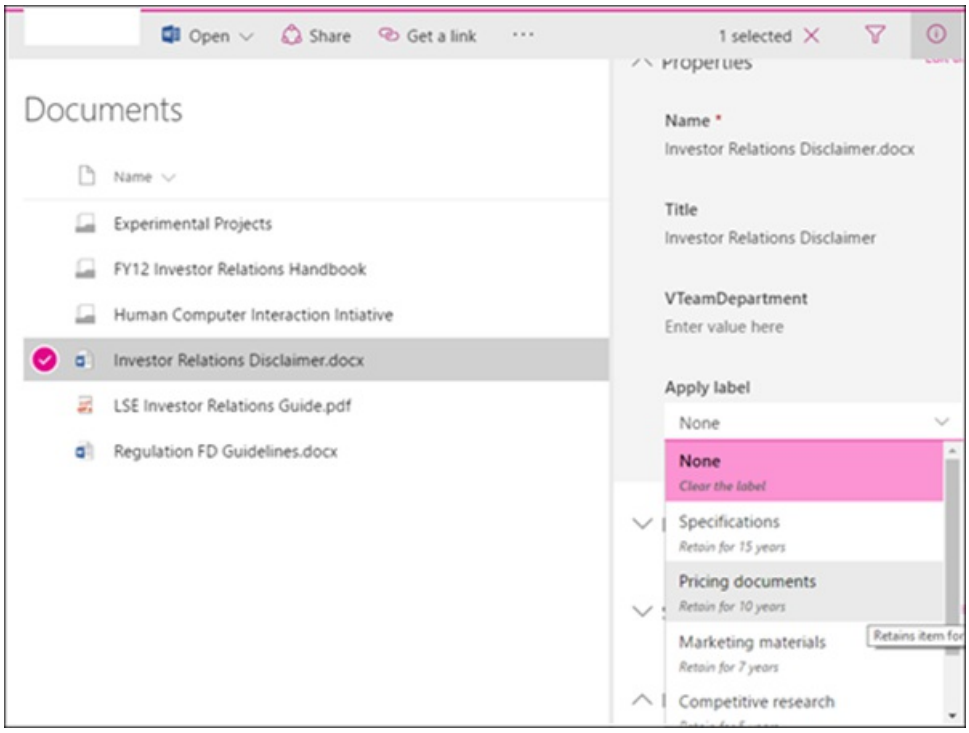

After a retention label is applied to an item, you can view it in the details pane when that item's selected.

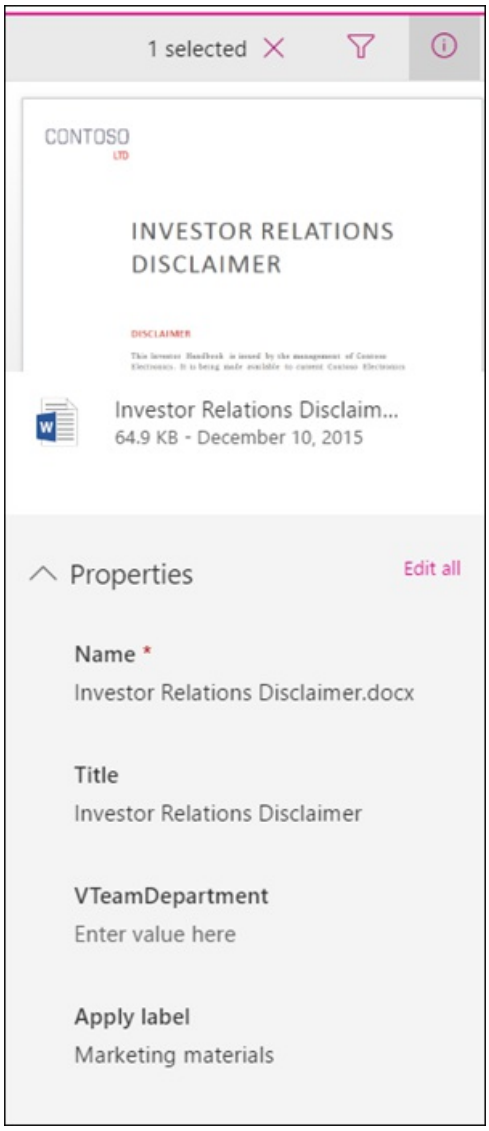

For SharePoint, but not OneDrive, you can create a view of the library that contains the **Labels** column or **Item is a Record** column. This view lets you see at a glance the retention labels assigned to all items and which items are records. Note, however, that you can't filter the view by the **Item is a Record** column. For instructions how to add columns, see Show or hide columns in a list or library.

Applying retention labels in Microsoft 365 groups

When you publish retention labels to Microsoft 365 groups (formerly Office 365 groups), the retention labels appear in both the group site and group mailbox in Outlook on the web. The experience of applying a retention label to content is identical to that for email and documents.

To retain content for a Microsoft 365 group, use the **Microsoft 365 Groups** location. Even though a Microsoft 365 group has an Exchange mailbox, a retention policy that includes the entire Exchange location won't include content in Microsoft 365 group mailboxes.

In addition, it's not possible to use the Exchange location to include or exclude a specific group mailbox. Although the Exchange location initially allows a group mailbox to be selected, when you try to save the retention policy, you receive an error that "RemoteGroupMailbox" is not a valid selection for the Exchange location.

First, create and configure the sensitivity labels that you want to make available for apps and other services. For example, the labels you want users to see and apply from Office apps.

Then, create one or more label policies that contain the labels and policy settings that you configure. It's the label policy that publishes the labels and settings for your chosen users and locations.

Applying a default retention label to all content in a SharePoint library, folder, or document set

This method requires retention labels to be published to a retention label policy.

In addition to letting people apply a retention label to individual documents, you can also apply a default retention label to a SharePoint library, folder, or document set. In this scenario, documents in that location can inherit your selected default retention label. Although the same label is applied, each document will be retained and deleted separately, according to the start of the retention period setting in the label.

For a document library, the default label configuration is done on the **Library settings** page for a document library. When you choose the default retention label, you can also choose to apply it to existing items in the library.

For example, if you have a retention label for marketing materials, and you know a specific document library contains only that type of content, you can make the **Marketing Materials** retention label the default label for all documents in that library.

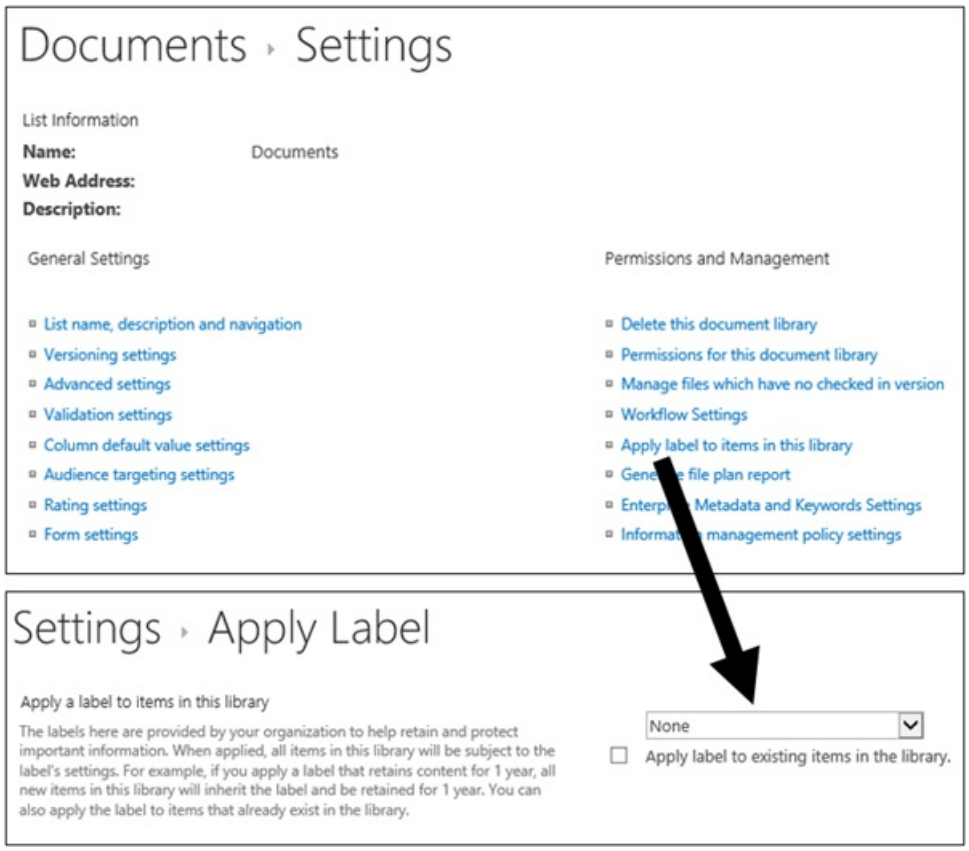

Label behavior when you use a default label for SharePoint

For standard retention labels that you apply as a default retention label to a library, folder, or document set:

- All new, unlabeled items in the container will have this retention label applied.
- For folders, the inheritance flows to any child folders and items inherit the label from their nearest folder.
- If you selected the option to apply the default label to existing items: Items that are already labeled retain their retention label, unless it was applied by a different default label.
- If you change the default retention label for the container: Existing retention labels applied to items in that container are changed only if you selected the option to apply the default label to existing items and those labels were applied by a default label.
- If you remove the default retention label for the container: Items retain their labels.
- If you move an item with a default retention label applied from one container to another container: The item keeps its existing default retention label, even if the new location has a different default retention label. Only if you then change the default label for this new location can the moved item inherit the default label from its current location.

When labels are applied that aren't standard retention labels but mark items as records (or regulatory records), these labels can only be manually changed or removed.

Automatically applying a retention label to email by using rules

In Outlook, you can create rules to apply a retention label.

For example, you can create a rule that applies a specific retention label to all messages sent to or from a specific distribution group.

To create a rule, right-click an item > **Rules** > **Create Rule** > **Advanced Options** > **Rules Wizard** > **apply retention policy**.

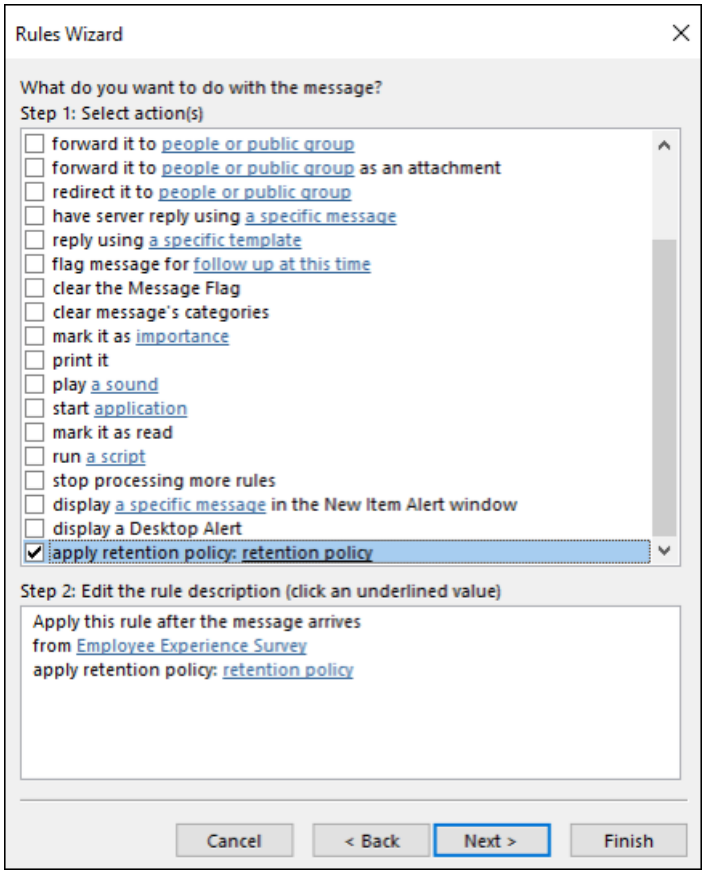

Although the UI refers to retention policies, it's your retention labels that display here and can be selected, not your retention policies.

Updating retention labels and their policies

When you edit a retention label or retention label policy, and the retention label or policy is already applied to content, your updated settings will automatically be applied to this content in addition to content that's newly identified.

Some settings can't be changed after the label or policy is created and saved, which include:

- The retention label and policy name, and the retention settings except the retention period. However, you can't change the retention period when the retention period is based on when items were labeled.
- The option to mark items as a record.

Locking the policy to prevent changes

If you need to ensure that no one can turn off the policy, delete the policy, or make it less restrictive, see Use Preservation Lock to restrict changes to retention policies and retention label policies.

Next steps

Event-based retention is another supported scenario for retention labels. For more information, see the following articles:

- Start retention when an event occurs
- Automate event-based retention
- Use retention labels to manage the lifecycle of documents stored in SharePoint

Automatically apply a retention label to retain or delete content

2/18/2021 • 11 minutes to read • Edit Online

Microsoft 365 licensing guidance for security & compliance.

NOTE

This scenario is not supported for regulatory records.

One of the most powerful features of retention labels is the ability to apply them automatically to content that matches specified conditions. In this case, people in your organization don't need to apply the retention labels. Microsoft 365 does the work for them.

Auto-applying retention labels are powerful because:

- You don't need to train your users on all of your classifications.
- You don't need to rely on users to classify all content correctly.
- Users no longer need to know about data governance policies - they can focus on their work.

You can apply retention labels to content automatically when that content contains sensitive information, keywords or searchable properties, or a match for trainable classifiers.

TIP

Now in preview, use searchable properties to identify Teams meeting recordings.

The processes to automatically apply a retention label based on these conditions:

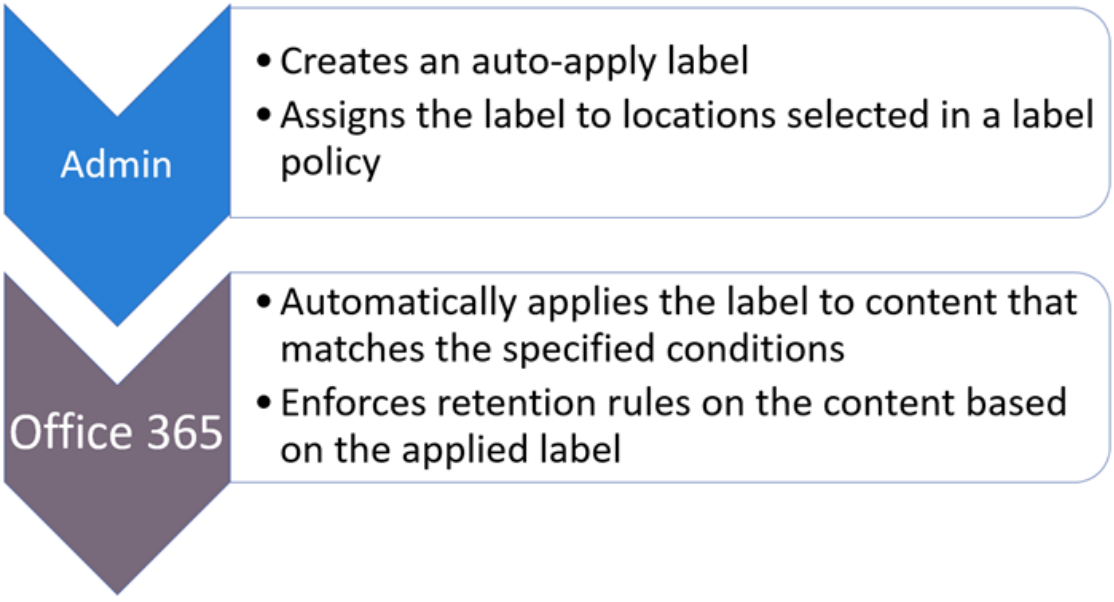

Use the following instructions for the two admin steps.

NOTE

Auto-policies use service-side labeling with conditions to automatically apply retention labels. You can also automatically apply a retention label with a label policy when you do the following:

- Apply a retention label to a document understanding model in SharePoint Syntex
- Apply a default retention label for SharePoint and Outlook
- Apply a retention label to email by using Outlook rules

For these scenarios, see Create and apply retention labels in apps.

Before you begin

The global admin for your organization has full permissions to create and edit retention labels and their policies. If you aren't signing in as a global admin, see Permissions required to create and manage retention policies and retention labels.

How to auto-apply a retention label

First, create your retention label. Then create an auto-policy to apply that label. If you have already created your retention label, skip to creating an auto-policy.

Navigation instructions depend on whether you're using records management or not. Instructions are provided for both scenarios.

Step 1: Create a retention label

1. In the Microsoft 365 compliance center, navigate to one of the following locations:

   - If you are using records management:
     - **Solutions** > **Records management** > **File plan** tab > **+ Create a label** > **Retention label**
   - If you are not using records management:
     - **Solutions** > **Information governance** > **Labels** tab > **+ Create a label**

   Don't immediately see your option? First select **Show all**.

2. Follow the prompts in the wizard. If you are using records management:

   - For information about the file plan descriptors, see Use file plan to manage retention labels
   - To use the retention label to declare records, select **Mark items as records**, or **Mark items as regulatory records**. For more information, see Configuring retention labels to declare records.

3. After you have created the label and you see the options to publish the label, auto-apply the label, or just save the label: Select **Auto-apply this label to a specific type of content**, and then select **Done** to start the Create auto-labeling wizard that takes you directly to step 2 in the following procedure.

To edit an existing label, select it, and then select the **Edit label** option to start the Edit retention wizard that lets you change the label descriptions and any eligible settings from step 2.

Step 2: Create an auto-apply policy

When you create an auto-apply policy, you select a retention label to automatically apply to content, based on the conditions that you specify.

1. In the Microsoft 365 compliance center, navigate to one of the following locations:

   - If you are using records management: **Information governance**:

    - **Solutions** > **Records management** > **Label policies** tab > **Auto-apply a label**
- If you are not using records management:
    - **Solutions** > **Information governance** > **Label policies** tab > **Auto-apply a label**

  Don't immediately see your option? First select **Show all**.

2. Follow the prompts in the Create auto-labeling wizard.

   For information about configuring the conditions that automatically apply the retention label, see the Configuring conditions for auto-apply retention labels section on this page.

   For information about the locations supported by retention labels, see the Retention labels and locations section.

To edit an existing auto-apply policy, select it to start the Edit retention policy wizard that lets you change the selected retention label and any eligible settings from step 2.

After content is labeled by using an auto-apply label policy, the applied label can't be automatically removed or changed by changing the content or the policy, or by a new auto-apply label policy. For more information, see Only one retention label at a time.

Configuring conditions for auto-apply retention labels

You can apply retention labels to content automatically when that content contains:

- Specific types of sensitive information
- Specific keywords or searchable properties that match a query you create
- A match for trainable classifiers

Auto-apply labels to content with specific types of sensitive information

> **WARNING**
> This configuration currently has a known limitation where all unlabeled emails always have the selected retention label applied when there is a match for your chosen sensitive information types. For example, even if you scope your auto-apply policy to specific users, or select locations other than Exchange for the policy, the label is always applied to unlabeled emails when there is a match.

When you create auto-apply retention label policies for sensitive information, you see the same list of policy templates as when you create a data loss prevention (DLP) policy. Each template is preconfigured to look for specific types of sensitive information. For example, the template shown here looks for U.S. ITIN, SSN, and passport numbers from the **Privacy** category, and **U.S Personally Identifiable Information (PII) Data** template:

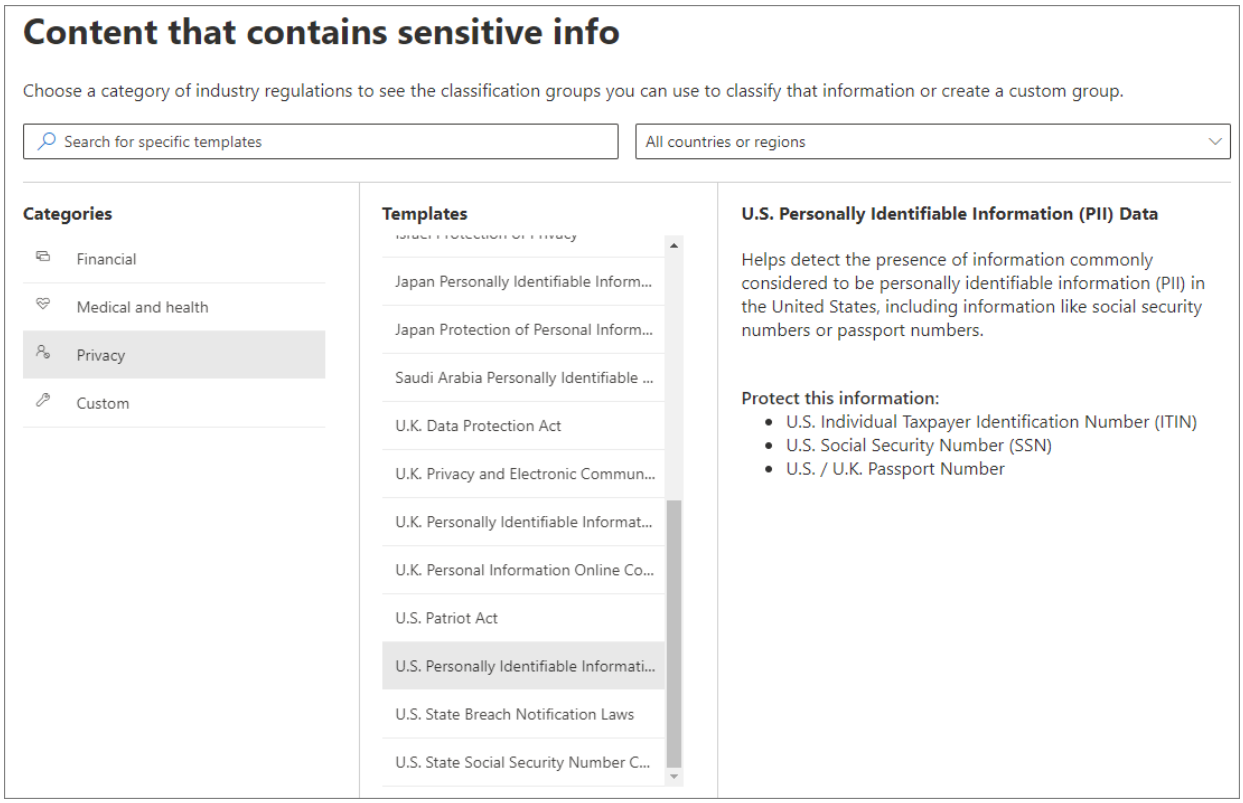

To learn more about the sensitivity information types, see Sensitive information type entity definitions.

After you select a policy template, you can add or remove any types of sensitive information, and you can change the instance count and match accuracy. In the example screenshot shown next, a retention label will be auto-applied only when:

- The type of sensitive information that's detected has a match accuracy (or confidence level) of at least 75. Many sensitive information types are defined with multiple patterns, where a pattern with a higher match accuracy requires more evidence to be found (such as keywords, dates, or addresses), while a pattern with a lower match accuracy requires less evidence. The lower the **min** match accuracy, the easier it is for content to match the condition.
- The content contains between 1 and 9 instances of any of these three sensitive information types. You can delete the **to** value so that it changes to **Any**.

For more information about these options, see the following guidance from the DLP documentation Tuning rules to make them easier or harder to match.

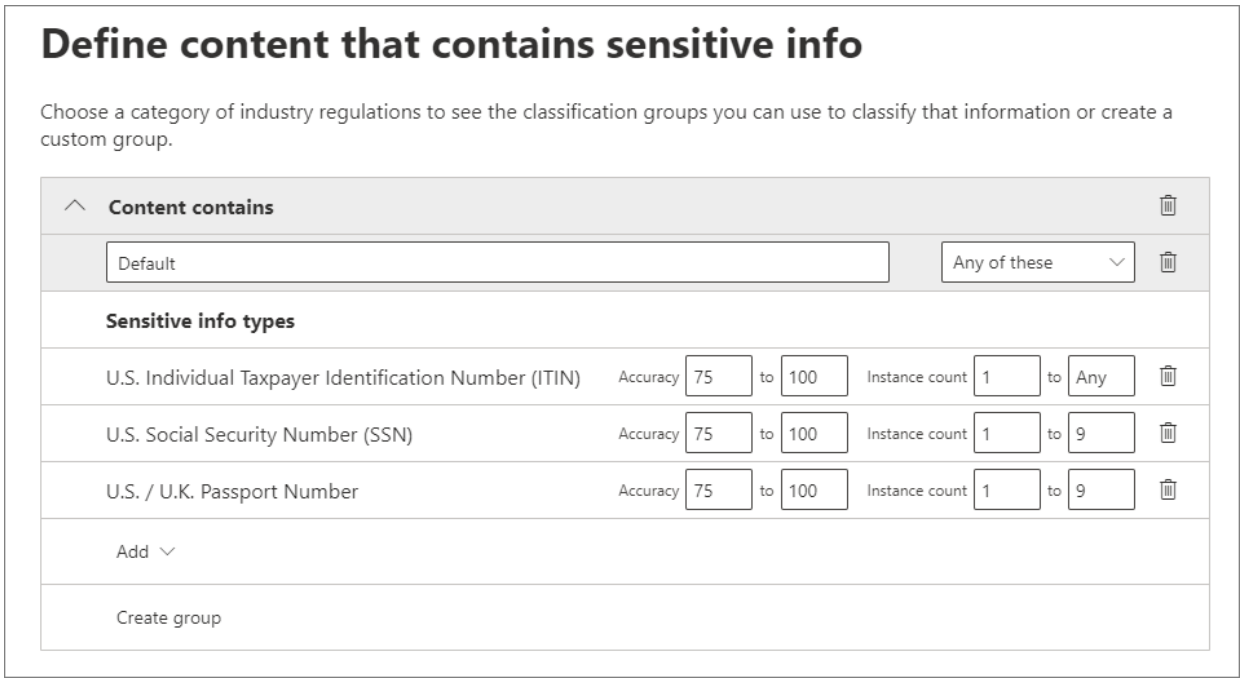

To consider when using sensitive information types to auto-apply retention labels:

- New and modified items can be auto-labeled.

Auto-apply labels to content with keywords or searchable properties

You can auto-apply labels to content by using a query that contains specific words, phrases, or values of searchable properties. You can refine your query by using search operators such as AND, OR, and NOT.

For more information about the query syntax that uses Keyword Query Language (KQL), see Keyword Query Language (KQL) syntax reference.

Query-based auto-apply policies use the same search index as eDiscovery content search to identify content. For more information about the searchable properties that you can use, see Keyword queries and search conditions for Content Search.

Some things to consider when using keywords or searchable properties to auto-apply retention labels:

- New, modified, and existing items will be auto-labeled for SharePoint, OneDrive, and Exchange.
- For SharePoint, crawled properties and custom properties aren't supported for these KQL queries and

you must use only predefined managed properties. However, you can use mappings at the tenant level with the predefined managed properties that are enabled as refiners by default (RefinableDate00-19, RefinableString00-99, RefinableInt00-49, RefinableDecimals00-09, and RefinableDouble00-09). For more information, see Overview of crawled and managed properties in SharePoint Server, and for instructions, see Create a new managed property.

- If you map a custom property to one of the refiner properties, wait 24 hours before you use it in your KQL query for a retention label.
- Although SharePoint managed properties can be renamed by using aliases, don't use these for KQL queries in your labels. Always specify the actual name of the managed property, for example, "RefinableString01".
- To search for values that contain spaces or special characters, use double quotation marks (`" "`) to contain the phrase; for example, `subject:"Financial Statements"` .
- Use the *DocumentLink* property instead of *Path* to match an item based on its URL.
- Suffix wildcard searches (such as `*cat`) or substring wildcard searches (such as `*cat*`) aren't supported. However, prefix wildcard searches (such as `cat*`) are supported.
- Be aware that partially indexed items can be responsible for not labeling items that you're expecting, or labeling items that you're expecting to be excluded from labeling when you use the NOT operator. For more information, see Partially indexed items in Content Search.

Examples queries:

WORKLOAD	EXAMPLE
Exchange	`subject:"Financial Statements"`
Exchange	`recipients:garthf@contoso.com`
SharePoint	`contenttype:document`
SharePoint	`site:https://contoso.sharepoint.com/sites/teams/procurement AND contenttype:document`
Exchange or SharePoint	`"customer information" OR "private"`

More complex examples:

The following query for SharePoint identifies Word documents or Excel spreadsheets when those files contain the keywords **password**, **passwords**, or **pw**:

```
(password OR passwords OR pw) AND (filetype:doc* OR filetype:xls*)
```

The following query for Exchange identifies any Word document or PDF that contains the word **nda** or the phrase **non disclosure agreement** when those documents are attached to an email:

```
(nda OR "non disclosure agreement") AND (attachmentnames:.doc* OR attachmentnames:.pdf)
```

The following query for SharePoint identifies documents that contain a credit card number:

```
sensitivetype:"credit card number"
```

The following query contains some typical keywords to help identify documents or emails that contain legal content:

```
ACP OR (Attorney Client Privilege*) OR (AC Privilege)
```

The following query contains typical keywords to help identify documents or emails for human resources:

```
(resume AND staff AND employee AND salary AND recruitment AND candidate)
```

Note that this final example uses the best practice of always including operators between keywords. A space between keywords (or twoproperty:valueexpressions) is the same as usingAND. By always adding operators, it's easier to see that this example query will identify only content that contains all these keywords, instead of content that contains any of the keywords. If your intention is to identify content that contains any of the keywords, specify OR instead of AND. As this example shows, when you always specify the operators, it's easier to correctly interpret the query.

Microsoft Teams meeting recordings

> **NOTE**
>
> The ability to retain and delete Teams meeting recordings is in preview and won't work before recordings are saved to OneDrive or SharePoint. For more information, see Use OneDrive for Business and SharePoint Online or Stream for meeting recordings.

To identify Microsoft Teams meeting recordings that are stored in users' OneDrive accounts or in SharePoint, specify the following for the **Keyword query editor**:

```
ProgID:Media AND ProgID:Meeting
```

Most of the time, meeting recordings are saved to OneDrive. But for channel meetings, they are saved in SharePoint.

Auto-apply labels to content by using trainable classifiers

When you choose the option for a trainable classifier, you can select one of the built-in classifiers, or a custom classifier. The built-in classifiers include **Resumes**, **SourceCode**, **Targeted Harassment**, **Profanity**, and **Threat**:

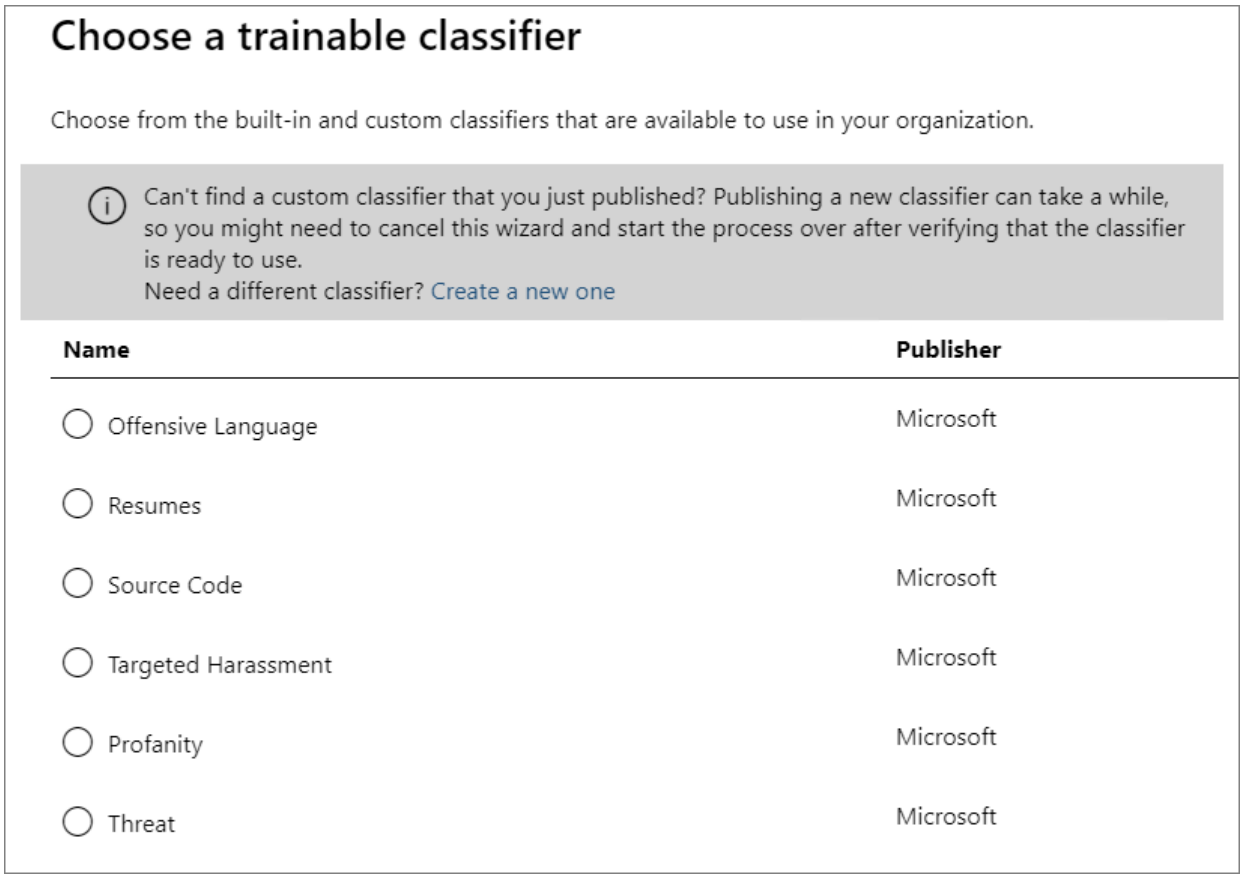

Caution

We are deprecating the **Offensive Language** built-in classifier because it has been producing a high number of false positives. Don't use this built-in classifier and if you are currently using it, you should move your business processes off it. We recommend using the **Targeted Harassment**, **Profanity**, and **Threat** built-in classifiers instead.

To automatically apply a label by using this option, SharePoint sites and mailboxes must have at least 10 MB of data.

For more information about trainable classifiers, see Learn about trainable classifiers.

TIP

If you use trainable classifiers for Exchange, see How to retrain a classifier in content explorer.

To consider when using trainable classifiers to auto-apply retention labels:

- New and modified items can be auto-labeled, and existing items from the last six months.

How long it takes for retention labels to take effect

When you auto-apply retention labels, it can take up to seven days for the retention labels to be applied to all existing content that matches the conditions.

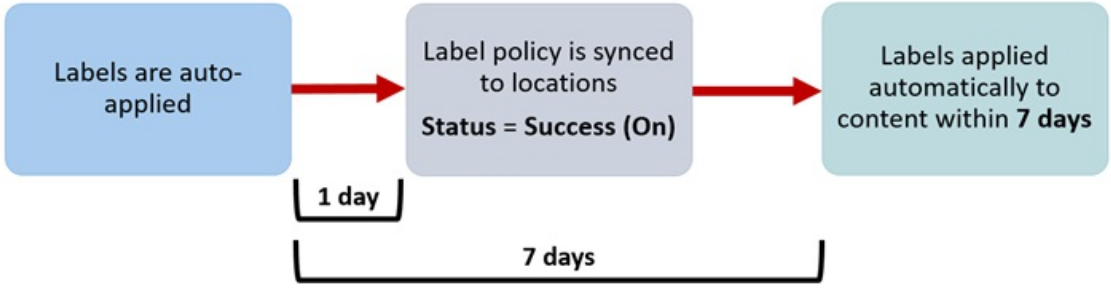

If the expected labels don't appear after seven days, check the **Status** of the auto-apply policy by selecting it from the **Label policies** page in the compliance center. If you see the status of **Off (Error)** and in the details for the locations see a message that it's taking longer than expected to deploy the policy (for SharePoint) or to try redeploying the policy (for OneDrive), try running the Set-RetentionCompliancePolicy PowerShell command to retry the policy distribution:

1. Connect to Security & Compliance Center PowerShell.
2. Run the following command:

```
Set-RetentionCompliancePolicy -Identity <policy name> -RetryDistribution
```

Updating retention labels and their policies

When you edit a retention label or auto-apply policy, and the retention label is already applied to content, your updated settings will automatically be applied to this content in addition to content that's newly identified.

Some settings can't be changed after the label or policy is created and saved, which include:

- The retention label and policy name, and the retention settings except the retention period. However, you can't change the retention period when the retention period is based on when items were labeled.
- The option to mark items as a record.

Locking the policy to prevent changes

If you need to ensure that no one can turn off the policy, delete the policy, or make it less restrictive, see Use Preservation Lock to restrict changes to retention policies and retention label policies.

Next steps

See Use retention labels to manage the lifecycle of documents stored in SharePoint for an example scenario that uses an auto-apply retention label policy with managed properties in SharePoint, and event-based retention to start the retention period.

Use Preservation Lock to restrict changes to retention policies and retention label policies

2/18/2021 • 2 minutes to read • Edit Online

Microsoft 365 licensing guidance for security & compliance.

Preservation Lock locks a retention policy or retention label policy so that no one—including a global admin—can turn off the policy, delete the policy, or make it less restrictive. This configuration might be needed for regulatory requirements and can help safeguard against rogue administrators.

When a retention policy is locked:

- No one can disable the policy or delete it
- Locations can be added but not removed
- You can extend the retention period but not decrease it

When a retention label policy is locked:

- No one can disable the policy or delete it
- Locations can be added but not removed
- Labels can be added but not removed

In summary, a locked policy can be increased or extended, but it can't be reduced or turned off.

> **IMPORTANT**
> Before you lock a retention policy or retention label policy, it's critical that you understand the impact and confirm whether it's required for your organization. For example, it might be needed to meet regulatory requirements. Administrators won't be able to disable or delete these policies after the preservation lock is applied.

Configure Preservation Lock after you've created a retention policy, or a retention label policy that you publish or auto-apply.

> **NOTE**
> Locking a label policy doesn't prevent an administrator from reducing the retention period in a label that is included in the locked policy. That requirement, with other restrictions, can be met when you configure a label to mark items as a regulatory record.

How to lock a retention policy or retention label policy

You must use PowerShell if you need to use Preservation Lock. Because administrators can't disable or delete a policy for retention after this lock is applied, enabling this feature is not available in the UI to safeguard against accidental configuration.

All policies for retention and with any configuration support Preservation Lock.

1. Connect to Security & Compliance Center PowerShell.
2. Find the name of the policy that you want to lock by running Get-RetentionCompliancePolicy. For

example:

```
PS C:\WINDOWS\system32> Get-RetentionCompliancePolicy
Name                 Workload                                                Enabled Mode
----                 --------                                                ------- ----
PII Retention Policy Exchange, SharePoint, OneDriveForBusiness, Skype, ModernGroup True    Enforce
Employee Records     Exchange, SharePoint, OneDriveForBusiness, Skype, ModernGroup True    Enforce
Personal Financial PII Exchange, SharePoint, OneDriveForBusiness, Skype, ModernGroup True    Enforce
```

3. To place a Preservation Lock on your policy, run the Set-RetentionCompliancePolicy cmdlet with the name of the policy, and the *RestrictiveRetention* parameter set to true:

   ```
   Set-RetentionCompliancePolicy -Identity "<Name of Policy>" -RestrictiveRetention $true
   ```

   For example:

   ```
   PS C:\WINDOWS\system32> Set-RetentionCompliancePolicy -Identity "Employee Records" -RestrictiveRetention $true
   ```

   When prompted, read and acknowledge the restrictions that come with this configuration by entering **Y**:

   ```
   Confirm
   You chose to make this a restrictive preservation policy. After you turn this policy on, you'll only be able to
   add new locations to search or extend the time frame to preserve the content. You won't be able to turn off the
   policy, delete it, or make any other changes.
   [Y] Yes  [A] Yes to All  [N] No  [L] No to All  [?] Help (default is "Y"):
   ```

A Preservation Lock is now placed on the policy. To confirm, run `Get-RetentionCompliancePolicy` again, but specify the policy name and display the policy parameters:

```
Get-RetentionCompliancePolicy -Identity "<Name of Policy>" |Fl
```

You should see **RestrictiveRetention** is set to **True**. For example:

```
PS C:\WINDOWS\system32> Get-RetentionCompliancePolicy -Identity "Employee Records" |Fl


RunspaceId                    : d2b39144-2e5d-4dea-ab59-2786944f9ff3
Type                          : Hold
TeamsPolicy                   : False
SharePointLocation            : {}
SharePointLocationException   : {}
RetentionRuleTypes            : {}
ExchangeLocation              : {}
ExchangeLocationException     : {}
PublicFolderLocation          : {}
SkypeLocation                 : {}
SkypeLocationException        : {}
ModernGroupLocation           : {}
ModernGroupLocationException  : {}
OneDriveLocation              : {}
OneDriveLocationException     : {}
TeamsChatLocation             : {}
TeamsChatLocationException    : {}
TeamsChannelLocation          : {}
TeamsChannelLocationException : {}
DynamicScopeLocation          : {}
RestrictiveRetention          : True
Workload                      : Exchange, SharePoint, OneDriveForBusiness, Skype, ModernGroup
Priority                      : 1
ObjectVersion                 : 533c3975-eae4-4520-aa83-08d83ee404f2
```

See also

Resources to help you meet regulatory requirements for information governance and records management

Create and publish retention labels by using PowerShell

11/2/2020 • 13 minutes to read • Edit Online

> *Microsoft 365 licensing guidance for security & compliance.*

After you've decided to use retention labels to help you keep or delete documents and emails in Microsoft 365, you might have realized that you have many and possibly hundreds of retention labels to create and publish. The recommended method to create retention labels at scale is by using file plan from the Microsoft 365 compliance center. However, you can also use PowerShell.

Use the information, template files and examples, and script in this article to help you bulk-create retention labels and publish them in retention label policies. Then, the retention labels can be applied by administrators and users.

The supplied instructions don't support retention labels that are auto-applied.

Overview:

1. In Excel, create a list of your retention labels and a list of their retention label policies.
2. Use PowerShell to create the retention labels and retention label policies in those lists.

Disclaimer

The sample scripts provided in this article aren't supported under any Microsoft standard support program or service. The sample scripts are provided AS IS without warranty of any kind. Microsoft further disclaims all implied warranties including, without limitation, any implied warranties of merchantability or of fitness for a particular purpose. The entire risk arising out of the use or performance of the sample scripts and documentation remains with you. In no event shall Microsoft, its authors, or anyone else involved in the creation, production, or delivery of the scripts be liable for any damages whatsoever (including, without limitation, damages for loss of business profits, business interruption, loss of business information, or other pecuniary loss) arising out of the use of or inability to use the sample scripts or documentation, even if Microsoft has been advised of the possibility of such damages.

Step 1: Create a .csv file for the retention labels

1. Copy the following sample .csv file for a template and example entries for four different retention labels, and paste them into Excel.
2. Convert the text to columns: **Data** tab > **Text to Columns** > **Delimited** > **Comma** > **General**
3. Replace the examples with entries for your own retention labels and settings. For more information about the parameter values, see New-ComplianceTag.
4. Save the worksheet as a .csv file in a location that's easy to find for a later step. For example: C:>Scripts\Labels.csv

Notes:

- If the .csv file contains a retention label with the same name as one that already exists, the script skips creating that retention label. No duplicate retention labels are created.

- Don't change or rename the column headers from the sample .csv file provided, or the script will fail.

Sample .csv file for retention labels

```
Name (Required),Comment (Optional),IsRecordLabel (Required),RetentionAction (Optional),RetentionDuration
(Optional),RetentionType (Optional),ReviewerEmail (Optional)
LabelName_t_1,Record - keep and delete - 2 years,$true,KeepAndDelete,730,CreationAgeInDays,
LabelName_t_2,Keep and delete tag - 7 years,$false,KeepAndDelete,2555,ModificationAgeInDays,
LabelName_t_3,5 year delete,$false,Delete,1825,TaggedAgeInDays,
LabelName_t_4,Record label tag - financial,$true,Keep,730,CreationAgeInDays,
```

Step 2: Create a .csv file for the retention label policies

1. Copy the following sample .csv file for a template and example entries for three different retention label policies, and paste them into Excel.
2. Convert the text to columns: **Data** tab > **Text to Columns** > **Delimited** > **Comma** > **General**
3. Replace the examples with entries for your own retention label policies and their settings. For more information about the parameter values for this cmdlet, see New-RetentionCompliancePolicy.
4. Save the worksheet as a .csv file in a location that's easy to find for a later step. For example: `<path>Policies.csv`

Notes:

- If the .csv file contains a retention label policy with the same name as one that already exists, the script skips creating that retention label policy. No duplicate retention label policies are created.
- Don't change or rename the column headers from the sample .csv file provided, or the script will fail.

Sample .csv file for retention policies

```
Policy Name (Required),PublishComplianceTag (Required),Comment (Optional),Enabled
(Required),ExchangeLocation (Optional),ExchangeLocationException (Optional),ModernGroupLocation
(Optional),ModernGroupLocationException (Optional),OneDriveLocation (Optional),OneDriveLocationException
(Optional),PublicFolderLocation (Optional),SharePointLocation (Optional),SharePointLocationException
(Optional),SkypeLocation (Optional),SkypeLocationException (Optional)
Publishing Policy Red1,"LabelName_t_1, LabelName_t_2, LabelName_t_3,
LabelName_t_4",N/A,$true,All,,All,,All,,,All,,,
Publishing Policy Orange1,"LabelName_t_1, LabelName_t_2",N/A,$true,All,,,,,,,,,
Publishing Policy Yellow1,"LabelName_t_3, LabelName_t_4",N/A,$false,All,,,,,,,,,
```

Step 3: Create the PowerShell script

1. Copy and paste the following PowerShell script into Notepad.
2. Save the file by using a file name extension of **.ps1** in a location that's easy to find. For example: `<path>CreateRetentionSchedule.ps1`

Notes:

- The script prompts you to provide the two source files that you created in the previous two steps:
  - If you don't specify the source file to create the retention labels, the script moves on to create the retention label policies.
  - If you don't specify the source file to create the retention label policies, the script creates the retention labels only.
- The script generates a log file that records each action it took and whether the action succeeded or failed.

See the final step for instructions how to locate this log file.

PowerShell script

```
<#
. Steps: Import and publish retention labels
    o Load retention labels csv file
    o Validate csv file input
    o Create retention labels
    o Create retention policies
    o Publish retention labels for the policies
    o Generate the log for retention labels and policies creation
    o Generate the csv result for the labels and policies created
. Syntax
    .\Publish-ComplianceTag.ps1 [-LabelListCSV <string>] [-PolicyListCSV <string>]
. Detailed Description
    1) [-LabelListCSV <string>]
    -LabelListCSV ".\SampleInputFile_LabelList.csv"
    Load compliance tag for creation.
    2) [-PolicyListCSV <string>]
    -PolicyListCSV ".\SampleInputFile_PolicyList.csv"
    Load compliance tag for creation.
#>
param (
    [Parameter(Mandatory = $true)]
    [string]$LabelListCSV = "",
    [Parameter(Mandatory = $true)]
    [string]$PolicyListCSV = "",
    [Switch]$ResultCSV
)
# -------------------
# File operation
# -------------------
Function FileExist
{
    Param(
        # File path needed to check
        [Parameter(Mandatory = $true)]
        [String]$FilePath,
        [Switch]$Warning
    )
    $inputFileExist = Test-Path $FilePath
    if (!$inputFileExist)
    {
        if ($Warning -eq $false)
        {
            WriteToLog -Type "Failed" -Message "[File: $FilePath] The file doesn't exist"
            throw
        }
        else
        {
            WriteToLog -Type "Warning" -Message "[File: $FilePath] The file doesn't exist"
        }
    }
    else
    {
        WriteToLog -Type "Succeed" -Message "[File: $FilePath] The file is found"
    }
}
# -------------------
# Log operation
# -------------------
Function WriteToLog
{
    Param(
        # Message want to write to log file
        [Parameter(Mandatory = $true)]
        [String]$Message,
```

```
        # "Succeed" or "Faild"
        [String]$Type = "Message"
    )
    $date = Get-Date -Format 'HH:mm:ss'
    $logInfo = $date + " - [$Type] " + $Message
    $logInfo | Out-File -FilePath $logfilePath -Append
    if ($Type -eq "Succeed") { Write-Host $logInfo -ForegroundColor Green }
    elseif ($Type -eq "Failed") { Write-Host $logInfo -ForegroundColor Red }
    elseif ($Type -eq "Warning") { Write-Host $logInfo -ForegroundColor Yellow }
    elseif ($Type -eq "Start") { Write-Host $logInfo -ForegroundColor Cyan }
    else { Write-Verbose $logInfo }
}
Function Create-Log
{
    Param(
        # Log folder Root
        [Parameter(Mandatory = $true)]
        [String]$LogFolderRoot,
        # The function Log file for
        [Parameter(Mandatory = $true)]
        [String]$LogFunction
    )
    $logFolderPath = "$LogFolderRoot\logfiles"
    $folderExist = Test-Path "$logFolderPath"
    if (!$folderExist)
    {
        $folder = New-Item "$logFolderPath" -type directory
    }
    $date = Get-Date -Format 'MMddyyyy_HHmmss'
    $logfilePath = "$logFolderPath\Log_{0}_{1}.txt" -f $LogFunction, $date
    Write-Verbose "Log file is written to: $logfilePath"
    $logfile = New-Item $logfilePath  -type file
    return $logfilePath
}
Function Create-ResultCSV
{
    Param(
        # Result folder Root
        [Parameter(Mandatory = $true)]
        [String]$ResultFolderRoot,
        # The function Result file for
        [Parameter(Mandatory = $true)]
        [String]$ResultFunction
    )
    $retFolderPath = "$ResultFolderRoot\logfiles"
    $folderExist = Test-Path "$retFolderPath"
    if (!$folderExist)
    {
        $folder = New-Item "$retFolderPath" -type directory
    }
    $date = Get-Date -Format 'MMddyyyy_HHmmss'
    $retfilePath = "$retFolderPath\Result_{0}_{1}.csv" -f $ResultFunction, $date
    Write-Verbose "Result file is written to: $retfilePath"
    $retfile = New-Item $retfilePath  -type file
    return $retfilePath
}
# --------------------
# Prepare Log File
# --------------------
$scriptPath = '.\'
$logfilePath = Create-Log -LogFolderRoot $scriptPath -LogFunction "Publish_Compliance_Tag"
if ($ResultCSV)
{
    $tagRetFile = Create-ResultCSV -ResultFolderRoot $scriptPath -ResultFunction "Tag_Creation"
    $tagPubRetFile = Create-ResultCSV -ResultFolderRoot $scriptPath -ResultFunction "Tag_Publish"
}
# --------------------
# Invoke Powershell cmdlet
# --------------------
```

```
Function InvokePowerShellCmdlet
{
    Param(
        [Parameter(Mandatory = $true)]
        [String]$CmdLet
    )
    try
    {
        WriteToLog -Type "Start" -Message "Execute Cmdlet : '$CmdLet'"
        return Invoke-Expression $CmdLet -ErrorAction SilentlyContinue
    }
    catch
    {
        WriteToLog -Type "Failed" "Failed to execute cmdlet!"
        WriteToLog -Type "Failed" $error[0]
        return $null
    }
}
# -------------------
# Create Compliance Tag
# -------------------
Function CreateComplianceTag
{
    Param(
        # File path needed to check
        [Parameter(Mandatory = $true)]
        [String]$FilePath
    )

    WriteToLog -Type "Start" "Start to create Compliance Tag"
    FileExist $FilePath

    # TODO Validate CSV file for the Header
    try
    {
        # Import csv
        $labels = Import-Csv $FilePath
        # Retrieve existing compliance tags
        $tags = InvokePowerShellCmdlet "Get-ComplianceTag"
        foreach($lab in $labels)
        {
            # Cmdlet parameters
            $para = [String]::Empty;
            $name = [String]::Empty;
            $cmdlet = 'New-ComplianceTag'
            if ([String]::IsNullOrEmpty($lab.'Name (Required)'))
            {
                WriteToLog -Type "Failed" -Message "Could not acquire table for writing."
                throw;
            }
            else
            {
                $name = $lab.'Name (Required)'
                $cmdlet += " -Name '" + $name + "'"
            }
            if (![String]::IsNullOrEmpty($lab.'Comment (Optional)'))
            {
                $para = $lab.'Comment (Optional)'
                $cmdlet += " -Comment '" + $para + "'"
            }
            if (![String]::IsNullOrEmpty($lab.'IsRecordLabel (Required)'))
            {
                $para = $lab.'IsRecordLabel (Required)'
                $cmdlet += " -IsRecordLabel " + $para
            }
            if (![String]::IsNullOrEmpty($lab.'RetentionAction (Optional)'))
            {
                $para = $lab.'RetentionAction (Optional)'
                $cmdlet += " -RetentionAction " + $para
```

```
            }
            if (![String]::IsNullOrEmpty($lab.'RetentionDuration (Optional)'))
            {
                $para = $lab.'RetentionDuration (Optional)'
                $cmdlet += " -RetentionDuration " + $para
            }
            if (![String]::IsNullOrEmpty($lab.'RetentionType (Optional)'))
            {
                $para = $lab.'RetentionType (Optional)'
                $cmdlet += " -RetentionType " + $para
            }
            if (![String]::IsNullOrEmpty($lab.'ReviewerEmail (Optional)'))
            {
                $emails = $lab.'ReviewerEmail (Optional)'.Split(",") | ForEach-Object { $_.Trim() }
                if (($emails -ne $null) -and ($emails.Count -ne 0))
                {
                    $eml = '@('
                    foreach($email in $emails)
                    {
                        $eml += "'{0}'," -f $email
                    }
                    $eml = $eml.Substring(0, $eml.Length - 1) + ')'

                    $cmdlet += " -ReviewerEmail " + $eml
                }
            }
            # If the tag already exists, skip for creation
            if (($tags -eq $null) -or ($tags | ? { $_.Name.ToLower() -eq $name.ToLower() }) -eq $null)
            {
                # Create compliance tag
                $msg = "Execute Cmdlet : {0}" -f $cmdlet

                $ret = InvokePowerShellCmdlet $cmdlet

                if ($ret -eq $null)
                {
                    WriteToLog -Type "Failed" $error[0]
                    break;
                }
            }
            else
            {
                WriteToLog -Type "Warning" -Message "The tag '$name' already exists! Skip for creation!"
            }
        }
    }
    catch
    {
        WriteToLog -Type "Failed" "Error in input"
    }
}
# -------------------
# Create Retention Compliance Policy
# -------------------
Function CreateRetentionCompliancePolicy
{
    Param(
        # File path needed to check
        [Parameter(Mandatory = $true)]
        [String]$FilePath
    )

    WriteToLog -Type "Start" "Start to Create Retention Policy"
    FileExist $FilePath
    try
    {
        # Import csv
        $list = Import-Csv -Path $FilePath
        # Retrieve existing retention compliance policy
```

```
# Retrieve existing retention compliance policy
$policies = InvokePowerShellCmdlet "Get-RetentionCompliancePolicy"
foreach($rp in $list)
{
    # Cmdlet parameters
    $para = [String]::Empty;
    $name = [String]::Empty;
    $rpid = [String]::Empty;
    $cmdlet = 'New-RetentionCompliancePolicy'
    if ([String]::IsNullOrEmpty($rp.'Policy Name (Required)'))
    {
        WriteToLog -Type "Failed" -Message "Could not acquire table for writing."
        throw;
    }
    else
    {
       $name = $rp.'Policy Name (Required)'
       $cmdlet += " -Name '" + $name + "'"
    }
    if ([String]::IsNullOrEmpty($rp.'Enabled (Required)'))
    {
        WriteToLog -Type "Failed" -Message "Could not acquire table for writing."
        throw;
    }
    else
    {
        $enabled = $rp.'Enabled (Required)'
        $cmdlet += " -Enabled " + $enabled
    }
    if (![String]::IsNullOrEmpty($rp.'ExchangeLocation (Optional)'))
    {
        $para = $rp.'ExchangeLocation (Optional)'
        $cmdlet += " -ExchangeLocation " + $para
    }

    if (![String]::IsNullOrEmpty($rp.'ExchangeLocationException (Optional)'))
    {
        $para = $rp.'ExchangeLocationException (Optional)'
        $cmdlet += " -ExchangeLocationException " + $para
    }
    if (![String]::IsNullOrEmpty($rp.'ModernGroupLocation (Optional)'))
    {
        $para = $rp.'ModernGroupLocation (Optional)'
        $cmdlet += " -ModernGroupLocation " + $para
    }
    if (![String]::IsNullOrEmpty($rp.'ModernGroupLocationException (Optional)'))
    {
        $para = $rp.'ModernGroupLocationException (Optional)'
        $cmdlet += " -ModernGroupLocationException " + $para
    }
    if (![String]::IsNullOrEmpty($rp.'OneDriveLocation (Optional)'))
    {
        $para = $rp.'OneDriveLocation (Optional)'
        $cmdlet += " -OneDriveLocation " + $para
    }
    if (![String]::IsNullOrEmpty($rp.'OneDriveLocationException (Optional)'))
    {
        $para = $rp.'OneDriveLocationException (Optional)'
        $cmdlet += " -OneDriveLocationException " + $para
    }
    if (![String]::IsNullOrEmpty($rp.'SharePointLocation (Optional)'))
    {
        $para = $rp.'SharePointLocation (Optional)'
        $cmdlet += " -SharePointLocation " + $para
    }
    if (![String]::IsNullOrEmpty($rp.'SharePointLocationException (Optional)'))
    {
        $para = $rp.'SharePointLocationException (Optional)'
        $cmdlet += " -SharePointLocationException " + $para
    }
```

```
            }
            if (![String]::IsNullOrEmpty($rp.'PublicFolderLocation (Optional)'))
            {
                $para = $rp.'PublicFolderLocation (Optional)'
                $cmdlet += " -PublicFolderLocation " + $para
            }
            if (![String]::IsNullOrEmpty($rp.'SkypeLocation (Optional)'))
            {
                $para = $rp.'SkypeLocation (Optional)'
                $cmdlet += " -SkypeLocation " + $para
            }
            if (![String]::IsNullOrEmpty($rp.'SkypeLocationException (Optional)'))
            {
                $para = $rp.'SkypeLocationException (Optional)'
                $cmdlet += " -SkypeLocationException " + $para
            }
            # If the policy already exists, skip for creation
            if (($policies -eq $null) -or ($policies | ? { $_.Name.ToLower() -eq $name.ToLower() }) -eq
$null)
            {
                # Create retention compliance policy
                $msg = "Execute Cmdlet : {0}" -f $cmdlet

                $ret = invokepowershellcmdlet $cmdlet

                if ($ret -eq $null)
                {
                    WriteToLog -Type "Failed" $error[0]
                    break;
                }
                $rpid = $ret.Guid
            }
            else
            {
                WriteToLog -Type "Warning" -Message "The policy '$name' already exists! Skip for creation!"
                $rpid = ($policies | ? { $_.Name.ToLower() -eq $name.ToLower() }).Guid
            }

            # Retrieve tag name for publishing
            $ts = $rp.'PublishComplianceTag (Required)'
            $tagList = $ts.Split(",") | ForEach-Object { $_.Trim() }

            WriteToLog -Type "Message" -Message "Publish Tags : '$ts'"

            PublishComplianceTag -PolicyGuid $rpid -TagName $tagList
        }
    }
    catch
    {
        WriteToLog -Type "Failed" "Error in input"
    }
}
# -------------------
# Publish Compliance Tag
# -------------------
Function PublishComplianceTag
{
    Param(
        [Parameter(Mandatory = $true)]
        [String]$PolicyGuid,
        [Parameter(Mandatory = $true)]
        [String[]]$TagNames
    )

    WriteToLog -Type "Start" "Start to Publish Compliance Tag"
    try
    {
        # Retrieve existing rule related to the given compliance policy
        $rule = InvokePowerShellCmdlet ("Get-RetentionComplianceRule -Policy {0}" -f $PolicyGuid)
```

```
        $tagGuids = New-Object System.Collections.ArrayList

        foreach ($tn in $TagNames)
        {
            $t = InvokePowerShellCmdlet ("Get-ComplianceTag {0}" -f $tn)
            $tagGuids.Add($t.Guid) | Out-Null
        }
        if ($rule -ne $null)
        {
            foreach ($r in $rule)
            {
                if ([String]::IsNullOrEmpty($r.PublishComplianceTag))
                {
                    continue;
                }
                else
                {
                    $tl = $r.PublishComplianceTag.Split(",")
                    if ($tagGuids.Contains([GUID]$tl[0]))
                    {
                        $tagGuids.Remove([GUID]$tl[0]);
                    }
                }
            }
        }

        foreach($t in $tagGuids)
        {
            # Publish compliance tag
            $cmdlet = "New-RetentionComplianceRule -Policy {0} -PublishComplianceTag {1}" -f $PolicyGuid, $t
            $ret = InvokePowerShellCmdlet $cmdlet

            if ($ret -eq $null)
            {
                WriteToLog -Type "Failed" $error[0]
                break;
            }
        }
    }
    catch
    {
        WriteToLog -Type "Failed" "Error in input"
    }
}
# -------------------
# Export All Labels Created in The Process
# -------------------
Function ExportCreatedComplianceTag
{
    Param(
        [Parameter(Mandatory = $true)]
        [String]$LabelFilePath
    )

    WriteToLog -Type "Start" "Start to Export Compliance Tag Created"
    try
    {
        # Import input csv
        $labels = Import-Csv $LabelFilePath
        # Create result table
        $tabName = "ResultTable"
        $table = New-Object system.Data.DataTable "$tabName"
        $col1 = New-Object system.Data.DataColumn Name,([string])
        $col2 = New-Object system.Data.DataColumn Comment,([string])
        $col3 = New-Object system.Data.DataColumn IsRecordLabel,([string])
        $col4 = New-Object system.Data.DataColumn RetentionAction,([string])
        $col5 = New-Object system.Data.DataColumn RetentionDuration,([string])
        $col6 = New-Object system.Data.DataColumn RetentionType,([string])
        $col7 = New-Object system.Data.DataColumn ReviewerEmail,([string])
```

```
        # Add the Columns
        $table.columns.add($col1)
        $table.columns.add($col2)
        $table.columns.add($col3)
        $table.columns.add($col4)
        $table.columns.add($col5)
        $table.columns.add($col6)
        $table.columns.add($col7)
        foreach($lab in $labels)
        {
            $t = InvokePowerShellCmdlet ("Get-ComplianceTag '{0}' " -f $lab.'Name (Required)')

            # Create a result row
            $row = $table.NewRow()
            $row['Name'] = $t.Name
            $row['Comment'] = $t.Comment
            $row['IsRecordLabel'] = $t.IsRecordLabel
            $row['RetentionAction'] = $t.RetentionAction
            $row['RetentionDuration'] = $t.RetentionDuration
            $row['RetentionType'] = $t.RetentionType
            $row['ReviewerEmail'] = $t.ReviewerEmail

            # Add the row to the table
            $table.Rows.Add($row)
        }
        $table | Export-Csv $tagRetFile -NoTypeInformation
    }
    catch
    {
        WriteToLog -Type "Failed" "Error in exporting results."
    }
}
# -------------------
# Export All Published Labels and Policies in The Process
# -------------------
Function ExportPublishedComplianceTagAndPolicy
{
    Param(
        [Parameter(Mandatory = $true)]
        [String[]]$PolicyFilePath
    )

    WriteToLog -Type "Start" "Start to Export Published Compliance Tag and Policy"
    try
    {
        # Import input csv
        $policies = Import-Csv $PolicyFilePath
        # Create result table
        $tabName = "ResultTable"
        $table = New-Object system.Data.DataTable "$tabName"
        $col1 = New-Object system.Data.DataColumn 'Policy Name',([string])
        $col2 = New-Object system.Data.DataColumn PublishComplianceTag,([string])
        $col3 = New-Object system.Data.DataColumn Comment,([string])
        $col4 = New-Object system.Data.DataColumn Enabled,([string])
        $col5 = New-Object system.Data.DataColumn ExchangeLocation,([string])
        $col6 = New-Object system.Data.DataColumn ExchangeLocationException,([string])
        $col7 = New-Object system.Data.DataColumn ModernGroupLocation,([string])
        $col8 = New-Object system.Data.DataColumn ModernGroupLocationException,([string])
        $col9 = New-Object system.Data.DataColumn OneDriveLocation,([string])
        $col10 = New-Object system.Data.DataColumn OneDriveLocationException,([string])
        $col11 = New-Object system.Data.DataColumn PublicFolderLocation,([string])
        $col12 = New-Object system.Data.DataColumn SharePointLocation,([string])
        $col13 = New-Object system.Data.DataColumn SharePointLocationException,([string])
        $col14 = New-Object system.Data.DataColumn SkypeLocation,([string])
        $col15 = New-Object system.Data.DataColumn SkypeLocationException,([string])

        # Add the Columns
        $table.columns.add($col1)
```

```
        $table.columns.add($col2)
        $table.columns.add($col3)
        $table.columns.add($col4)
        $table.columns.add($col5)
        $table.columns.add($col6)
        $table.columns.add($col7)
        $table.columns.add($col8)
        $table.columns.add($col9)
        $table.columns.add($col10)
        $table.columns.add($col11)
        $table.columns.add($col12)
        $table.columns.add($col13)
        $table.columns.add($col14)
        $table.columns.add($col15)
        foreach($policy in $policies)
        {
            $t = InvokePowerShellCmdlet ("Get-RetentionCompliancePolicy '{0}' -DistributionDetail" -f
$policy.'Policy Name (Required)')

            # Create a result row
            $row = $table.NewRow()
            $row['Policy Name'] = $t.Name

            $rules = InvokePowerShellCmdlet ("Get-RetentionComplianceRule -Policy {0}" -f $t.Guid)
            $tagList = [String]::Empty
            foreach($rule in $rules)
            {
                if ([String]::IsNullOrEmpty($rule.PublishComplianceTag) -eq $False)
                {
                    $tName = $rule.PublishComplianceTag.Split(',')[1]
                    $tagList = [String]::Concat($tagList, $tName, ",")
                }
            }
            if (![String]::IsNullOrEmpty($tagList))
            {
                $tagList = $tagList.Substring(0, $tagList.LastIndexOf(','))
            }
            $row['PublishComplianceTag'] = $tagList
            $row['Comment'] = $t.Comment
            $row['Enabled'] = $t.Enabled
            $row['ExchangeLocation'] = $t.ExchangeLocation
            $row['ExchangeLocationException'] = $t.ExchangeLocationException
            $row['ModernGroupLocation'] = $t.ModernGroupLocation
            $row['ModernGroupLocationException'] = $t.ModernGroupLocationException
            $row['OneDriveLocation'] = $t.OneDriveLocation
            $row['OneDriveLocationException'] = $t.OneDriveLocationException
            $row['PublicFolderLocation'] = $t.PublicFolderLocation
            $row['SharePointLocation'] = $t.SharePointLocation
            $row['SharePointLocationException'] = $t.SharePointLocationException
            $row['SkypeLocation'] = $t.SkypeLocation
            $row['SkypeLocationException'] = $t.SkypeLocationException

            # Add the row to the table
            $table.Rows.Add($row)
        }
        $table | Export-Csv $tagPubRetFile -NoTypeInformation
    }
    catch
    {
        WriteToLog -Type "Failed" "Error in exporting results."
    }
}
# Create compliance tag
CreateComplianceTag -FilePath $LabelListCSV
# Create retention policy and publish compliance tag with the policy
CreateRetentionCompliancePolicy -FilePath $PolicyListCSV
# Export to result csv
if ($ResultCSV)
{
```

```
    ExportCreatedComplianceTag -LabelFilePath $LabelListCSV
    ExportPublishedComplianceTagAndPolicy -PolicyFilePath $PolicyListCSV
}
```

Step 4: Run the PowerShell script

First, Connect to Security & Compliance Center PowerShell.

Then, run the script that creates and publishes the retention labels:

1. In your Security & Compliance Center PowerShell session, enter the path, followed by the characters `.\` and the file name of the script, and then press ENTER to run the script. For example:

   ```
   <path>.\CreateRetentionSchedule.ps1
   ```

2. The script prompts you for the locations of the .csv files that you created in the previous steps. Enter the path, followed by the characters `.\` and file name of the .csv file, and then press ENTER. For example, for the first prompt:

   ```
   <path>.\Labels.csv
   ```

Step 5: View the log file with the results

Use the log file that the script created to check the results and identify any failures that need resolving.

You can find the log file at the following location, although the digits in the example file name vary.

```
<path>.\Log_Publish_Compliance_Tag_01112018_151239.txt
```

Start retention when an event occurs

2/18/2021 • 13 minutes to read • Edit Online

Microsoft 365 licensing guidance for security & compliance.

When you retain content, the retention period is often based on the age of the content. For example, you might retain documents for seven years after they're created and then delete them. But when you configure retention labels, you can also base a retention period on when a specific type of event occurs. The event triggers the start of the retention period, and all content with a retention label applied for that type of event get the label's retention actions enforced on them.

Examples for using event-based retention:

- **Employees leaving the organization** Suppose that employee records must be retained for 10 years from the time an employee leaves the organization. After 10 years elapse, all documents related to the hiring, performance, and termination of that employee must be disposed. The event that triggers the 10-year retention period is the employee leaving the organization.
- **Contract expiration** Suppose that all records related to contracts must be retained for five years from the time the contract expires. The event that triggers the five-year retention period is the expiration of the contract.
- **Product lifetime** Your organization might have retention requirements related to the last manufacturing date of products for content such as technical specifications. In this case, the last manufacturing date is the event that triggers the retention period.

Event-based retention is typically used as part of a records-management process. This means that:

- Retention labels based on events also usually mark items as a record, as a part of a records management solution. For more information, see Learn about records management.
- A document that's been declared a record but whose event trigger has not yet happened is retained indefinitely (records can't be permanently deleted), until an event triggers that document's retention period.
- Retention labels based on events usually trigger a disposition review at the end of the retention period, so that a records manager can manually review and dispose of the content. For more information, see Disposition of content.

A retention label based on an event has the same capabilities as any retention label in Microsoft 365. For more information, see Learn about retention policies and retention labels.

Understanding the relationship between event types, labels, events, and asset IDs

To successfully use event-based retention, it's important to understand the relationship between event types, retention labels, events, and asset IDs as illustrated in the diagrams and the explanation that follows:

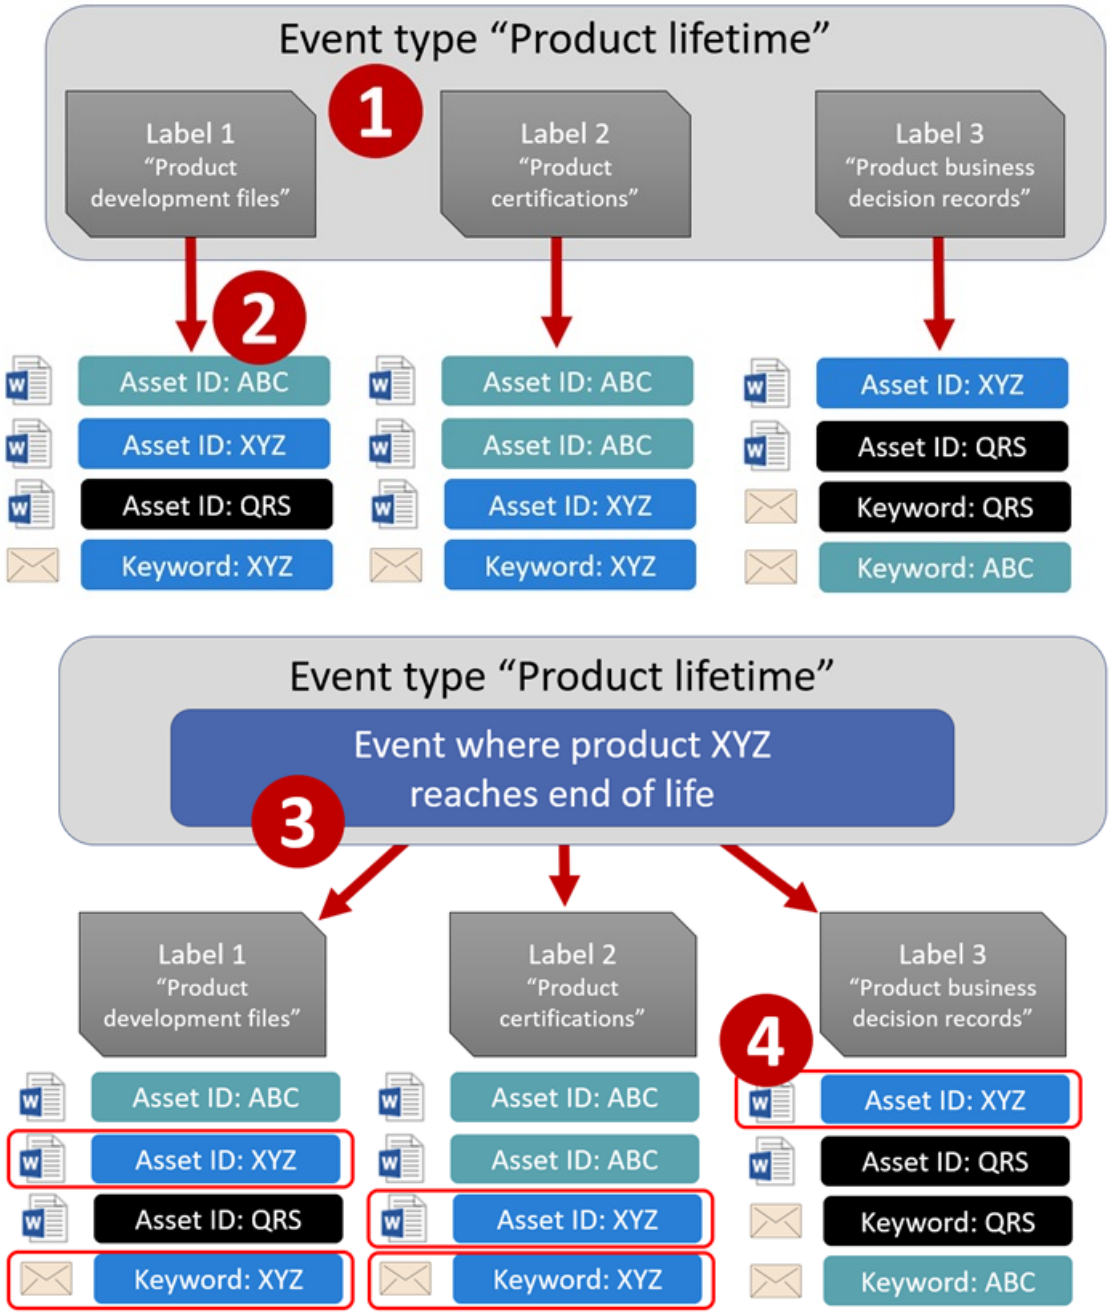

1. You create retention labels for different types of content and then associate them with a type of event. For example, retention labels for different types of product files and records are associated with an event type named Product Lifetime because those records must be retained for 10 years from the time the product reaches its end of life.

2. Users (typically records managers) apply those retention labels to content and (for documents in SharePoint and OneDrive) enter an asset ID for each item. In this example, the asset ID is a product name or code used by the organization. Then, each product's records are assigned a retention label, and each record has a property that contains an asset ID. The diagram represents **all the content** for all product records in an organization, and each item bears the asset ID of the product whose record it is.

3. Product Lifetime is the event type; a specific product reaching end of life is an event. When an event of that event type occurs—in this case, when a product reaches its end of life—you create an event that specifies:

   - An asset ID (for SharePoint and OneDrive documents)
   - Keywords (for Exchange items). In this example, the organization uses a product code in messages containing product records, so the keyword for Exchange items is functionally the same as the asset ID for SharePoint and OneDrive documents.

- The date when the event occurred. This date is used as the start of the retention period. This date can be the current, a past, or a future date.

4. After you create an event, that event date is synchronized to all the content that has a retention label of that event type and that contains the specified asset ID or keyword. Like any retention label, this synchronization can take up to seven days. In the previous diagram, all the items circled in red have their retention period triggered by this event. In other words, when this product reaches its end of life, that event triggers the retention period for that product's records.

It's important to understand that if you don't specify an asset ID or keywords for an event, **all content** with a retention label of that event type will have its retention period triggered by the event. This means that in the previous diagram, all content would start being retained. This might not be what you intend.

Finally, remember that each retention label has its own retention settings. In this example, they all specify 10 years, but it's possible for an event to trigger retention labels where each label has a different retention period.

How to set up event-driven retention

High-level workflow for event-driven retention:

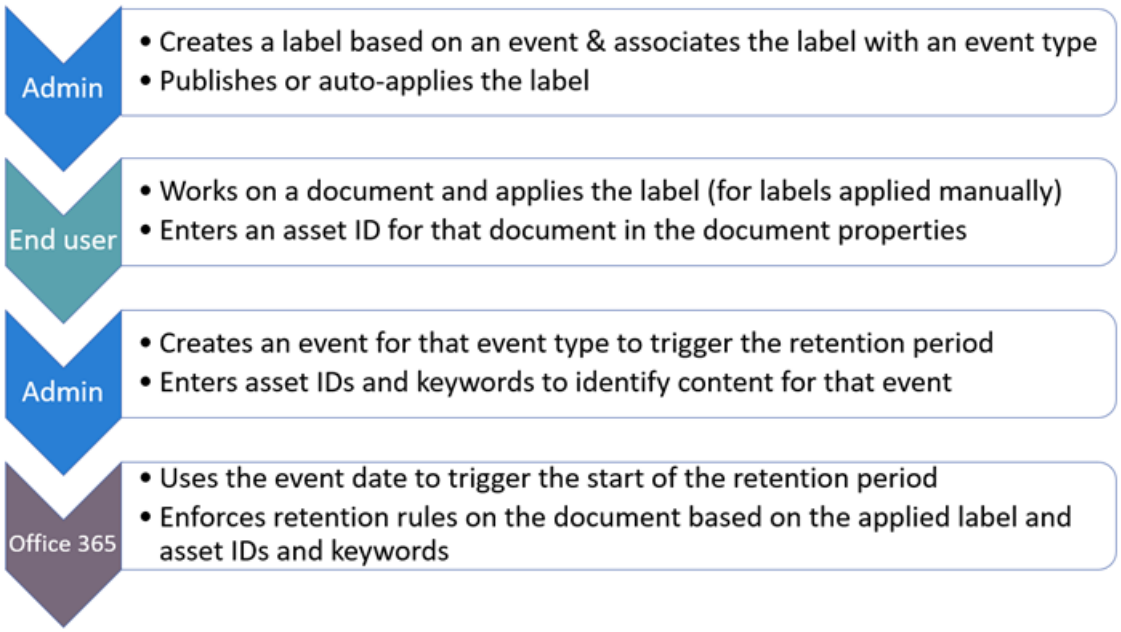

TIP

See Use retention labels to manage the lifecycle of documents stored in SharePoint for a detailed scenario about using managed properties in SharePoint to auto-apply retention labels and implement event-driven retention.

Step 1: Create a label whose retention period is based on an event

To create and configure your retention label, use the instructions from Create and configure retention labels. But specific to event-based retention, on the **Define retention settings** page of the Create retention label wizard, after **Start the retention period based on**, select one of the default event types from the dropdown list, or create your own by selecting **Create new event type**:

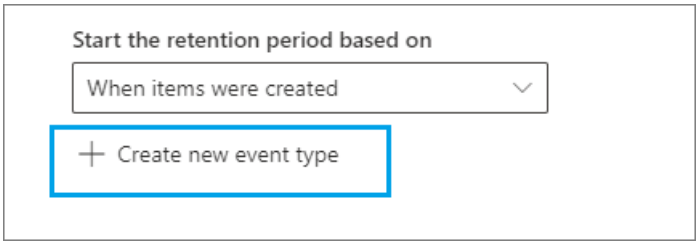

An event type is simply a general description of an event that you want to associate with a retention label.

The default event types have **(event type)** after their name in the dropdown list for easier identification, and you can also see and create event type from the **Records management** > **Events** tab > **Manage event types**.

Event-based retention requires retention settings that:

- Retain the content.
- Delete the content automatically or trigger a disposition review at the end of the retention period.

Event-based retention is typically used for content that's declared a record, so this is a good time to check whether you also need to select the option that marks content as a record.

If you're using an existing event type rather than creating a new event type, skip to step 3.

> **NOTE**
> After you choose an event type and save the retention label, the event type cannot be changed.

Step 2: Create a new event type for your label

For the retention settings, if you selected **Create new event type**, enter a name and description for your event type. Then select **Next**, **Submit**, and **Done**.

Back on the **Define retention settings** page, for **Start the retention period based on**, use the dropdown list to select the event type that you created.

Step 3: Publish or auto-apply the event-based retention labels

Just like any retention label, you need to publish or auto-apply an event-based label, for it to be manually or automatically applied to content:

- Create retention labels and apply them in apps
- Apply a retention label to content automatically

Step 4: Enter an asset ID

After an event-based label is applied to content, you can enter an asset ID for each item. For example, your organization might use:

- Product codes that you can use to retain content for only a specific product.
- Project codes that you can use to retain content for only a specific project.
- Employee IDs that you can use to retain content for only a specific person.

Asset ID is simply another document property that's available in SharePoint and OneDrive. Your organization might already use other document properties and IDs to classify content. If so, you can also use those properties and values when you create an event—see step 6 that follows. The important point is that you must use some *property:value* combination in the document properties to associate that item with an event type.

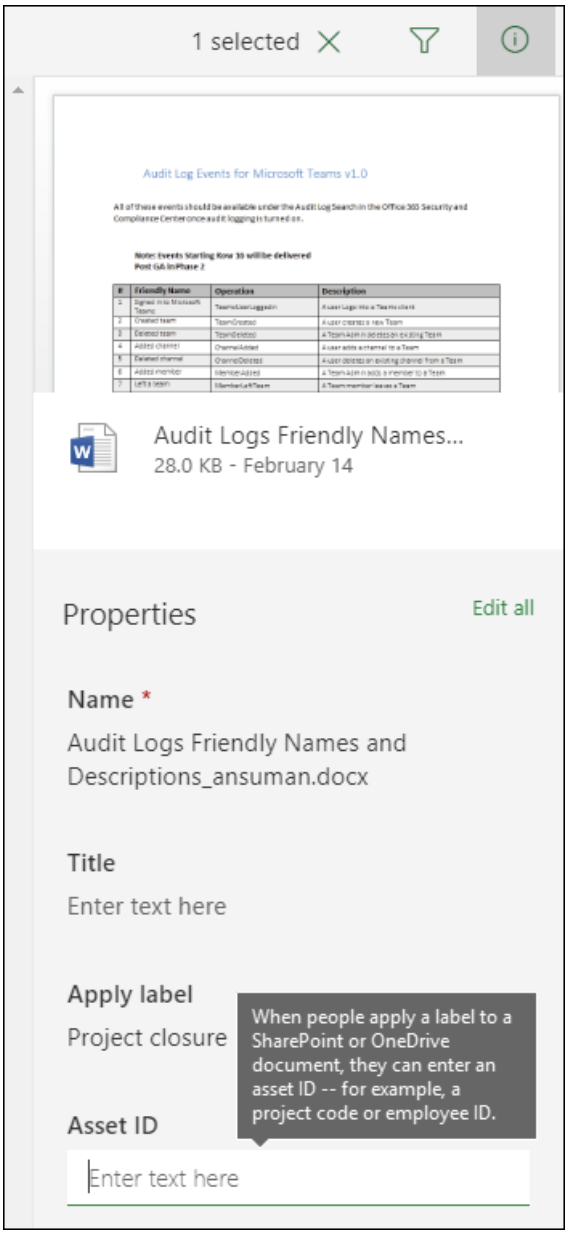

Step 5: Create an event

When a particular instance of that event type occurs, such as a product reaches its end of life, go to the **Records management** > **Events** page in the Microsoft 365 compliance center, and select **+ Create** to create an event. You trigger the event by creating it, here.

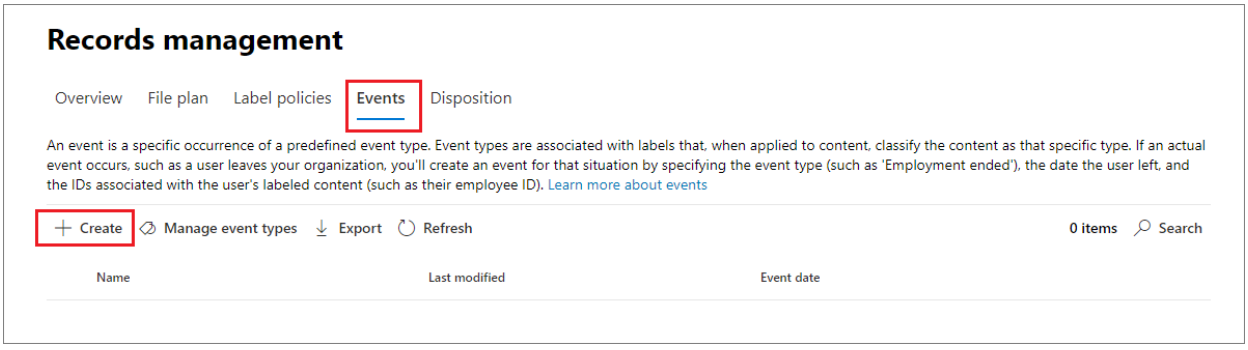

Up to one million events are supported per tenant.

Step 6: Choose the same event type used by the label in step 2

When you create the event, choose the same event type specified in the retention label settings in step 2. For example, if you selected **Product Lifetime** as your event type for the label settings, select **Product Lifetime** when you create the event. Only content with retention labels applied to it of that event type will have its retention period triggered.

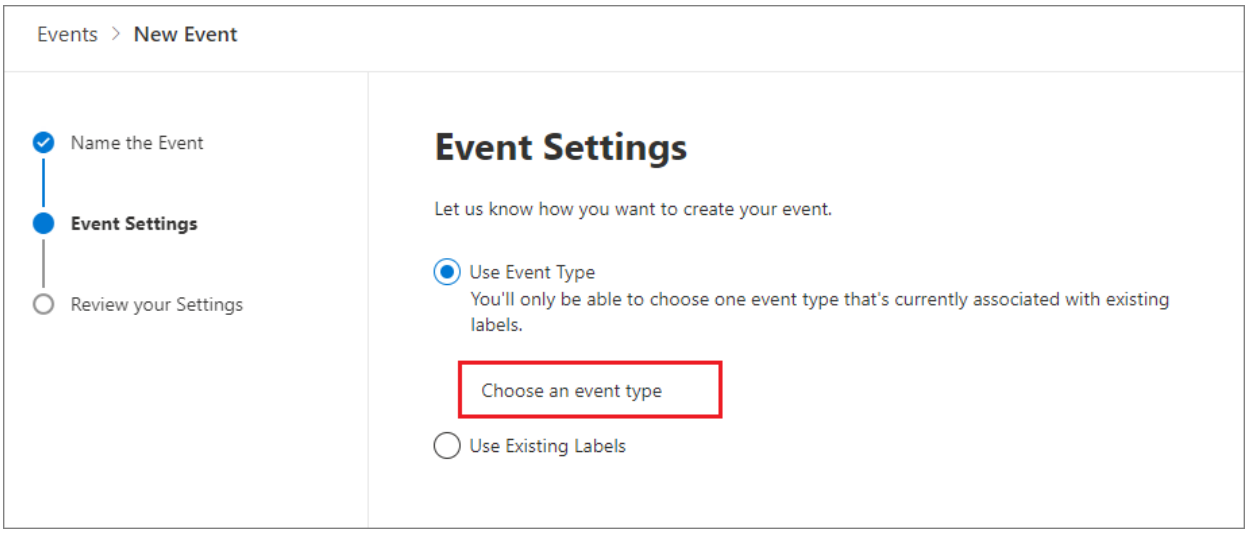

Alternatively, if you need to create an event for multiple retention labels that have different event types, select the **Choose Existing Labels** option. Then, select the labels that are configured for the event types you want to associate with this event.

Step 7: Enter keywords or an asset ID

Now you narrow the scope of the content by specifying asset IDs for SharePoint and OneDrive content, or keywords for Exchange content. For asset IDs, retention will be enforced only on content with the specified *property:value* pair. If an asset ID is not entered, all content with labels of that event type get the same retention date applied to them.

For example: If you're using the Asset ID property, enter `ComplianceAssetID:<value>` in the box for asset IDs shown below.

Your organization might have applied other properties and IDs to the documents related to this event type. For example, if you need to detect a specific product's records, the ID might be a combination of your custom property ProductID and the value "XYZ". In this case, you'd enter `ProductID:XYZ` in the box for asset IDs shown in the following picture.

For Exchange items, use keywords. You can use a query by using search operators such as AND, OR, and NOT. For more information, see Keyword queries and search conditions for Content Search.

Finally, choose the date when the event occurred; this date is used as the start of the retention period. After you create an event, that event date is synchronized to all the content with a retention label of that event type, asset ID, and keywords. As with any retention label, this synchronization can take up to seven days.

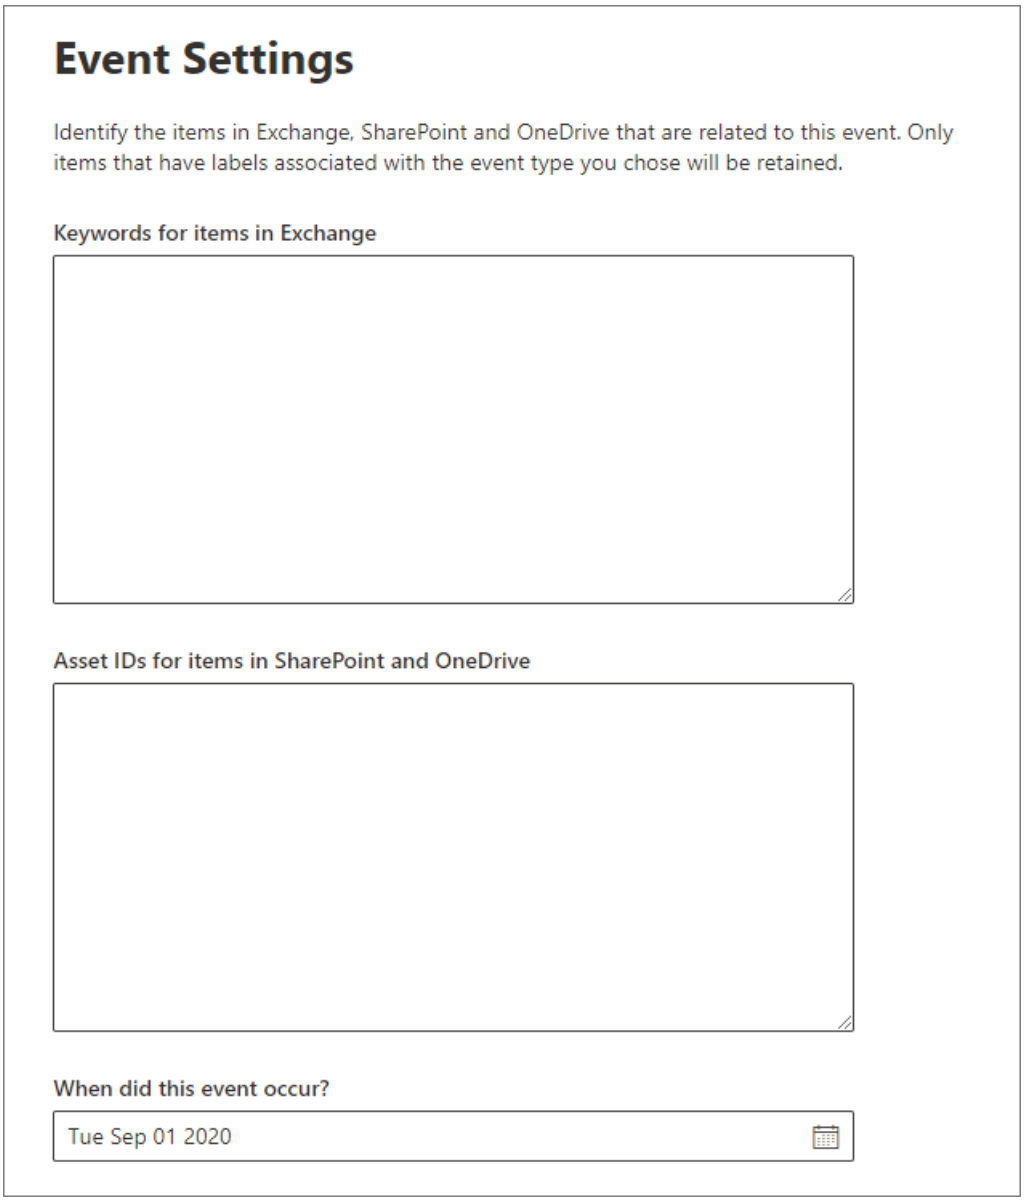

After creating an event, the retention settings take effect for the content that's already labeled and indexed. If the retention label is added to new content after the event is created, you must create a new event with the same details.

Deleting an event doesn't cancel the retention settings that are now in effect for the content that's already labeled. To do that, create a new event with the same details, but leave the date blank.

Use Content Search to find all content with a specific label or asset ID

After retention labels are assigned to content, you can use content search to find all content that's classified with a specific retention label or that contains a specific asset ID:

- To find all content with a specific retention label, choose the **Retention label** condition, and then enter the complete label name or part of the label name and use a wildcard.
- To find all content with a specific asset ID, enter the **ComplianceAssetID** property and a value, using the format `ComplianceAssetID:<value>`.

For more information, see Keyword queries and search conditions for Content Search.

Automate events by using PowerShell

You can use a PowerShell script to automate event-based retention from your business applications. The

PowerShell cmdlets available for event-based retention:

- Get-ComplianceRetentionEventType
- New-ComplianceRetentionEventType
- Remove-ComplianceRetentionEventType
- Set-ComplianceRetentionEventType
- Get-ComplianceRetentionEvent
- New-ComplianceRetentionEvent

Automate events by using a REST API

You can use a REST API to automatically create the events that trigger the start of the retention time.

A REST API is a service endpoint that supports sets of HTTP operations (methods), which provide create/retrieve/update/delete access to the service's resources. For more information, see Components of a REST API request/response. By using the Microsoft 365 REST API, events can be created and retrieved using the POST and GET methods.

There are two options for using the REST API:

- **Microsoft Power Automate or a similar application** to trigger the occurrence of an event automatically. Microsoft Power Automate is an orchestrator for connecting to other systems, so you don't need to write a custom solution. For more information, see the Power Automate website.
- **PowerShell or an HTTP client to call the REST API** to create events by using PowerShell (version 6 or later), which is part of a custom solution.

Before you use the REST API, as a global administrator, confirm the URL to use for the retention event call. To do this, run a GET retention event call by using the REST API URL:

```
https://ps.compliance.protection.outlook.com/psws/service.svc/ComplianceRetentionEvent
```

Check the response code. If it's 302, get the redirected URL from the Location property of the response header and use that URL instead of `https://ps.compliance.protection.outlook.com/psws/service.svc/ComplianceRetentionEvent` in the instructions that follow.

The events that get automatically created can be confirmed by viewing them in the Microsoft 365 compliance center > **Records management** > **Events**.

Use Microsoft Power Automate to create the event

Create a flow that creates an event using the Microsoft 365 REST API:

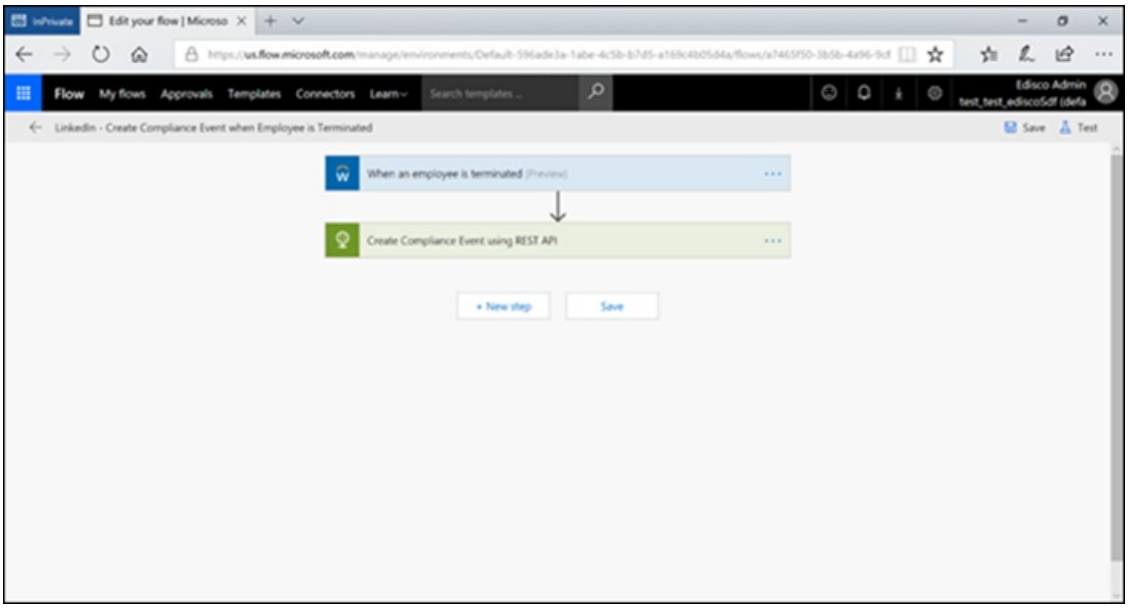

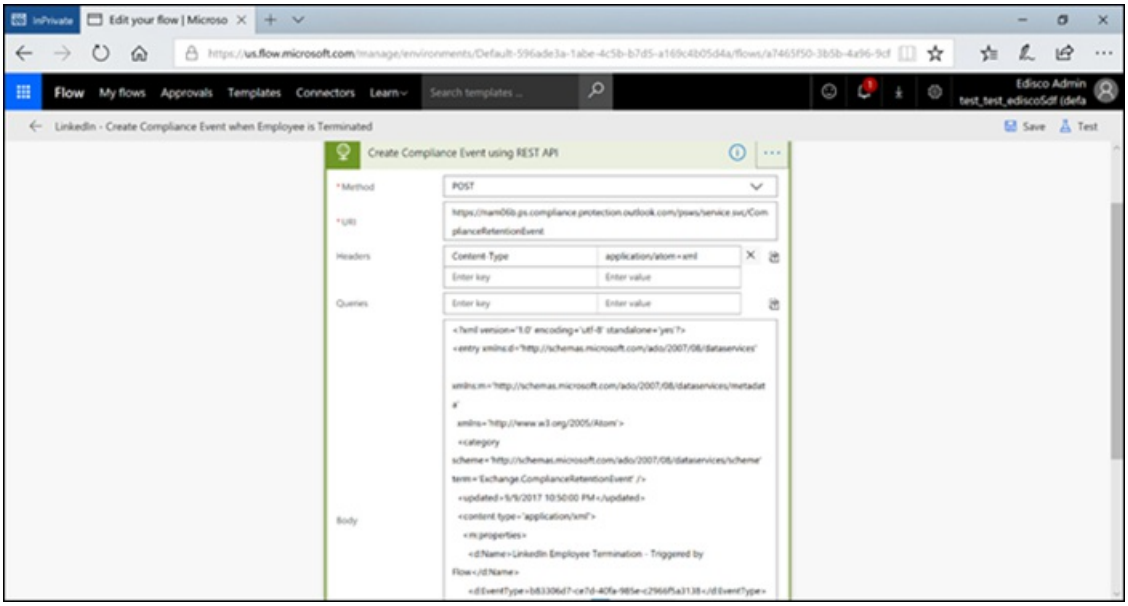

Create an event

Sample code to call the REST API:

- **Method**: POST
- **URL**: `https://ps.compliance.protection.outlook.com/psws/service.svc/ComplianceRetentionEvent`
- **Headers**: Key = Content-Type, Value = application/atom+xml
- **Body**:

```
<?xml version='1.0' encoding='utf-8' standalone='yes'?>

<entry xmlns:d='http://schemas.microsoft.com/ado/2007/08/dataservices'

xmlns:m='http://schemas.microsoft.com/ado/2007/08/dataservices/metadata'

xmlns='http://www.w3.org/2005/Atom'>

<category scheme='http://schemas.microsoft.com/ado/2007/08/dataservices/scheme'
term='Exchange.ComplianceRetentionEvent' />

<updated>9/9/2017 10:50:00 PM</updated>

<content type='application/xml'>

<m:properties>

<d:Name>Employee Termination </d:Name>

<d:EventType>99e0ae64-a4b8-40bb-82ed-645895610f56</d:EventType>

<d:SharePointAssetIdQuery>1234</d:SharePointAssetIdQuery>

<d:EventDateTime>2018-12-01T00:00:00Z </d:EventDateTime>

</m:properties>

</content>

</entry>
```

- **Authentication**: Basic
- **Username**: "Complianceuser"
- **Password**: "Compliancepassword"

Available parameters

PARAMETERS	DESCRIPTION	NOTES
<d:Name> </d:Name>	Provide a unique name for the event,	Cannot contain trailing spaces or the following characters: % * \ & < > | # ? , : ;
<d:EventType> </d:EventType>	Enter event type name (or Guid),	Example: "Employee termination". Event type has to be associated with a retention label.
<d:SharePointAssetIdQuery> </d:SharePointAssetIdQuery>	Enter "ComplianceAssetId:" + employee ID	Example: "ComplianceAssetId:12345"
<d:EventDateTime> </d:EventDateTime>	Event Date and Time	Format: yyyy-MM-ddTHH:mm:ssZ, Example: 2018-12-01T00:00:00Z

Response codes

RESPONSE CODE	DESCRIPTION
302	Redirect

RESPONSE CODE	DESCRIPTION
201	Created
403	Authorization Failed
401	Authentication Failed

Get events based on a time range

- **Method**: GET
- **URL**:
  ```
  https://ps.compliance.protection.outlook.com/psws/service.svc/ComplianceRetentionEvent?BeginDateTime=2019-01-11&EndDateTime=2019-01-16
  ```
- **Headers**: Key = Content-Type, Value = application/atom+xml
- **Authentication**: Basic
- **Username**: "Complianceuser"
- **Password**: "Compliancepassword"

Response codes

RESPONSE CODE	DESCRIPTION
200	OK, A list of events in atom+ xml
404	Not found
302	Redirect
401	Authorization Failed
403	Authentication Failed

Get an event by ID

- **Method**: GET
- **URL**:
  ```
  https://ps.compliance.protection.outlook.com/psws/service.svc/ComplianceRetentionEvent('174e9a86-74ff-4450-8666-7c11f7730f66')
  ```
- **Headers**: Key = Content-Type, Value = application/atom+xml
- **Authentication**: Basic
- **Username**: "Complianceuser"
- **Password**: "Compliancepassword"

Response codes

RESPONSE CODE	DESCRIPTION
200	OK, The response body contains the event in atom+xml
404	Not found

RESPONSE CODE	DESCRIPTION
302	Redirect
401	Authorization Failed
403	Authentication Failed

Get an event by name

- **Method**: GET
- **URL**: `https://ps.compliance.protection.outlook.com/psws/service.svc/ComplianceRetentionEvent`
- **Headers**: Key = Content-Type, Value = application/atom+xml
- **Authentication**: Basic
- **Username**: "Complianceuser"
- **Password**: "Compliancepassword"

Response codes

RESPONSE CODE	DESCRIPTION
200	OK, The response body contains the event in atom+xml
404	Not found
302	Redirect
401	Authorization Failed
403	Authentication Failed

Use PowerShell or any HTTP client to create the event

PowerShell must be version 6 or later.

In a PowerShell session, run the following script:

```
param([string]$baseUri)

$userName = "UserName"

$password = "Password"

$securePassword = ConvertTo-SecureString $password -AsPlainText -Force

$credentials = New-Object System.Management.Automation.PSCredential($userName, $securePassword)

$EventName="EventByRESTPost-$(([Guid]::NewGuid()).ToString('N'))"

Write-Host "Start to create an event with name: $EventName"

$body = "<?xml version='1.0' encoding='utf-8' standalone='yes'?>

<entry xmlns:d='http://schemas.microsoft.com/ado/2007/08/dataservices'

xmlns:m='http://schemas.microsoft.com/ado/2007/08/dataservices/metadata'

xmlns='http://www.w3.org/2005/Atom'>
```

```
<category scheme='http://schemas.microsoft.com/ado/2007/08/dataservices/scheme'
term='Exchange.ComplianceRetentionEvent' />

<updated>7/14/2017 2:03:36 PM</updated>

<content type='application/xml'>

<m:properties>

<d:Name>$EventName</d:Name>

<d:EventType>e823b782-9a07-4e30-8091-034fc01f9347</d:EventType>

<d:SharePointAssetIdQuery>'ComplianceAssetId:123'</d:SharePointAssetIdQuery>

</m:properties>

</content>

</entry>"

$event = $null

try

{

$event = Invoke-RestMethod -Body $body -Method 'POST' -Uri "$baseUri/ComplianceRetentionEvent" -ContentType
"application/atom+xml" -Authentication Basic -Credential $credentials -MaximumRedirection 0

}

catch

{

$response = $_.Exception.Response

if($response.StatusCode -eq "Redirect")

{

$url = $response.Headers.Location

Write-Host "redirected to $url"

$event = Invoke-RestMethod -Body $body -Method 'POST' -Uri $url -ContentType "application/atom+xml" -
Authentication Basic -Credential $credentials -MaximumRedirection 0

}

}

$event | fl *
```

Use retention labels to manage the lifecycle of documents stored in SharePoint

11/2/2020 • 16 minutes to read • Edit Online

Microsoft 365 licensing guidance for security & compliance.

This article describes how you can manage the lifecycle of documents that are stored in SharePoint by using automatically applied retention labels and event-based retention.

The auto-apply functionality uses SharePoint metadata for document classification. The example in this article is for product-related documents, but the same concepts can be used for other scenarios. For example, in the oil and gas industry, you could use it to manage the lifecycle of documents about physical assets such as oil platforms, well logs, or production licenses. In the financial services industry, you could manage bank account, mortgage, or insurance contract documents. In the public sector, you could manage construction permits or tax forms.

In this article, we'll look at the information architecture and definition of the retention labels. Then we'll classify documents by auto-applying the labels. And finally we'll generate the events that initiate the retention period.

Information architecture

Our scenario is a manufacturing company that uses SharePoint to store all the documents about the products that the company develops. These documents include product specifications, agreements with suppliers, and user manuals. When these documents are stored in SharePoint through Enterprise Content Management policies, document metadata is defined, which is used to classify them. Each document has the following metadata properties:

- **Doc Type** (such as product specification, agreement, or user manual)
- **Product Name**
- **Status** (draft or final)

This metadata forms a base content type called *Production Document* for all the documents.

Columns

Name	Type	Status	Source
Name	File	Required	Document
Title	Single line of text	Optional	Item
Doc Type	Managed Metadata	Required	
Product Name	Managed Metadata	Required	
Status	Choice	Required	

NOTE

The **Doc Type** and **Status** properties are used by retention policies later in this scenario to classify and auto-apply retention labels.

We might have several content types that represent different types of documents, but let's focus on the product documentation.

In this scenario, we use the Managed Metadata service and the Term Store to create a term set for *Doc Type* and another one for *Product Name*. For each term set, we create a term for each value. It would look like something like this in Term Store for your SharePoint organization:

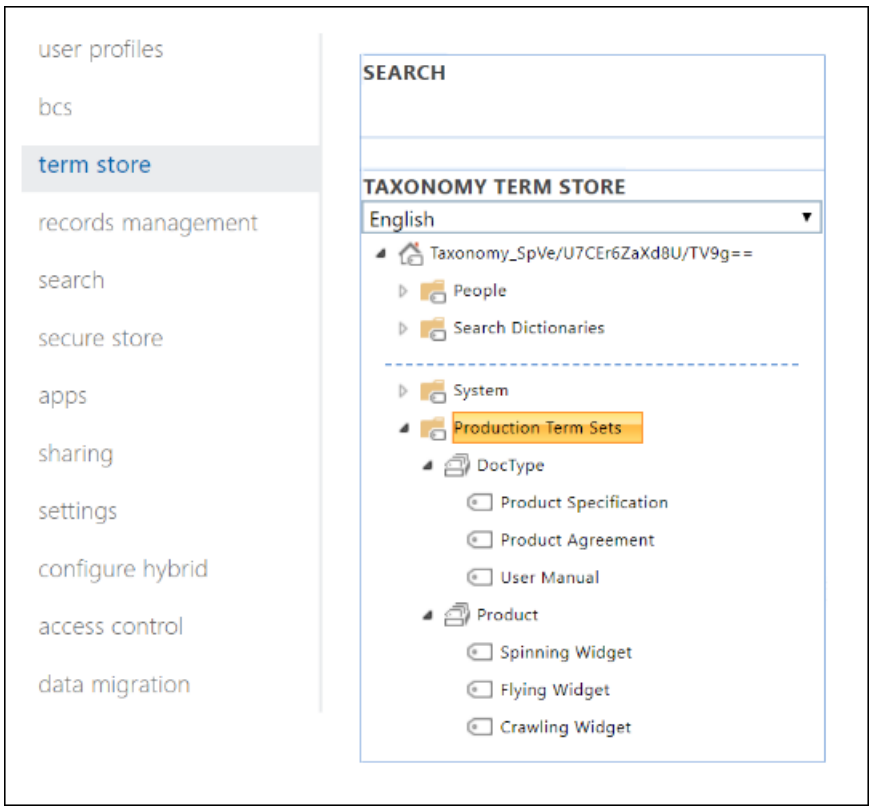

Content Type can be created and published by using the Content Type Hub. You can also create and publish a content type by using site provisioning tools, such as the PnP provisioning framework or site design JSON schema.

Each product has a dedicated SharePoint site that contains one document library that has the right content types enabled. All documents are stored in this document library.

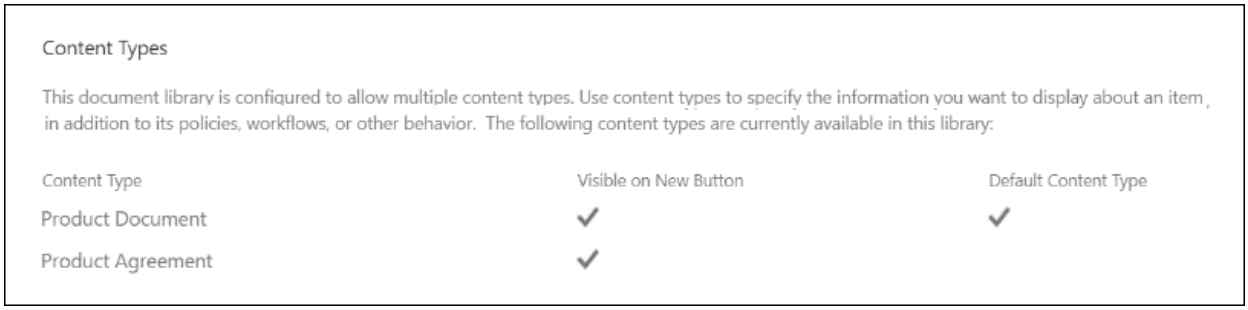

> **NOTE**
> Instead of having a SharePoint site per product, the manufacturing company in this scenario could use a Microsoft Team per product to support collaboration among members of the team, such as through persistent chat, and use the **Files** tab in Teams for document management. In this article we only focus on documents, so, we'll only use a site.

Here's a view of the document library for the Spinning Widget product:

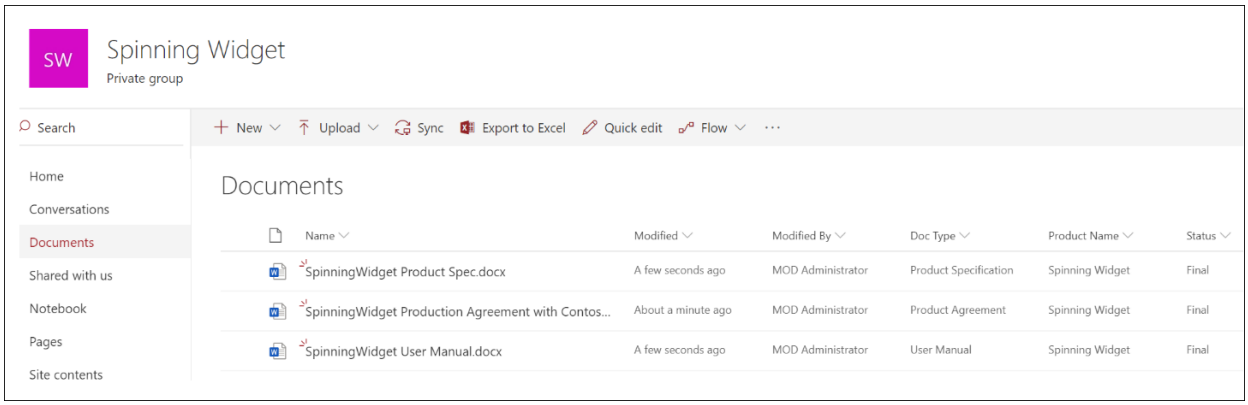

Now that we have the basic information architecture in place for document management, let's look at the retention and disposal strategy for the documents that use the metadata and how we classify those documents.

Retention and disposition

The manufacturing company's compliance and data governance policies dictate how data is preserved and disposed of. Product-related documents must be kept for as long as the product is manufactured and for a certain additional period. The additional period differs for product specifications, agreements, and user manuals. The following table indicates the retention and disposition requirements:

DOCUMENT TYPE	RETENTION	DISPOSITION
Product specifications	5 years after production stops	Delete
Product agreements	10 years after production stops	Review
User manuals	5 years after production stops	Delete
All other types of documents	Don't actively retain	Delete when document is older than 3 years A document is considered older than 3 years if it hasn't been modified within the last 3 years.

We use the Microsoft 365 compliance center to create the following retention labels:

- Product Specification
- Product Agreement
- User Manual

In this article, we only show how to create and auto-apply the Product Specification retention label. To implement the complete scenario, you would also create and auto-apply retention labels for the other two document types.

Settings for the Product Specification retention label

Here's the file plan for the Product Specification retention label:

- **Name:** Product Specification
- **Description for users:** Retain for 5 years after production stops.
- **Description for admins:** Retain for 5 years after production stops, auto delete, event-based retention,

event type is *Product Cessation*.

- **Retention action:** Retain and delete.
- **Retention duration:** 5 years (1,825 days).
- **Record label**: Configure the retention label to mark items as a record, which means the labeled documents can't then be modified or deleted by users.
- **File plan descriptors:** For simplifying the scenario, no optional file descriptors are provided.

The following screenshot shows the settings when you create the Product Specification retention label in the Microsoft 365 compliance center. You can create the *Product Cessation* event type when you create the retention label. See the procedure in the following section.

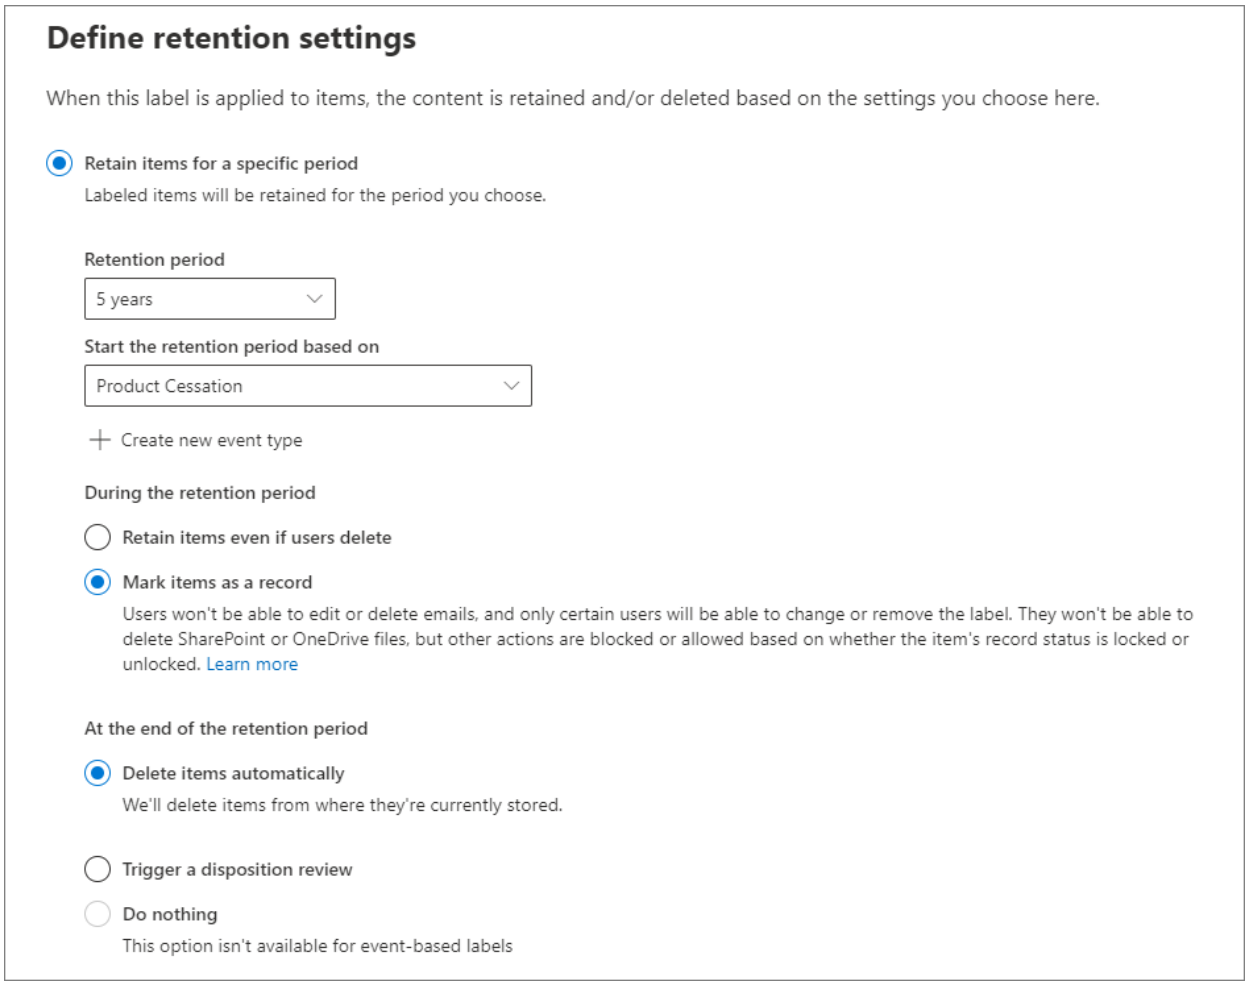

> **NOTE**
>
> To avoid a 5-year wait for document deletion, set the retention duration to ***1 day*** if you're recreating this scenario in a test environment.

Create an event type when you create a retention label

1. On the **Define retention settings** page of the Create retention label wizard, after **Start the retention period based on**, select **Create new event type**:

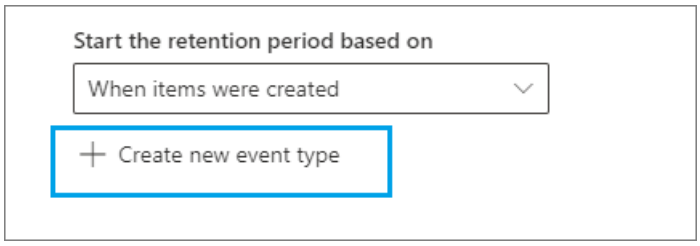

2. On the **Name your event type** page, enter **Product Cessation** and an optional description. Then select **Next**, **Submit**, and **Done**.

3. Back on the **Define retention settings** page, for **Start the retention period based on**, use the dropdown box to select the **Product Cessation** event type that you created.

   Here's what the settings look like for the Product Specification retention label:

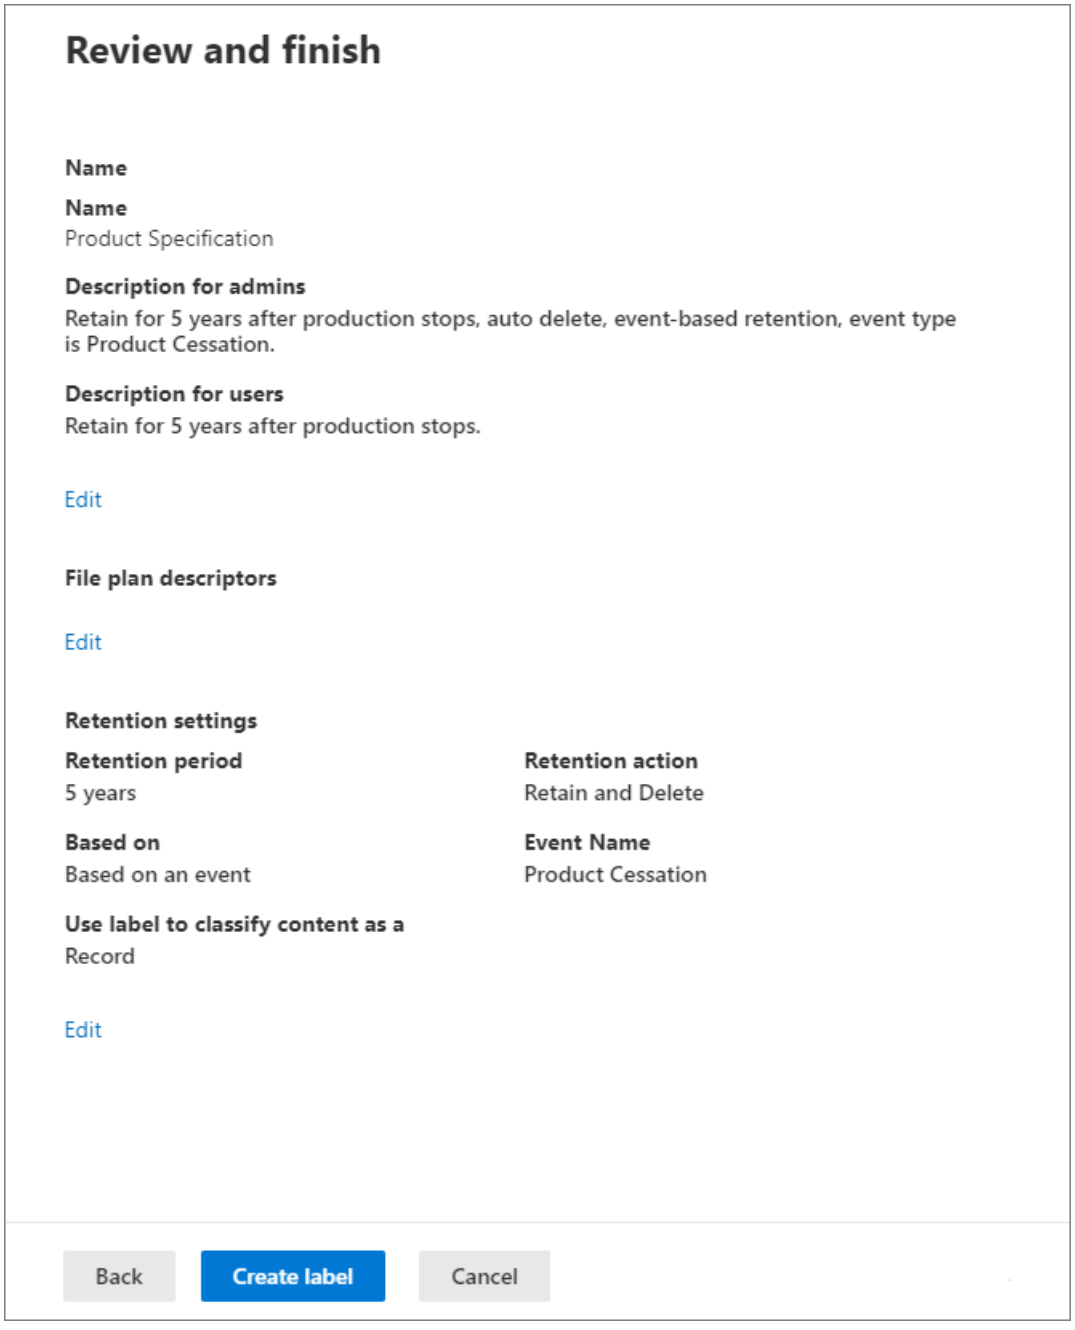

4. Select **Create label**, and on the next page when you see the options to publish the label, auto-apply the label, or just save the label: Select **Just save the label for now**, and then select **Done**.

> **TIP**
>
> For more detailed steps, see Create a label whose retention period is based on an event.

Now let's look at how we'll auto-apply the retention label to product-specification content.

Auto-apply retention labels to documents

We're going to use Keyword Query Language (KQL) to auto-apply the retention labels that we created. KQL is the language that's used to build search queries. In KQL, you can search by using keywords or managed properties. For more information, see Keyword Query Language (KQL) syntax reference.

Basically, we want to tell Microsoft 365 to "apply the **Product Specification** retention label to all documents that have a **Status** of **Final** and a **Doc Type** of **Product Specification**." Recall that **Status** and **Doc Type** are the site columns that we defined for the Product Documentation content type in the Information architecture section. To do this, we need to configure the search schema.

When SharePoint indexes content, it automatically generates crawled properties for each site column. For this scenario, we're interested in the **Doc Type** and **Status** properties. We need documents in the library that are the right content type and have the site columns filled in for search to create the crawled properties.

In the SharePoint admin center, open the Search configuration, and select **Manage Search Schema** to view and configure the crawled properties.

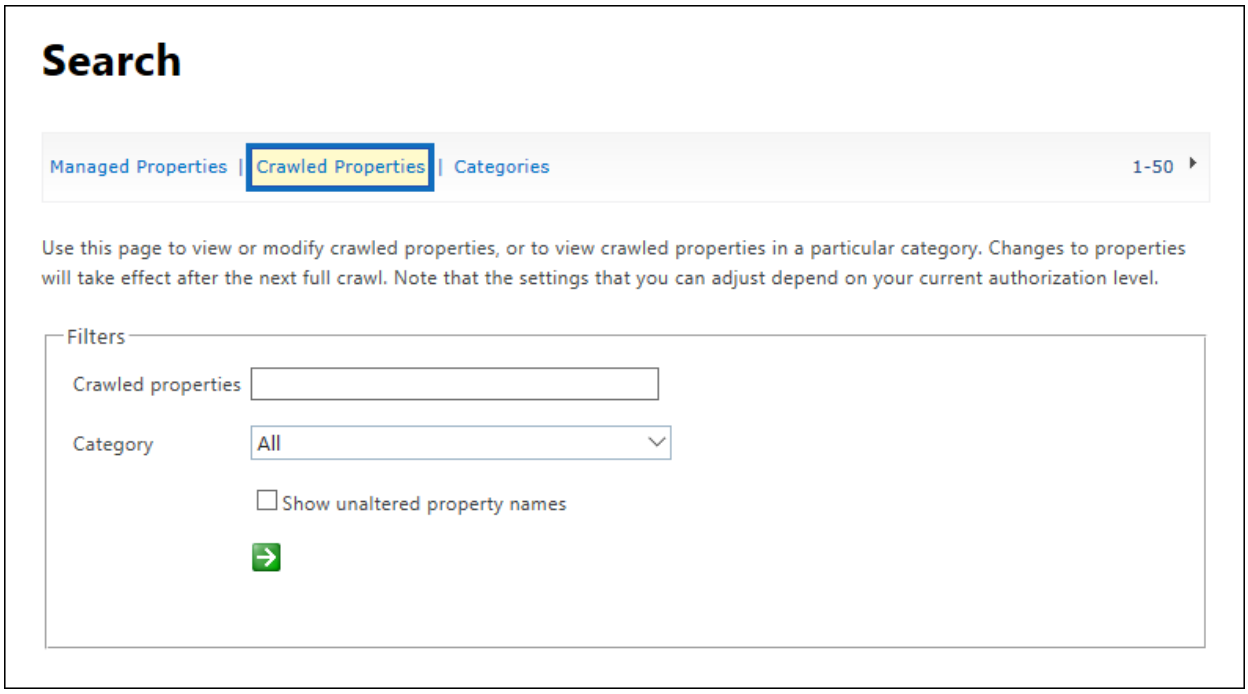

If we type ***status*** in the **Crawled properties** box and select the green arrow, we should see a result like this:

The **ows__Status** property (notice the double underscore) is the one that interests us. It maps to the **Status** property of the Production Document content type.

Now, if we type ***ows_doc*** and select the green arrow, we should see something like this:

The **ows_Doc_x0020_Type** property is the second property that interests us. It maps to the **Doc Type** property of the Production Document content type.

> **TIP**
>
> To identify the name of a crawled property for this scenario, go to the document library that contains the production documents. Then go to the library settings. For **Columns**, select the name of the column (for example, **Status** or **Doc Type**) to open the site column page. The *Field* parameter in the URL for that page contains the name of the field. This field name, prefixed with "ows_", is the name of the crawled property. For example, the URL
>
> `https://tenantname.sharepoint.com/sites/SpinningWidget/_layouts/15/FldEdit.aspx?List=%7BC38C2F45-3BD6-4C3B-AA3B-EF5DF6B3D172%7D&Field=_Status`
>
> corresponds to the *ows__Status* crawled property.

If the crawled properties you're looking for don't appear in the Manage Search Schema section in SharePoint admin center:

- Maybe the documents haven't been indexed. You can force a reindex of the library by going to **Document library settings** > **Advanced Settings**.
- If the document library is in a modern site, make sure that the SharePoint admin is also a site collection admin.

For more information about crawled and managed properties, see Automatically created managed properties in SharePoint Server.

Map crawled properties to pre-defined managed properties

KQL can't use crawled properties in search queries. It has to use a managed property. In a typical search scenario, we create a managed property and map it to the crawled property that we need. However, for auto-applying retention labels, you can only specify pre-defined managed properties in KQL, not custom managed properties. There's a set of predefined managed properties in the system for string *RefinableString00* to *RefinableString199* that you can use. For a complete list, see Default unused managed properties. These default managed properties are typically used for defining search refiners.

For the KQL query to automatically apply the correct retention label to product document content, we map the crawled properties *ows_Doc_x0020_Type and ows__Status* to two refinable managed properties. In our test environment for this scenario, **RefinableString00** and **RefinableString01** aren't being used. We determined this by looking at **Managed Properties** in **Manage Search Schema** in the SharePoint admin center.

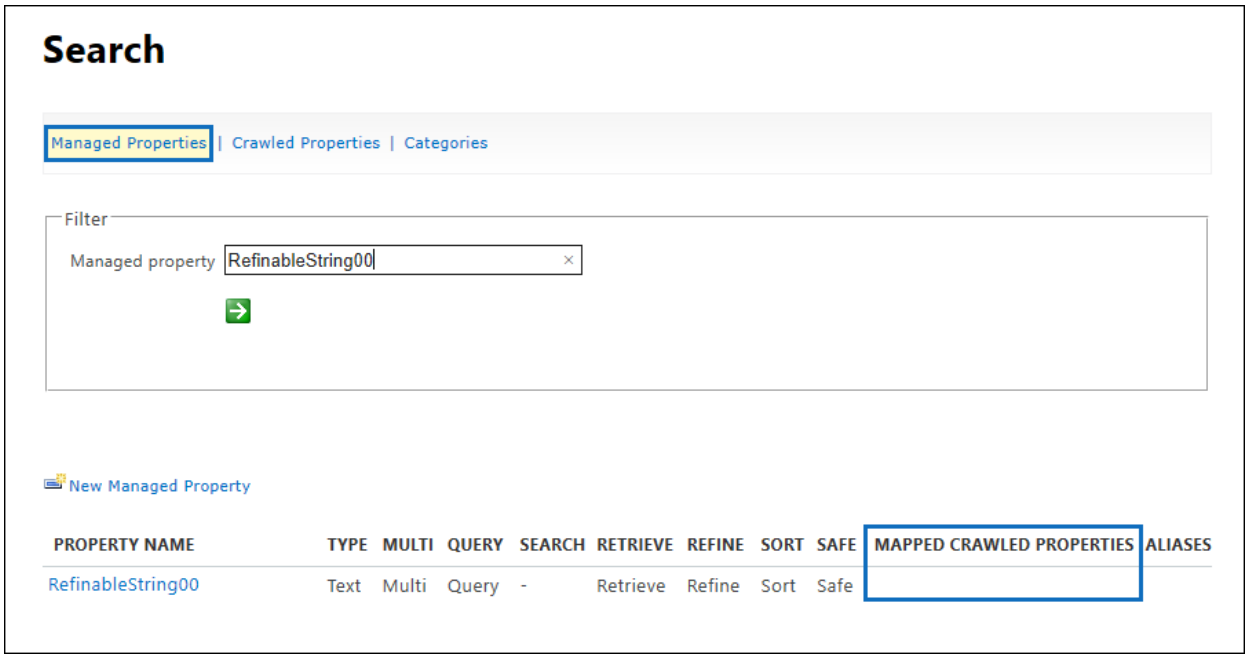

Notice that the **Mapped Crawled Properties** column in the previous screenshot is empty.

To map the **ows_Doc_x0020_Type** crawled property, follow these steps:

1. In the **Managed property** filter box, type ***RefinableString00*** and select the green arrow.
2. In the results list, select the **RefinableString00** link, and then scroll down to the **Mappings to crawled properties** section.
3. Select **Add a Mapping**, and then type ***ows_Doc_x0020_Type*** in the **Search for a crawled property name** box in the **Crawled property selection** window. Select **Find**.
4. In the results list, select **ows_Doc_x0020_Type** and then select **OK**.

   In the **Mapped Crawled Properties** section, you should see something similar to this screenshot:

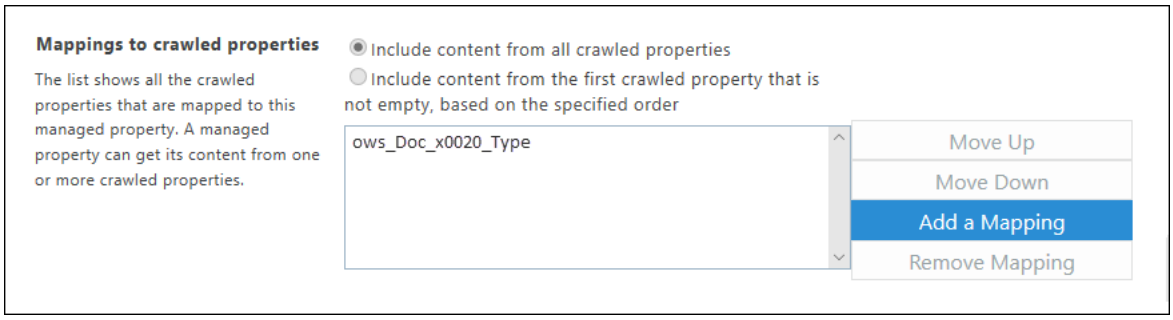

5. Scroll to the bottom of the page and select **OK** to save the mapping.

Repeat these steps to map **RefinableString01** and **ows__Status**.

Now you should have two managed properties mapped to the two crawled properties:

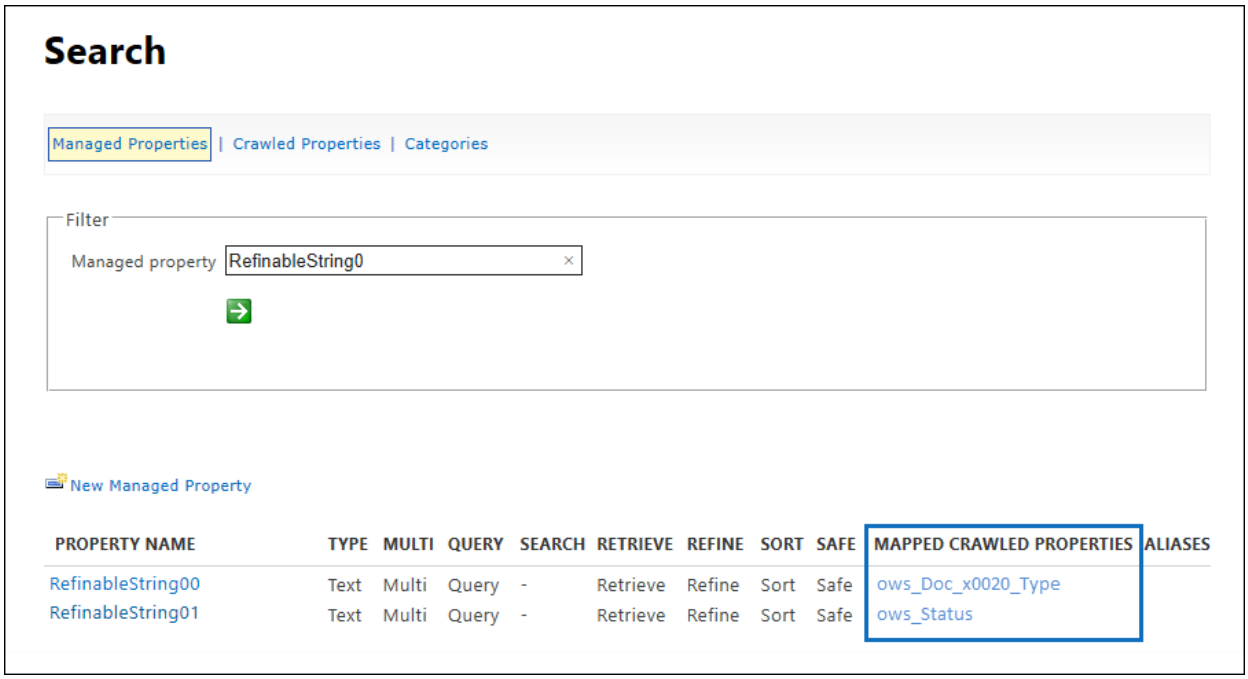

Let's verify that our setup is correct by running an enterprise search. In a browser, go to *https://<your_tenant>.sharepoint.com/search*. In the search box, type ***RefinableString00:"Product Specification"*** and press enter. This search should return all documents that have a **Product Specification** of ***Doc Type***.

Now in the search box, type **RefinableString00:"Product Specification" AND RefinableString01:Final** and press enter. This should return all documents that have **Product Specification** of ***Doc Type*** and a **Status** of ***Final***.

Create auto-apply label policies

Now that we've verified that the KQL query is working, let's create an auto-apply label policy that uses a KQL query to automatically apply the Product Specification retention label to the appropriate documents.

1. In the compliance center, go to **Records management** > **Label policies** > **Auto-apply a label**.

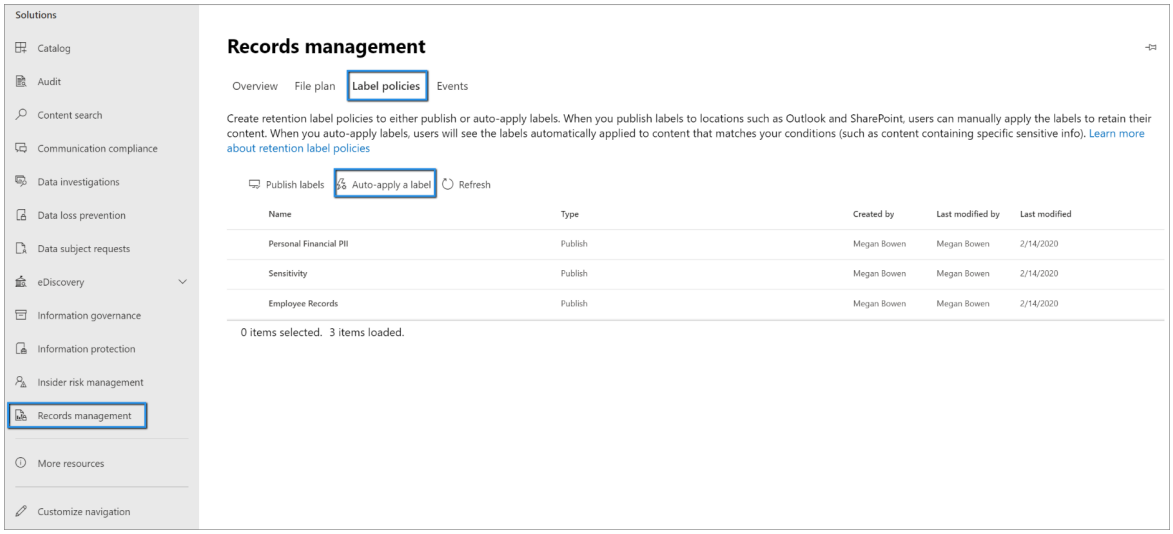

2. In the Create auto-labeling policy wizard, on the **Name your auto-labeling policy** page, enter a name such as **Auto-apply Product Specification label**, and an optional description. Then select **Next**.
3. On the **Choose the type of content you want to apply this label to** page, select **Apply label to content that contains specific words or phrases, or properties**, and then select **Next**.

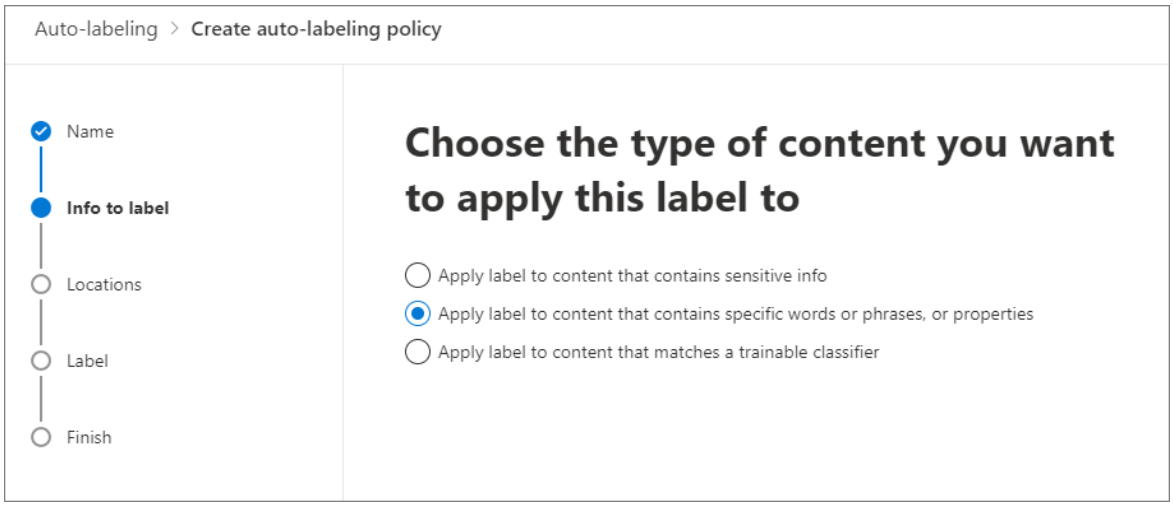

This option lets us provide the same KQL search query that we tested in the previous section. The query returns all Product Specification documents that have a status of *Final*. When we use this same query in the auto-apply label policy, the Product Specification retention label will be automatically applied to all documents that match it.

4. On the **Apply label to content matching this query** page, type **RefinableString00:"Product Specification" AND RefinableString01:Final**, and then select **Next**.

5. On the **Choose locations to apply the policy** page, you select the content locations that you want to apply the policy to. For this scenario, we apply the policy only to SharePoint locations, because all the production documents are stored in SharePoint document libraries. Toggle the status for **Exchange email**, **OneDrive accounts**, and **Microsoft 365 Groups** to **Off**. Make sure that the status for SharePoint sites is set to **On** before you select **Next**:

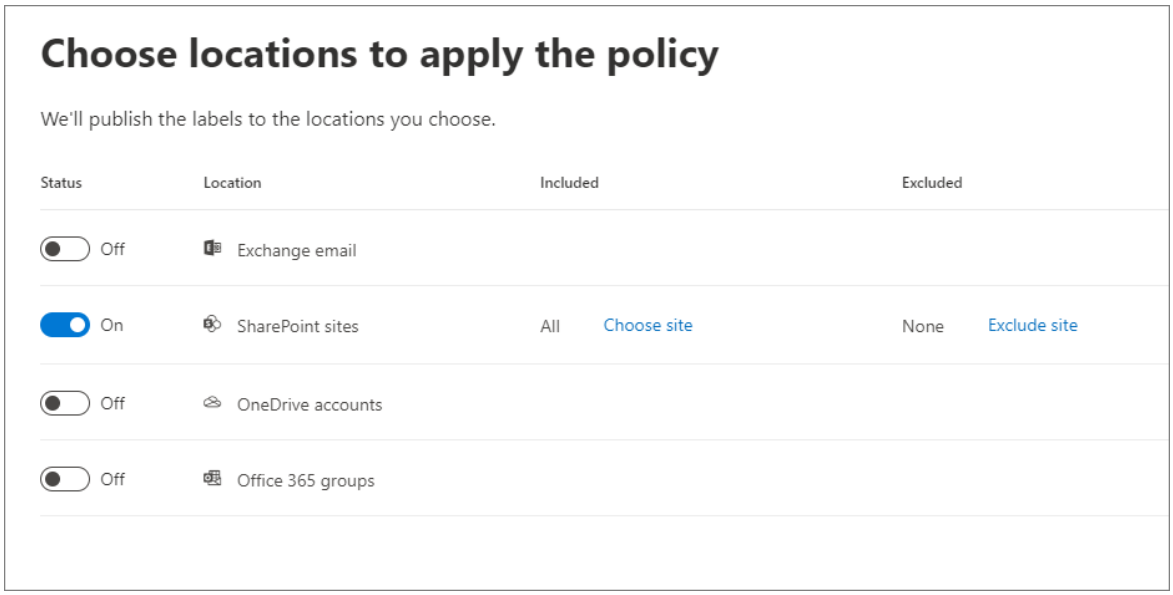

TIP

Instead of applying the policy to all SharePoint sites, you can select **Choose site** and add the URLs for specific SharePoint sites.

6. On the **Choose a label to auto-apply** page, select **Add label**.
7. From the list of retention labels, select **Product Specification**. Then select **Add** and **Next**.
8. Review your settings:

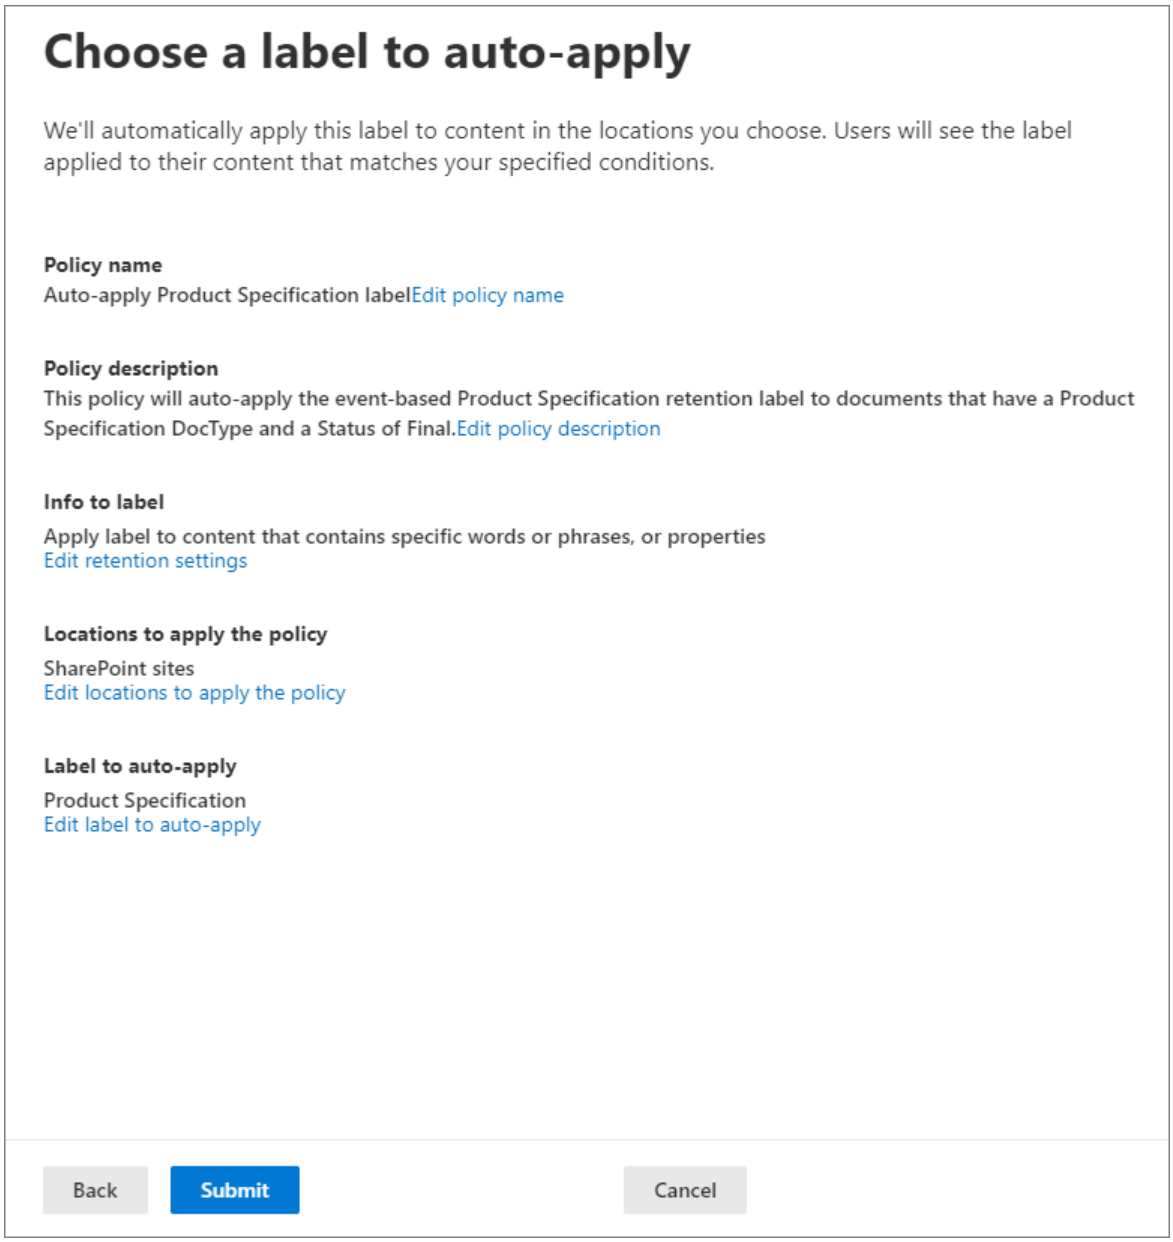

9. Select **Submit** to create the auto-apply label policy.

> **NOTE**
>
> It takes up to 7 days to automatically apply the Product Specification label to all documents that match the KQL search query.

Verify that the retention label was automatically applied

After 7 days, use activity explorer in the compliance center to verify that the auto-apply label policy that we created did automatically apply the retention labels to the product documents.

Also look at the properties of the documents in the Document Library. In the information panel, you can see that the retention label is applied to a selected document.

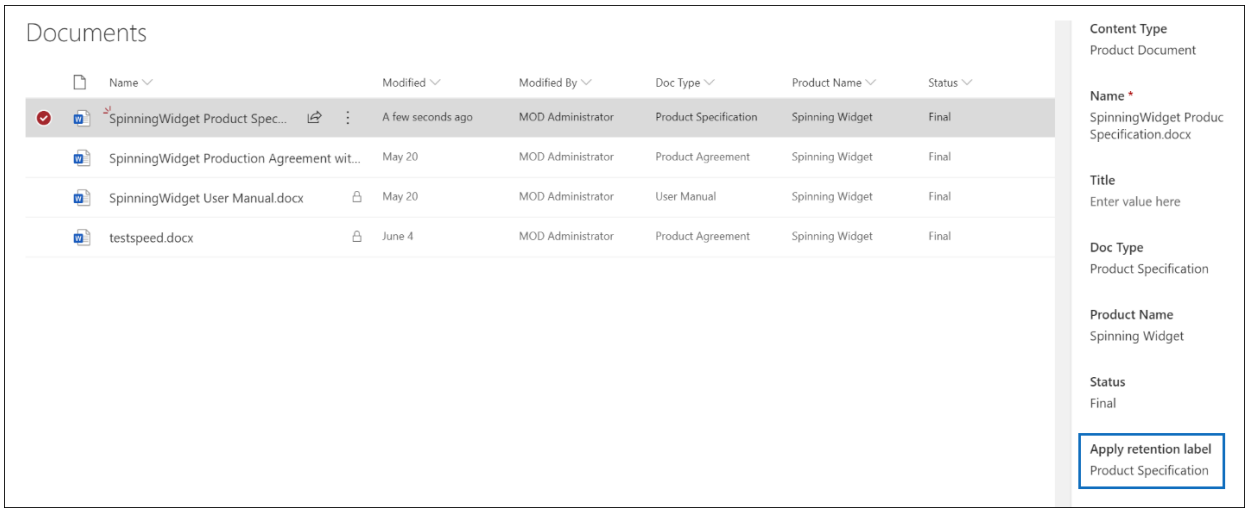

Because the retention labels were auto-applied to documents, those documents are protected from deletion because the retention label was configured to declare the documents as *records*. As an example of this protection, we get the following error message when we try to delete one of these documents:

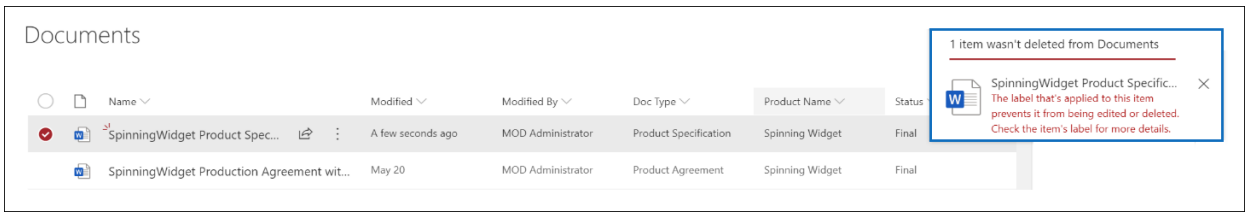

Generate the event that triggers the retention period

Now that the retention labels are applied, let's focus on the event that will indicate the end of production for a particular product. This event triggers the beginning of the retention period that's defined in the retention labels. For example, for product specification documents, the 5-year retention period begins when the "end of production" event is triggered.

You can manually create the event in the Microsoft 365 compliance center by going to **Records Managements** > **Events**. You would choose the event type, set the correct asset IDs, and enter a date for the event. For more information, see Start retention when an event occurs.

But for this scenario, we'll automatically generate the event from an external production system. The system is a simple SharePoint list that indicates whether a product is in production. A Power Automate flow that's associated with the list will trigger the event. In a real-world scenario, you could use various systems to generate the event, such as an HR or CRM system. Power Automate contains many ready-to-use interactions and building block for Microsoft 365 workloads, such as Microsoft Exchange, SharePoint, Teams, and Dynamics 365, plus third-party apps such as Twitter, Box, Salesforce, and Workdays. This feature makes it easy to integrate Power Automate with various systems. For more information, see Automate event-driven retention.

The following screenshot shows the SharePoint list that will be used the trigger the event:

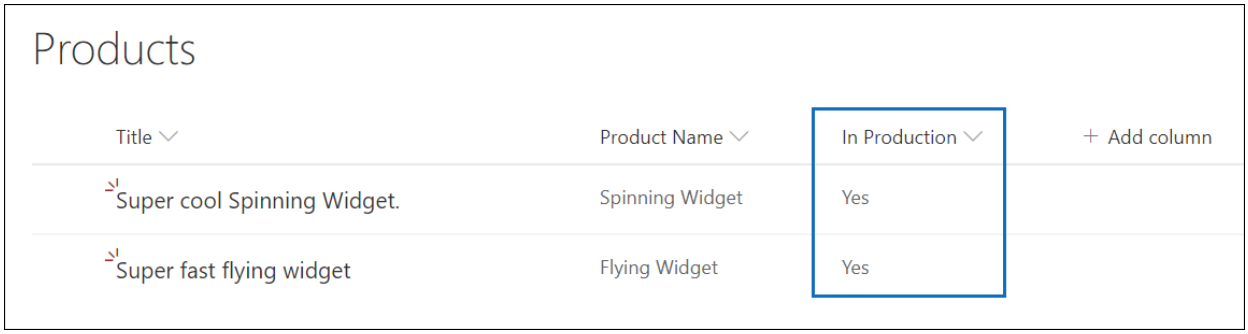

There are two products currently in production, as indicated by the ***Yes*** in the **In Production** column. When the

value in this column is set to ***No*** for a product, the flow associated with the list will automatically generate the event. The event triggers the start of the retention period for the retention label that was auto-applied to the corresponding product documents.

For this scenario, we use the following flow to trigger the event:

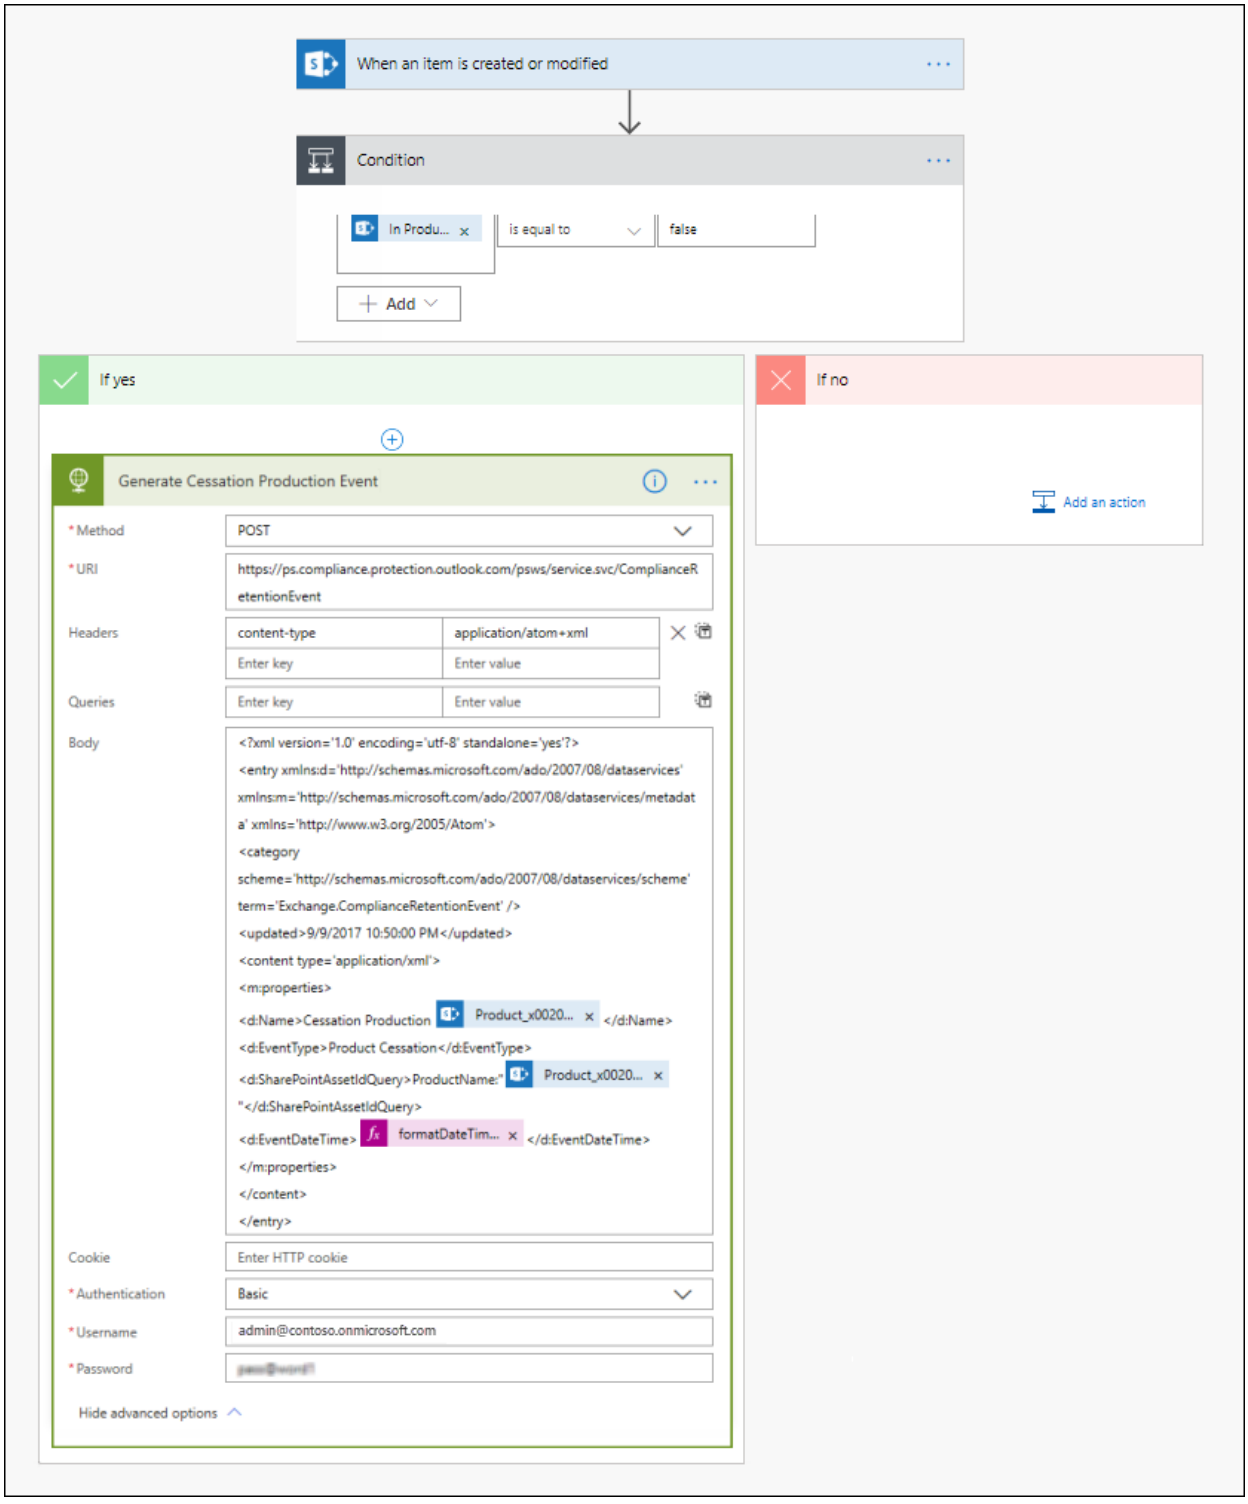

To create this flow, start from a SharePoint connector and select the **When an item is created or modified** trigger. Specify the site address and list name. Then add a condition based on when the **In Production** list column value is set to ***No*** (or equal to *false* on the condition card). Then add an action based on the built-in HTTP template. Use the values in the following section to configure the HTTP action. You can copy the values for the **URI** and **Body** properties from the following section and paste them into the template.

- **Method**: POST
- **URI**: `https://ps.compliance.protection.outlook.com/psws/service.svc/ComplianceRetentionEvent`

- **Headers**: Key = Content-Type, Value = application/atom+xml
- **Body**:

```
<?xml version='1.0' encoding='utf-8' standalone='yes'>
<entry xmlns:d='http://schemas.microsoft.com/ado/2007/08/dataservices'
xmlns:m='http://schemas.microsoft.com/ado/2007/08/dataservices/metadata'
xmlns='https://www.w3.org/2005/Atom'>
<category scheme='http://schemas.microsoft.com/ado/2007/08/dataservices/scheme'
term='Exchange.ComplianceRetentionEvent'>
<updated>9/9/2017 10:50:00 PM</updated>
<content type='application/xml'>
<m:properties>
<d:Name>Cessation Production @{triggerBody()?['Product_x0020_Name']?['Value']}</d:Name>
<d:EventType>Product Cessation&lt;</d:EventType>
<d:SharePointAssetIdQuery>ProductName:&quot;@{triggerBody()?['Product_x0020_Name']?['Value']}
<d:SharePointAssetIdQuery>
<d:EventDateTime>@{formatDateTime(utcNow(),'yyyy-MM-dd')}</d:EventDateTime>
</m:properties>
</content&gt>
</entry>
```

This list describes the parameters in the **Body** property of the action that must be configured for this scenario:

- **Name**: This parameter specifies the name of the event that will be created in the Microsoft 365 compliance center. For this scenario, the name is "Cessation Production *xxx*", where *xxx* is the value of the **ProductName** managed property that we created earlier.
- **EventType**: The value for this parameter corresponds to the event type that the created event will apply to. This event type was defined when you created the retention label. For this scenario, the event type is "Product Cessation."
- **SharePointAssetIdQuery**: This parameter defines the asset ID for the event. Event-based retention needs a unique identifier for the document. We can use asset IDs to identify the documents that a particular event applies to or, as in this scenario, the metadata column **Product Name**. To do this, we need to create a new **ProductName** managed property that can be used in the KQL query. (Alternatively, we could use **RefinableString00** instead of creating a new managed property). We also need to map this new managed property to the **ows_Product_x0020_Name** crawled property. Here's a screenshot of this managed property.

- **EventDateTime**: This parameter defines the date that the event occurs. Use the current date format:

  formatDateTime(utcNow(),'yyyy-MM-dd')

Putting it all together

Now the retention label is created and auto-applied, and the flow is configured and created. When the value in the **In Production** column for the Spinning Widget product in the Products list is changed from ***Yes*** to ***No***, the flow is triggered to create the event. To see this event in the compliance center, go to **Records management** > **Events**.

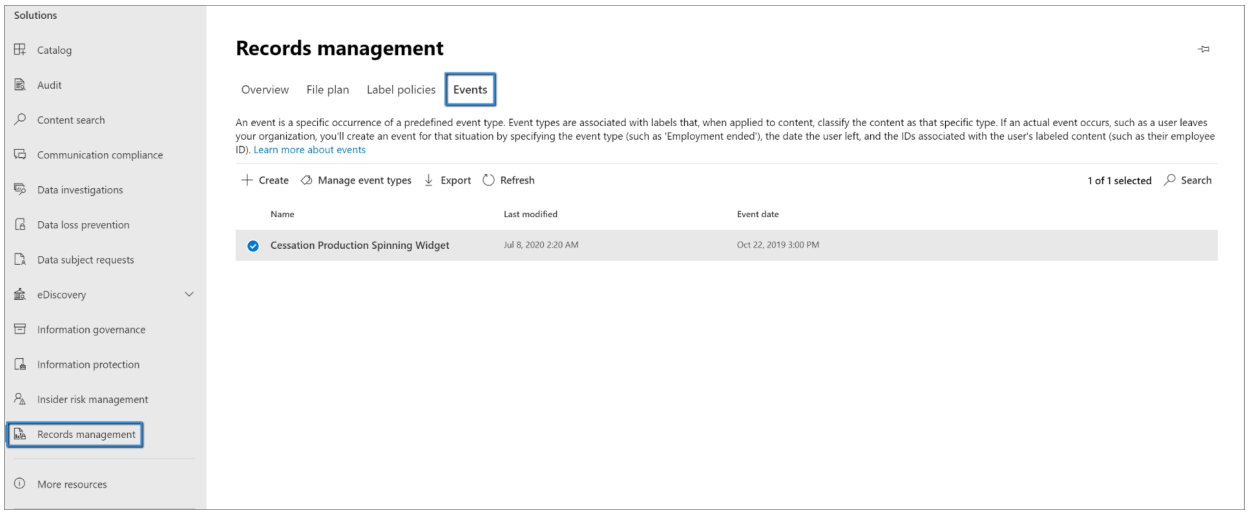

Select the event to view the details on the flyout page. Notice that even though the event is created, the event status shows that no SharePoint sites or documents have been processed.

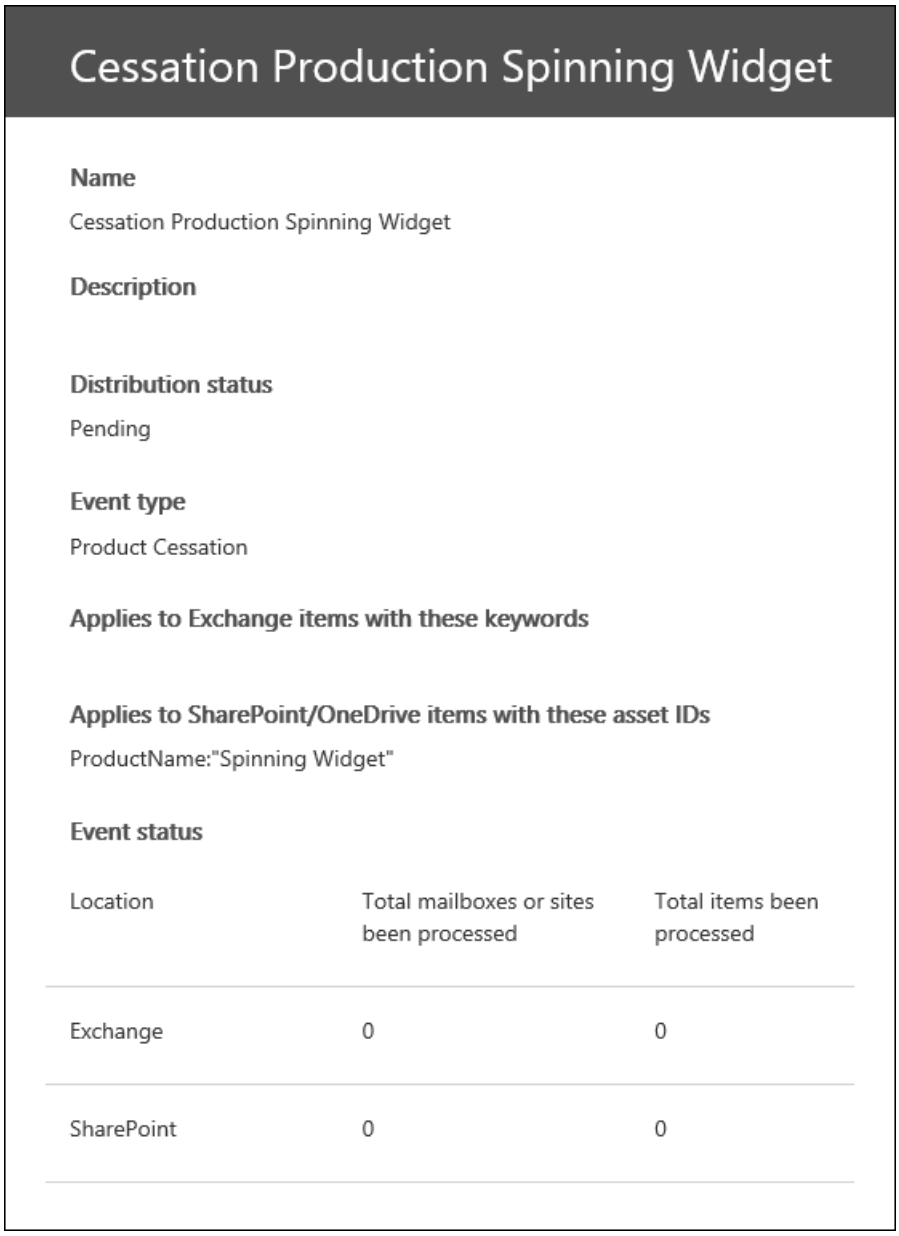

But after a delay, the event status shows that a SharePoint site and a SharePoint document have been processed.

This shows that the retention period for the label applied to the Spinning Widget product document has been initiated, based on the event date of the *Cessation Production Spinning Widget* event. Assuming that you implemented the scenario in your test environment by configuring a one-day retention period, you can go to the document library for your product documents a few days after the event was created and verify that the document was deleted (after the deletion job in SharePoint has run).

More about asset IDs

As the Start retention when an event occurs article explains, it's important to understand the relationship between event types, retention labels, events, and asset IDs. The asset ID is simply a document property in SharePoint and OneDrive. It helps you identify the documents whose retention period will be triggered by the event. By default, SharePoint has an **Asset Id** property that you can use for event-driven retention:

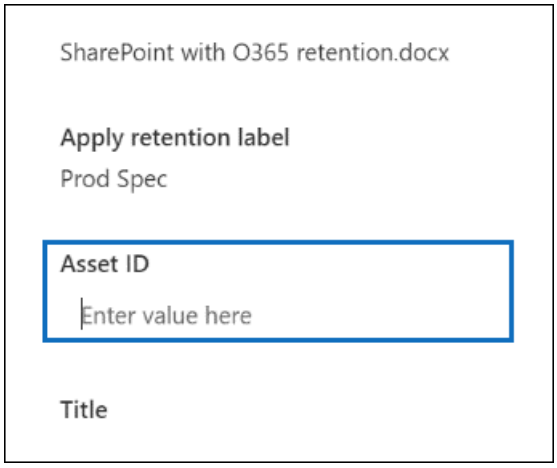

As the following screenshot shows, the asset ID managed property is called **ComplianceAssetId**.

Instead of using the default **Asset Id** property as we do in this scenario, you can use any other property. But it's important to understand that if you don't specify an asset ID or keywords for an event, all the content that has a label of that event type will get its retention period triggered by the event.

Using advanced search in SharePoint

In the previous screenshot, you can see that there's another managed property related to retention labels called **ComplianceTag** that's mapped to a crawled property. The **ComplianceAssetId** managed property is also mapped to a crawled property. This means that you can use these managed properties in advanced search to retrieve all documents that have been tagged with a retention label.

Disposition of content

2/18/2021 • 6 minutes to read • Edit Online

> *Microsoft 365 licensing guidance for security & compliance*.

Use the **Disposition** tab from **Records Management** in the Microsoft 365 compliance center to manage disposition reviews and view records that have been automatically deleted at the end of their retention period.

Prerequisites for viewing content dispositions

To manage disposition reviews and confirm that records have been deleted, you must have sufficient permissions and auditing must be enabled.

Permissions for disposition

To successfully access the **Disposition** tab in the Microsoft 365 compliance center, users must have the **Disposition Management** admin role. From December 2020, this role is now included in the **Records Management** default admin role group.

> **NOTE**
> By default, a global admin isn't granted the **Disposition Management** role.

To grant users just the permissions they need for disposition reviews without granting them permissions to view and configure other features for retention and records management, create a custom role group (for example, named "Disposition Reviewers") and grant this group the Disposition Management role.

Additionally, to view the contents of items during the disposition process, add users to the following two role groups: **Content Explorer Content Viewer** and **Content Explorer List Viewer**. If users don't have the permissions from these role groups, they can still select a disposition review action to complete the disposition review, but must do so without being able to view the item's contents from the compliance center.

For instructions to configure these permissions, see Give users access to the Office 365 Security & Compliance Center.

Enable auditing

Make sure that auditing is enabled at least one day before the first disposition action. For more information, see Search the audit log in the Office 365 Security & Compliance Center.

Disposition reviews

When content reaches the end of its retention period, there are several reasons why you might want to review that content to decide whether it can be safely deleted ("disposed"). For example, you might need to:

- Suspend the deletion of relevant content in the event of litigation or an audit.
- Remove content from the disposition list to store in an archive, if that content has research or historical value.
- Assign a different retention period to the content, perhaps because the original retention settings were a temporary or provisional solution.

- Return the content to clients or transfer it to another organization.

When a disposition review is triggered at the end of the retention period:

- The people you choose receive an email notification that they have content to review. These reviewers can be individual users or mail-enabled security groups. Note that notifications are sent on a weekly basis.
- The reviewers go to the **Disposition** tab in the Microsoft 365 compliance center to review the content and decide whether to permanently delete it, extend its retention period, or apply a different retention label.

A disposition review can include content in Exchange mailboxes, SharePoint sites, OneDrive accounts, and Microsoft 365 groups. Content awaiting a disposition review in those locations is deleted only after a reviewer chooses to permanently delete the content.

> **NOTE**
> A mailbox must have at least 10 MB data to support disposition reviews.

You can see an overview of all pending dispositions in the **Overview** tab. For example:

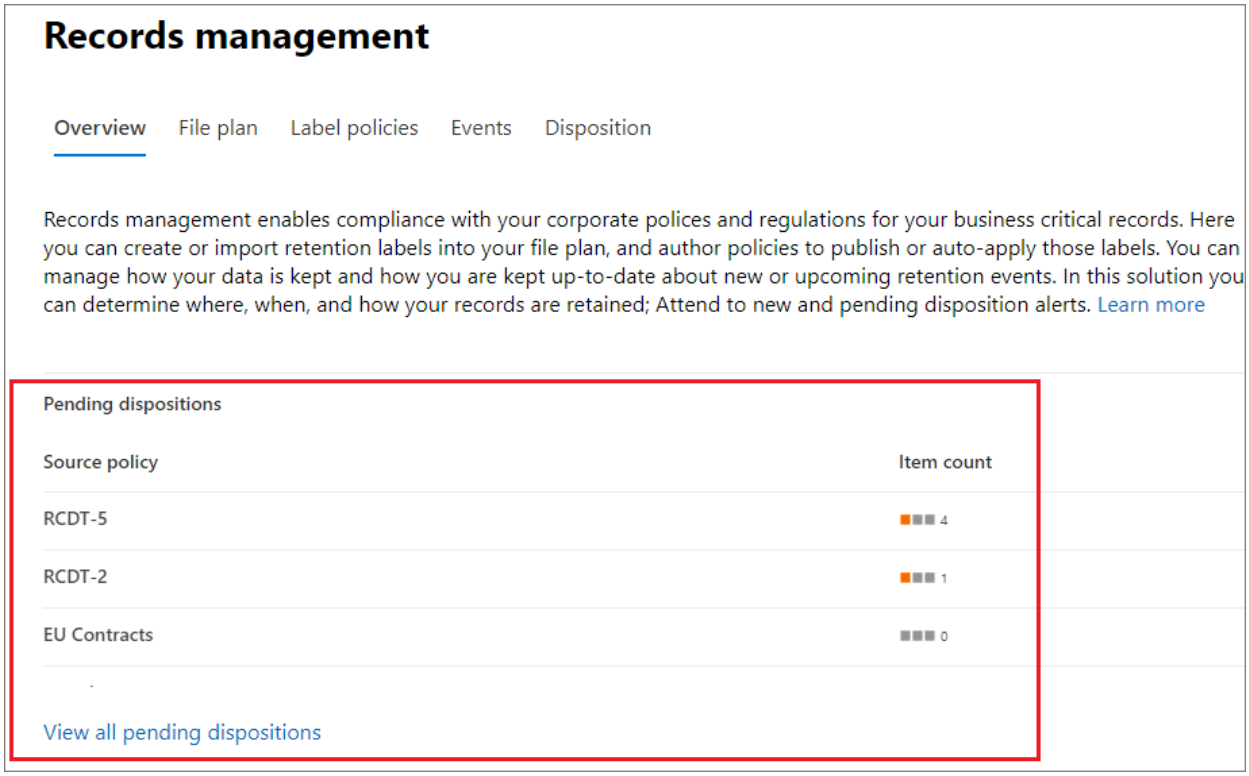

When you select the **View all pending dispositions**, you're taken to the **Disposition** page. For example:

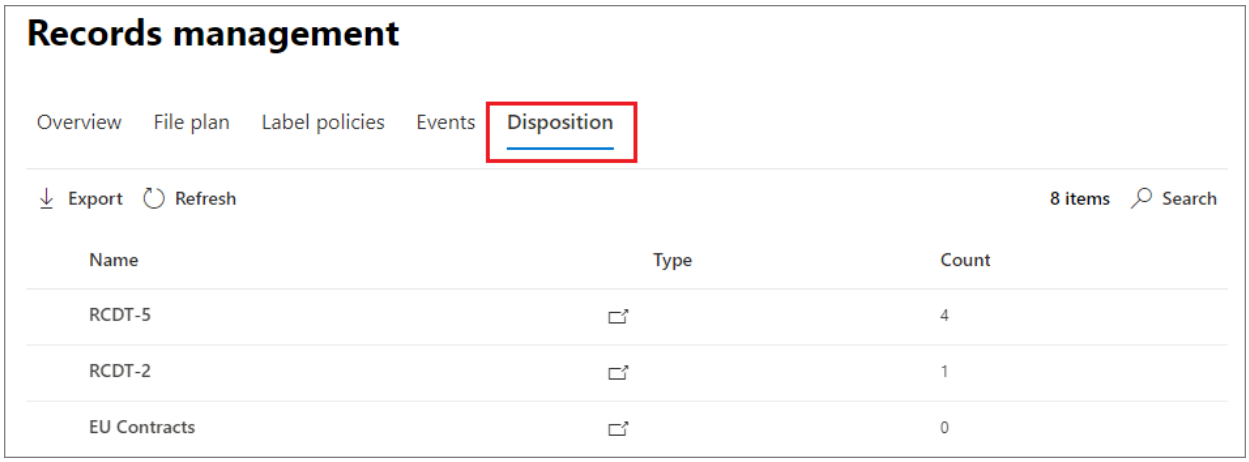

Workflow for a disposition review

The following diagram shows the basic workflow for a disposition review when a retention label is published and then manually applied by a user. Alternatively, a retention label configured for a disposition review can be auto-applied to content.

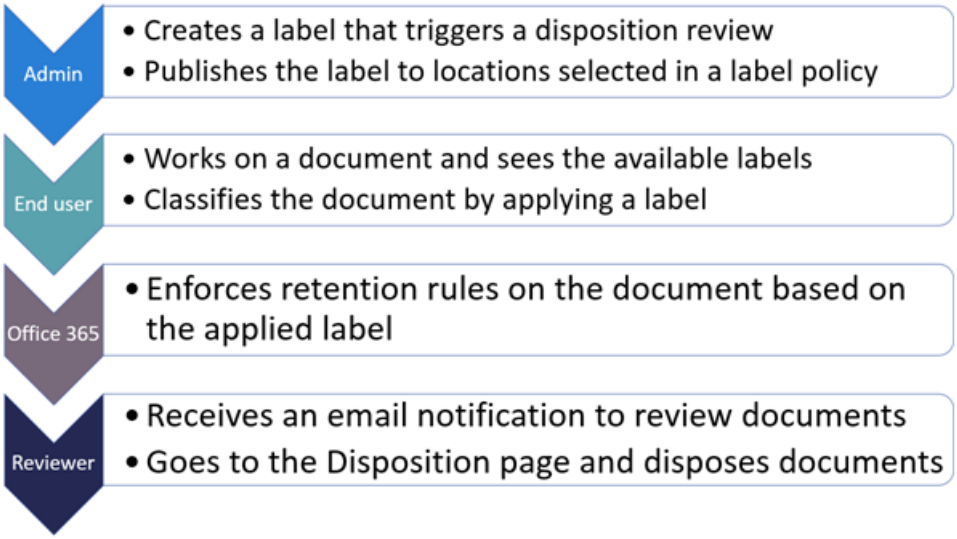

Triggering a disposition review at the end of the retention period is a configuration option that's available only with a retention label. This option is not available for a retention policy. For more information about these two retention solutions, see Learn about retention policies and retention labels.

From the **Define retention settings** page for a retention label:

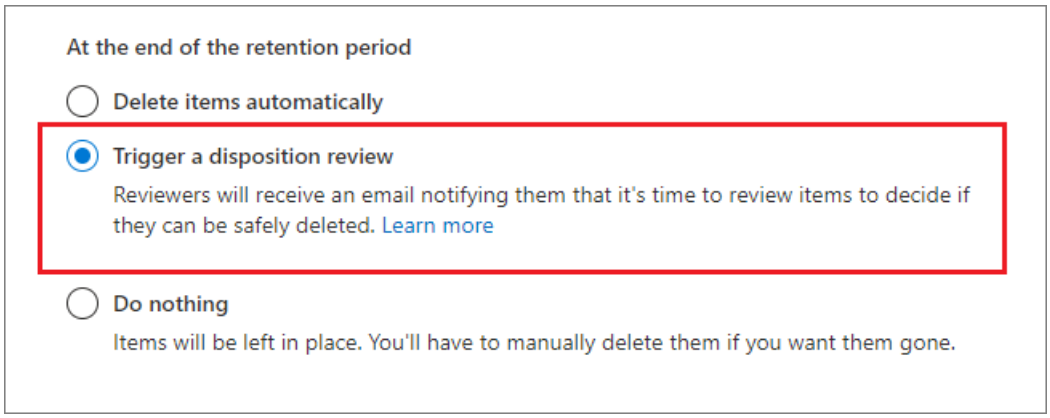

After you select this **Trigger a disposition review** option, you specify the disposition reviewers on the next page of the wizard:

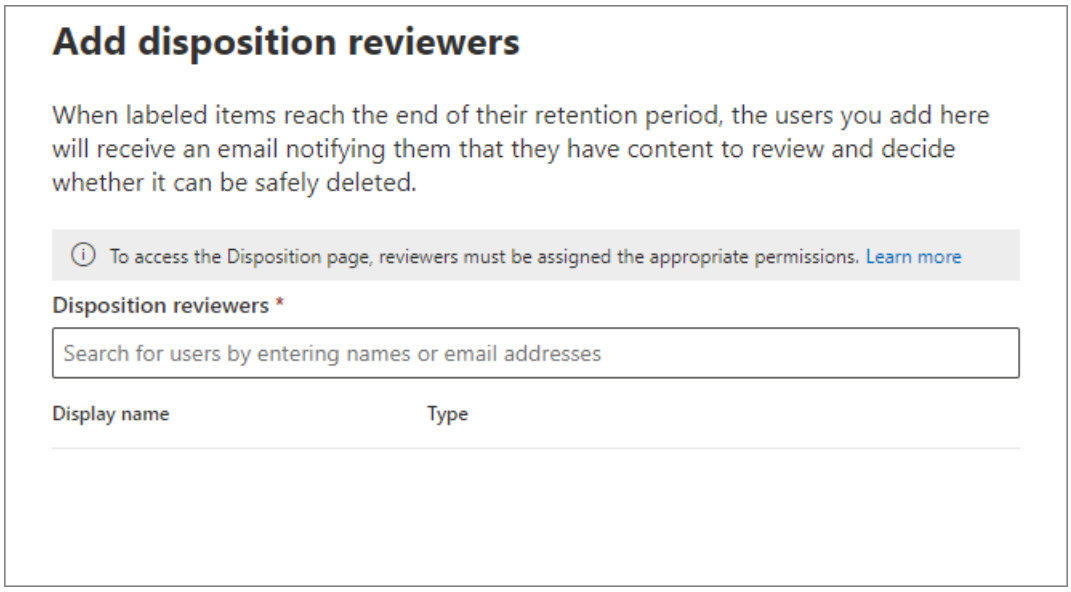

For the reviewers, specify a user or mail-enabled security group. Microsoft 365 groups (formerly Office 365 groups) are not supported for this option.

Viewing and disposing of content

When a reviewer is notified by email that content is ready to review, they go to the **Disposition** tab from **Records Management** in the Microsoft 365 compliance center. The reviewers can see how many items for each retention label are awaiting disposition, and then select a retention label to see all content with that label.

After you select a retention label, you then see all pending dispositions for that label from the **Pending disposition** tab. Select one or more items where you can then choose an action and enter a justification comment:

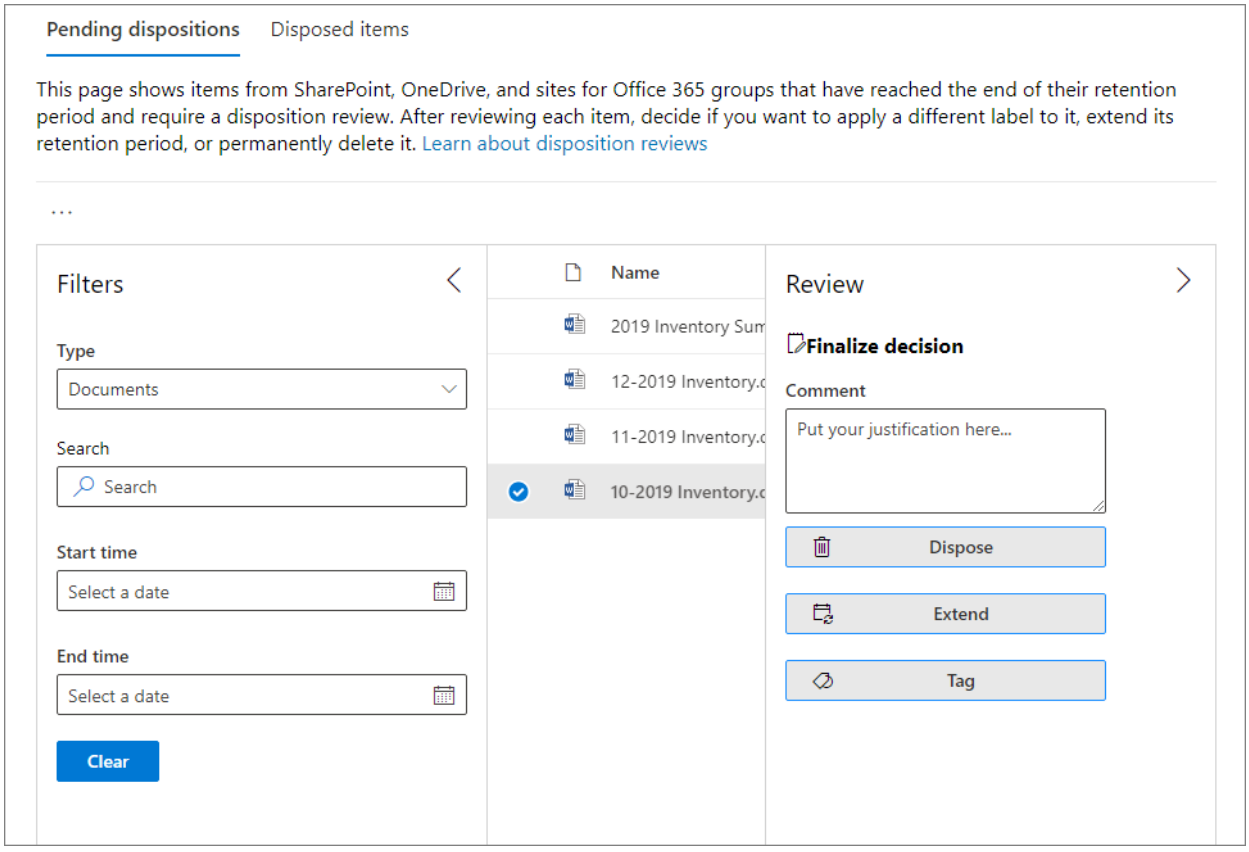

As you can see from the picture, the actions supported are:

- Permanently delete the item
- Extend the retention period

- Apply a different retention label

Providing you have permissions to the location and the content, you can use the link in the **Location** column to view documents in their original location. During a disposition review, the content never moves from its original location, and it's never deleted until the reviewer chooses to do so.

The email notifications are sent automatically to reviewers on a weekly basis. This scheduled process means that when content reaches the end of its retention period, it might take up to seven days for reviewers to receive the email notification that content is awaiting disposition.

All disposition actions can be audited and the justification text entered by the reviewer is saved and displayed in the **Comment** column on the **Disposed items** page.

How long until disposed content is permanently deleted

Content awaiting a disposition review is deleted only after a reviewer chooses to permanently delete the content. When the reviewer chooses this option, the content in the SharePoint site or OneDrive account becomes eligible for the standard cleanup process described in How retention settings work with content in place.

Disposition of records

Use the **Disposition** tab from the **Records Management** page to identify records that are now deleted, either automatically or after a disposition review. These items display **Records Disposed** in the **Type** column. For example:

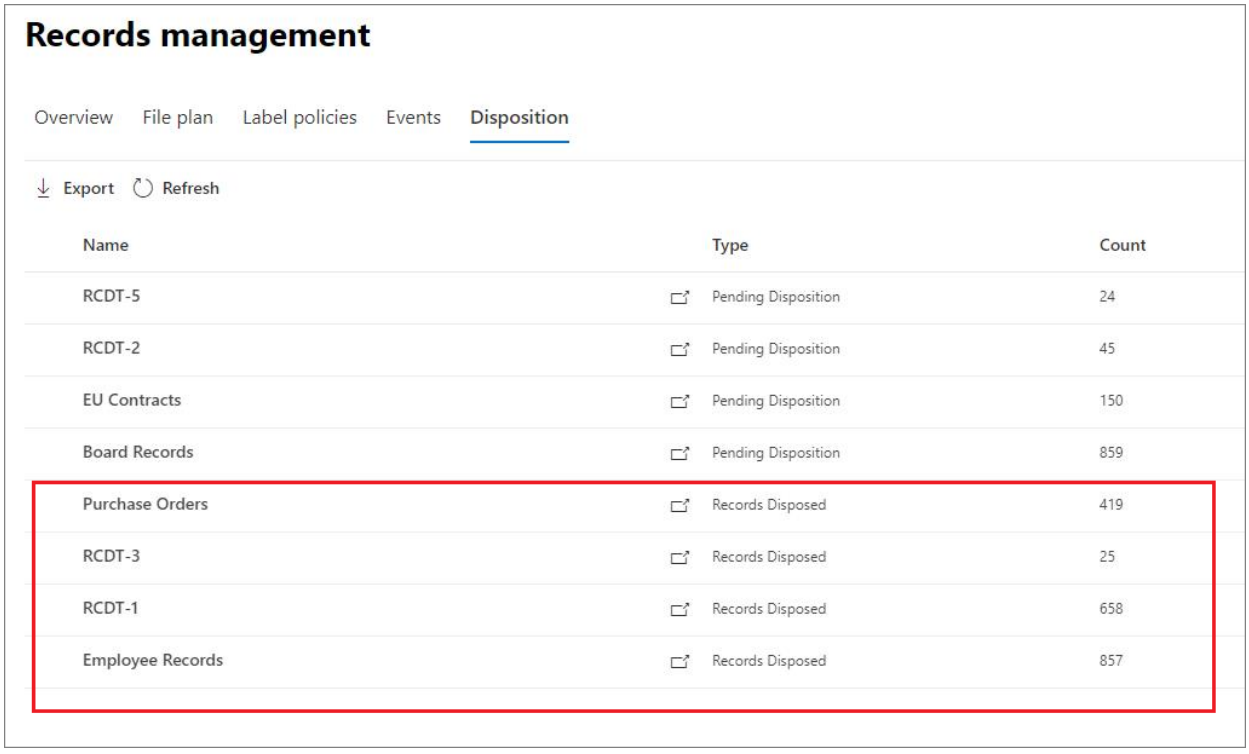

Items that are shown in the **Disposed Items** tab for record labels are kept for up to seven years after the item was disposed, with a limit of one million items per record for that period. If you see the **Count** number nearing this limit of one million, and you need proof of disposition for your records, contact Microsoft Support.

> **NOTE**
> This functionality is based on information from the unified audit log and therefore requires auditing to be enabled and searchable so the corresponding events are captured.

For auditing, search for **Deleted file marked as a record** in the **File and page activities** category. This audit

event is applicable to documents and emails.

Filter and export the views

When you select a retention label from the **Disposition** page, the **Pending disposition** tab (if applicable) and the **Disposed items** tab let you filter the views to help you more easily find items.

For pending dispositions, the time range is based on the expiration date. For disposed items, the time range is based on the deletion date.

You can export information about the items in either view as a .csv file that you can then sort and manage using Excel:

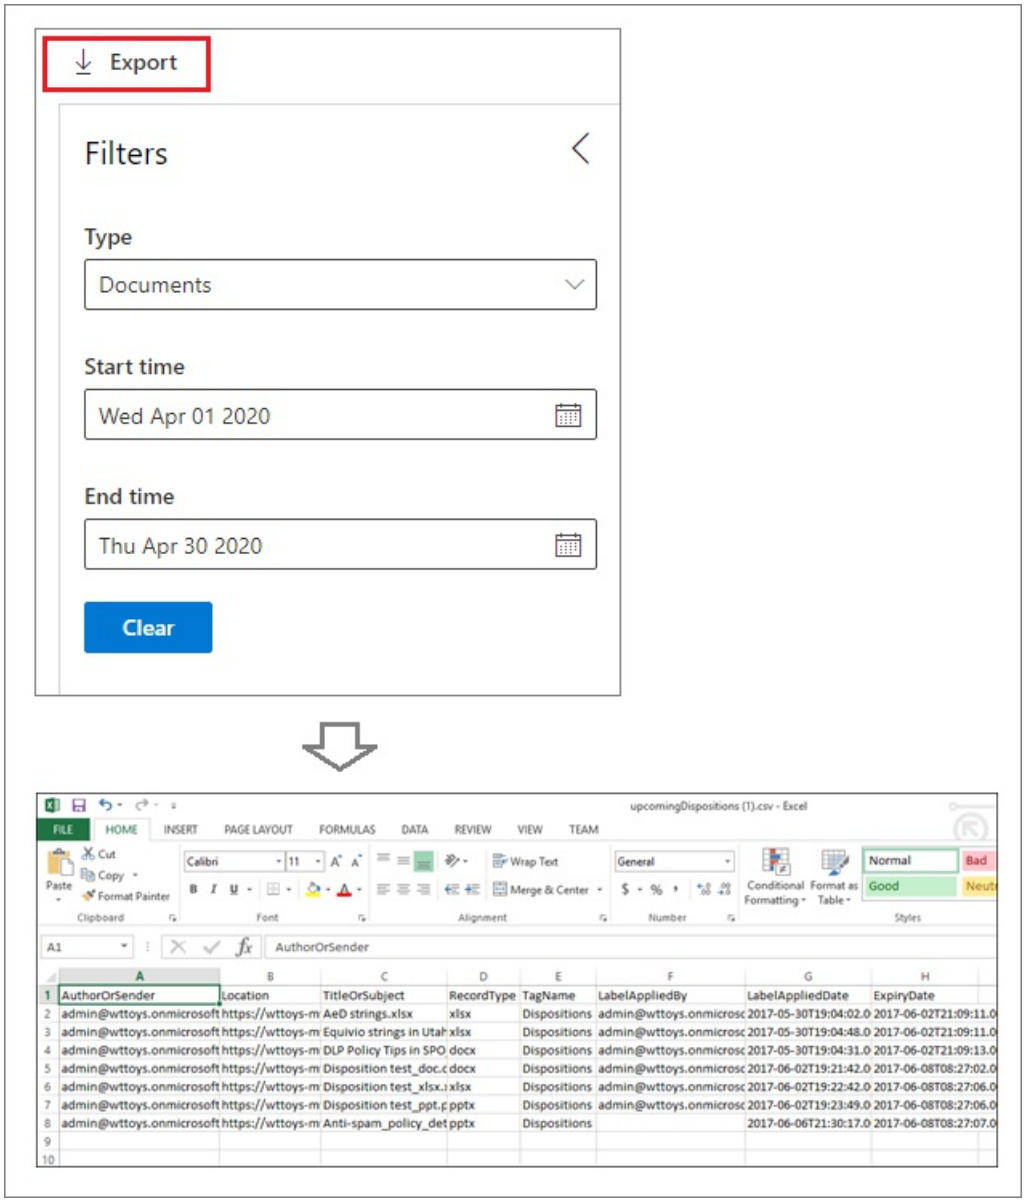

Overview of importing your organization's PST files

11/2/2020 • 20 minutes to read • Edit Online

NOTE

This article is for administrators. Are you trying to import PST files to your own mailbox? See Import email, contacts, and calendar from an Outlook .pst file.

You can use the Import service in the Security & Compliance Center to quickly bulk-import PST files to Exchange Online mailboxes in your organization. There are two ways you can import PST files to Office 365:

- **Network upload** - Upload the PST files over the network to a temporary Azure Storage location in the Microsoft cloud. Then you use the Office 365 Import service to import the PST data to mailboxes in your organization.
- **Drive shipping** - Copy the PST files to a BitLocker-encrypted hard drive and then physically ship the drive to Microsoft. When Microsoft receives the hard drive, data center personnel upload the data to a temporary Azure Storage location in the Microsoft cloud. Then you use the Office 365 Import service to import the data to mailboxes in your organization.

Step-by-step instructions

See one of the following topics for detailed, step-by-step instructions for bulk-importing your organization's PST files to Office 365.

- Use network upload to import PST files to Office 365
- Use drive shipping to import PST files

How importing PST files works

Here's an illustration and description of the complete PST import process. The illustration shows the primary workflow and highlights the differences between the network upload and drive shipping methods.

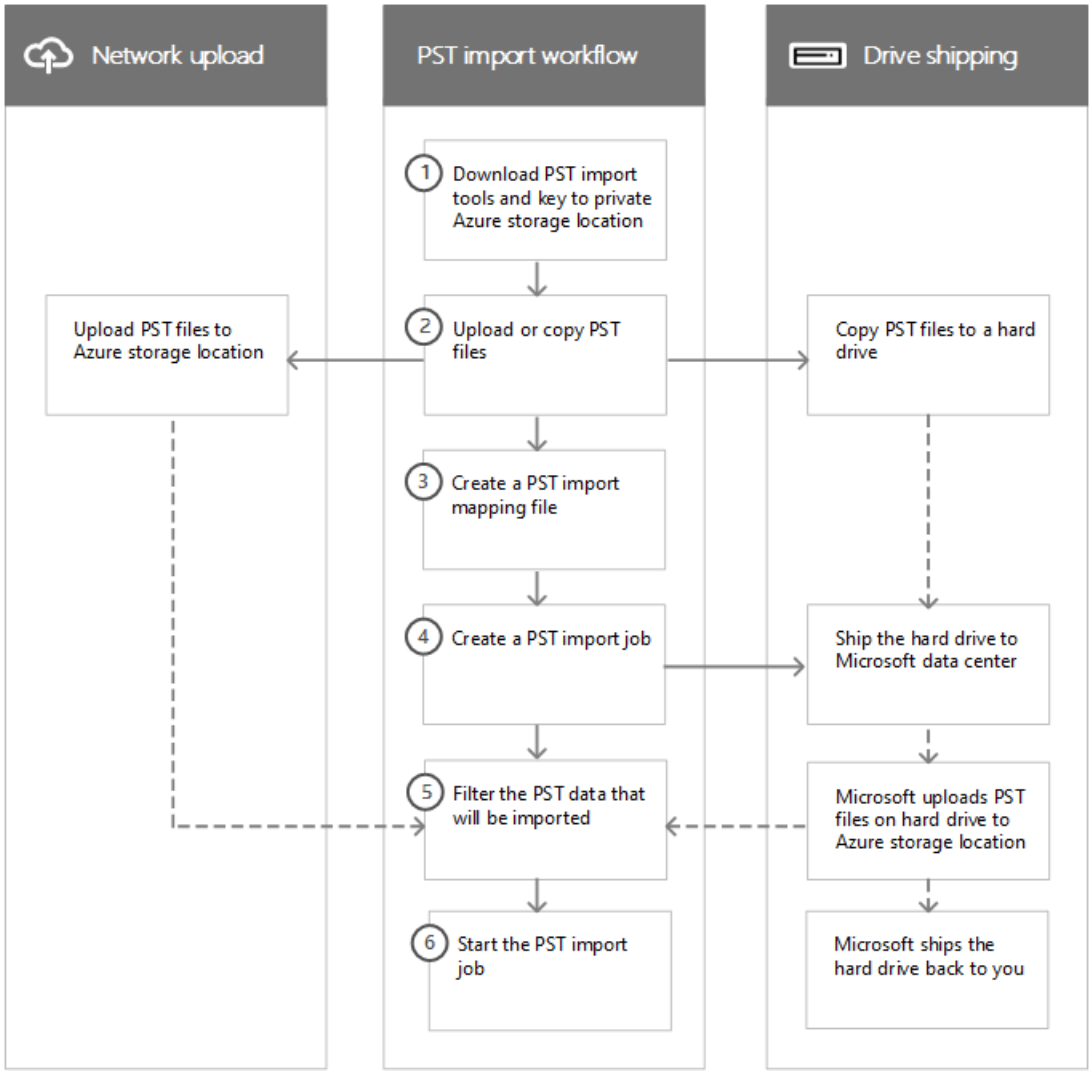

1. **Download the PST import tools and key to private Azure Storage location** - The first step is to download the tool and access key used to upload the PST files or copy them to a hard drive. You obtain these from the **Import** page in the Security & Compliance Center. The key provides you (or Microsoft data center personnel in the case of drive shipping) with the necessary permissions to upload PST files to a private and secure Azure Storage location. This access key is unique to your organization and helps prevent unauthorized access to your PST files after they're uploaded to the Microsoft cloud. Importing PST files to Microsoft 365 doesn't require your organization to have a separate Azure subscription.

2. **Upload or copy the PST files** - The next step depends on whether you're using network upload or drive shipping to import PST files. In both cases, you'll use the tool and secure storage key that you obtained in the previous step.

   - **Network upload:** The AzCopy.exe tool (downloaded in step 1) is used to upload and store your PST files in an Azure Storage location in the Microsoft cloud. The Azure Storage location that you upload your PST files to is located in the same regional Microsoft datacenter as your organization.

     To upload them, the PST files that you want to import have to be located in a file share or file server in your organization.

   - **Drive shipping:** The WAImportExport.exe tool (downloaded in step 1) is used to copy your PST files to the hard drive. This tool encrypts the hard drive with BitLocker and then copies the PSTs to the hard drive. Like network upload, the PST files that you want to copy to the hard drive have to be located in a file share or file server in your organization.

3. **Create a PST import mapping file** - After the PST files have been uploaded to the Azure Storage location or copied to a hard drive, the next step is to create a comma-separated value (CSV) file that specifies which user mailboxes the PST files will be imported to (and a PST file can be imported to a

user's primary mailbox or their archive mailbox). The Office 365 Import service will use the information to import the PST files.

4. **Create a PST import job** - The next step is to create a PST import job on the **Import PST files** page in the Security & Compliance Center and submit the PST import mapping file created in the previous step. For network upload (because the PST files have been uploaded to Azure) Microsoft 365 analyzes the data in the PST files and then gives you an opportunity to set filters that control what data actually gets imported to the mailboxes specified in the PST import mapping file.

   For drive shipping, a few additional things happen at this point in the process.

   - You physically ship the hard drive to a Microsoft data center (the shipping address for the Microsoft data center is displayed when the import job is created).
   - When Microsoft receives the hard drive, data center personnel will upload the PST files on the hard drive to the Azure Storage location for your organization. As previously explained, your PST files are uploaded to a Azure Storage location that resides in the same regional Microsoft datacenter where your organization is located.

     > **NOTE**
     > The PST files on the hard drive are uploaded to Azure within 7 to 10 business days after Microsoft receives the hard drive.

     Like the network upload process, Microsoft 365 then analyzes the data in the PST files and gives you an opportunity to set filters that control what data actually gets imported to the mailboxes specified in the PST import mapping file.
   - Microsoft ships the hard drive back to you.

5. **Filter the PST data that will be imported to mailboxes** - After the import job is created (and after the PST files from a drive shipping job are uploaded to the Azure Storage location) Microsoft 365 analyzes the data in the PST files (safely and securely) by identifying the age of the items and the different message types included in the PST files. When the analysis is completed and the data is ready to import, you have the option to import all the data contained in the PST files or you can trim the data that's imported by setting filters that control what data gets imported.

6. **Start the PST import job** - After the import job is started, Microsoft 365 uses the information in the PST import mapping file to import the PSTs files from the he Azure Storage location to user mailboxes. Status information about the import job (including information about each PST file being imported) is displayed on the **Import PST files** page in the Security & Compliance Center. When the import job is finished, the status for the job is set to **Complete**.

Why import email data to Microsoft 365?

- It's a good way to import your organization's archival messaging data to Microsoft 365.
- You can use the Intelligent Import feature to filter the items in PST files that actually get imported to the target mailboxes. This lets you trim the data that's imported by setting filters that control what data gets imported.
- Importing email data to Microsoft 365 helps address compliance needs of your organization by letting you:
  - Enable archive mailboxes and unlimited archiving to give users additional mailbox storage space.
  - Place mailboxes on Litigation Hold to retain content.

    - Use the Content Search tool to search for mailbox content.
    - Use eDiscovery cases to manage your organization's legal investigations
    - Use retention policies in the Security & Compliance Center to control how long mailbox content is retained, and then delete content after the retention period expires.
    - Use Communication compliance policies to examine messages to make sure they are compliant with message standards and add a classification type.
- Importing data to Microsoft 365 helps protect against data loss. Email data that's imported to Microsoft 365 inherits the high availability features of Exchange Online.
- Email data is available to users from all devices because it's stored in the cloud.

Importing SharePoint data to Microsoft 365

You can also import files and documents to SharePoint sites and OneDrive accounts in your organization. For more information, see the following articles:

- Migrate to SharePoint Online
- Introducing the SharePoint Migration Tool
- Migrate to SharePoint Online using PowerShell
- Migrate your file share content to SharePoint Online using the Azure Data Box

Frequently asked questions about importing PST files

Here are some frequently asked questions about using the Office 365 Import service to bulk-import PST files to Microsoft 365 mailboxes.

- Using network upload to import PST files
- Using drive shipping to import PST files

Using network upload to import PST files

What permissions are required to create import jobs in the Office 365 Import Service?

You have to be assigned the Mailbox Import Export role in Exchange Online to import PST files to Microsoft 365 mailboxes. By default, this role isn't assigned to any role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group. Or you can create a new role group, assign the Mailbox Import Export role, and then add yourself or other users as a member. For more information, see the "Add a role to a role group" or the "Create a role group" sections in Manage role groups in Exchange Online.

Additionally, to create import jobs in the Security & Compliance Center, one of the following must be true:

- You have to be assigned the Mail Recipients role in Exchange Online. By default, this role is assigned to the Organization Management and Recipient Management roles groups.

  Or

- You have to be a global administrator in your organization.

> **TIP**
> Consider creating a new role group in Exchange Online that's specifically intended for importing PST files to Office 365. For the minimum level of privileges required to import PST files, assign the Mailbox Import Export and Mail Recipients roles to the new role group, and then add members.

Where is network upload available?

Network upload is currently available in these regions: United States, Canada, Brazil, the United Kingdom, France, Germany, Switzerland, Norway, Europe, India, East Asia, Southeast Asia, Japan, Republic of Korea, Australia, and United Arab Emirates (UAE). Network upload will be available in more regions soon.

What is the pricing for importing PST files by using network upload?

Using network upload to import PST files is free.

This also means that after PST files are deleted from the Azure Storage area, they're no longer displayed in the list of files for a completed import job in the Microsoft 365 admin center. Although an import job might still be listed on the **Import data to Office 365** page, the list of PST files might be empty when you view the details of older import jobs.

What version of the PST file format is supported for importing to Office 365?

There are two versions of the PST file format: ANSI and Unicode. We recommend importing files that use the Unicode PST file format. However, files that use the ANSI PST file format, such as those for languages that use a double-byte character set (DBCS), can also be imported to Office 365. For more information about importing ANSI PST files, see Step 4 in Use network upload to import PST files to Office 365.

Additionally, PST files from Outlook 2007 and later versions can be imported to Office 365.

After I upload my PST files to the Azure Storage area, how long are they kept in Azure before they're deleted?

When you use the network upload method to import PST files, you upload them to an Azure blob container named `ingestiondata`. If there are no import jobs in progress on the **Import PST files** page in the Security & Compliance Center), then all PST files in the `ingestiondata` container in Azure are deleted 30 days after the most recent import job was created in the Security & Compliance Center. That also means you have to create a new import job in the Security & Compliance Center (described in Step 5 in the network upload instructions) within 30 days of uploading PST files to Azure.

This also means that after PST files are deleted from the Azure Storage area, they're no longer displayed in the list of files for a completed import job in the Security & Compliance Center. Although an import job might still be listed on the **Import PST files** page in the Security & Compliance Center, the list of PST files might be empty when you view the details of older import jobs.

How long does it take to import a PST file to a mailbox?

It depends on the capacity of your network, but it typically takes several hours for each terabyte (TB) of data to be uploaded to the Azure Storage area for your organization. After the PST files are copied to the Azure Storage area, a PST file is imported to a Microsoft 365 mailbox at a rate of at least 24 GB per day. If this rate doesn't meet your needs, you might consider other methods to get email data into Office 365. For more information, see Ways to migrate multiple email accounts to Office 365.

If different PST files are imported to different target mailboxes, the import process occurs in parallel; in other words, each PST/mailbox pair is imported simultaneously. Likewise, if multiple PST files are imported to the same mailbox, they will be simultaneously imported.

How does the PST import process handle duplicate email items?

The PST import process checks for duplicate items and doesn't copy the items from a PST file to the mailbox or archive if a matching item exists in the target folder in the target mailbox or target archive. If you reimport the same PST file and specify a different target folder (using the TargetRootFolder property in the PST import mapping file) than the one you specified in a previous import job, all items in the PST file will be reimported.

Is there a message size limit when importing PST files?

Yes. If a PST file contains a mailbox item that is larger than 150 MB, the item will be skipped during the import process.

Are message properties, such as when the message was sent or received, the list of recipients and other properties, preserved when PST files are imported to a Microsoft 365 mailbox?

Yes. The original message metadata isn't changed during the import process.

Is there a limit to the number of levels in a folder hierarchy for a PST file that I want to import to a mailbox?

Yes. You can't import a PST file that has 300 or more levels of nested folders.

Can I use network upload to import PST files to an inactive mailbox in Office 365?

Yes, this capability is now available.

Can I use network upload to import PST files to an online archive mailbox in an Exchange hybrid deployment?

Yes, this capability is now available.

Can I use network upload to import PST files to public folders in Exchange Online?

No, you can't import PST files to public folders.

Using drive shipping to import PST files

What permissions are required to create import jobs in the Office 365 Import Service?

You have to be assigned the Mailbox Import Export role to import PST files to Microsoft 365 mailboxes. By default, this role isn't assigned to any role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group. Or you can create a new role group, assign the Mailbox Import Export role, and then add yourself or other users as a member. For more information, see the "Add a role to a role group" or the "Create a role group" sections in Manage role groups in Exchange Online.

Additionally, to create import jobs in the Security & Compliance Center, one of the following must be true:

- You have to be assigned the Mail Recipients role in Exchange Online. By default, this role is assigned to the Organization Management and Recipient Management roles groups.

  Or

- You have to be a global administrator in your organization.

> **TIP**
> Consider creating a new role group in Exchange Online that's specifically intended for importing PST files to Office 365. For the minimum level of privileges required to import PST files, assign the Mailbox Import Export and Mail Recipients roles to the new role group, and then add members.

Where is drive shipping available?

Drive shipping is currently available in the United States, Canada, Brazil, the United Kingdom, Europe, India, East Asia, Southeast Asia, Japan, Republic of Korea, and Australia. Drive shipping will be available in more regions soon.

> **NOTE**
> At this time, drive shipping to import PST files is not available in Germany and Switzerland. This FAQ will be updated when drive shipping is available in these countries.

What commercial licensing agreements support drive shipping?

Drive shipping to import PST files to Microsoft 365 is available through a Microsoft Enterprise Agreement (EA). Drive shipping isn't available through a Microsoft Products and Services Agreement (MPSA).

What is the pricing for using drive shipping to import PST files to Microsoft 365?

The cost to use drive shipping to import PST files to Microsoft 365 mailboxes is $2 USD per GB of data. For example, if you ship a hard drive that contains 1,000 GB (1 TB) of PST files, the cost is $2,000 USD. You can work with a partner to pay the import fee. For information about finding a partner, see Find your Microsoft partner or reseller.

What kind of hard drives are supported for drive shipping?

Only 2.5-inch solid-state drives (SSDs) or 2.5 inch or 3.5 inch SATA II/III internal hard drives are supported for use with the Office 365 Import service. You can use hard drives up to 10 TB. For import jobs, only the first data volume on the hard drive will be processed. The data volume must be formatted with NTFS. When copying data to a hard drive, you can attach it directly using a 2.5 inch SSD or 2.5 inch or 3.5 inch SATA II/III connector or you can attach it externally using an external 2.5 inch SSD or 2.5 inch or 3.5 inch SATA II/III USB adaptor.

> **IMPORTANT**
> External hard drives that come with an built-in USB adaptor aren't supported by the Office 365 Import service. Additionally, the disk inside the casing of an external hard drive can't be used. Please don't ship external hard drives.

How many hard drives can I ship for a single import job?

You can ship a maximum of 10 hard drives for a single import job.

After I ship my hard drive, how long does it take to get to the Microsoft datacenter?

That depends on a few things, such as your proximity to the Microsoft data center and what kind of shipping option you used to ship your hard drive (such as, next-day delivery, two-day delivery, or ground-delivery). With most shippers, you can use the tracking number to track the status of your delivery.

After my hard drive arrives at the Microsoft datacenter, how long does it take to upload my PST files to Azure?

After your hard drive is received at the Microsoft data center, it will take between 7 to 10 business days to upload the PST files to the Azure Storage location for your organization. The PST files will be uploaded to an Azure blob container named `ingestiondata`.

How long does it take to import a PST file to a mailbox?

After the PST files are uploaded to the Azure Storage area, Microsoft 365 analyzes the data in the PST files (in a safe and secure manner) to identify the age of the items and the different message types included in the PST files. When this analysis is complete, you'll have the option to import all the data in the PST files or set filters to that control what data gets imported. After you start the import job, a PST file is imported to a Microsoft 365

mailbox at a rate of at least 24 GB per day. If this rate doesn't meet your needs, you might consider other methods to get email data into Microsoft 365. For more information, see Ways to migrate multiple email accounts to Microsoft 365.

If different PST files are imported to different target mailboxes, the import process occurs in parallel; in other words, each PST/mailbox pair is imported simultaneously. Likewise, if multiple PST files are imported to the same mailbox, they will be simultaneously imported.

After Microsoft uploads my PST files to Azure, how long are they kept in Azure before they're deleted?

All PST files in the Azure Storage location for your organization (in blob container named `ingestiondata`), are deleted 30 days after the most recent import job was created on the **Import PST files** page in the Security & Compliance Center.

This also means that after PST files are deleted from the Azure Storage area, they're no longer displayed in the list of files for a completed import job in the Security & Compliance Center. Although an import job might still be listed on the **Import PST files** page in the Security & Compliance Center, the list of PST files might be empty when you view the details of older import jobs.

What version of the PST file format is supported for importing to Microsoft 365?

There are two versions of the PST file format: ANSI and Unicode. We recommend importing files that use the Unicode PST file format. However, files that use the ANSI PST file format, such as those for languages that use a double-byte character set (DBCS), can also be imported to Microsoft 365. For more information about importing ANSI PST files, see Step 3 in Use drive shipping to import your organization PST files to Microsoft 365.

Additionally, PST files from Outlook 2007 and later versions can be imported to Microsoft 365.

Is there a message size limit when importing PST files?

Yes. If a PST file contains a mailbox item that is larger than 150 MB, the item will be skipped during the import process.

How does the PST import process handle duplicate email items?

The PST import process checks for duplicate items and doesn't copy the items from a PST file to the mailbox or archive if a matching item exists in the target folder in the target mailbox or target archive. If you reimport the same PST file and specify a different target folder (using the TargetRootFolder property in the PST import mapping file) than the one you specified in a previous import job, all items in the PST file will be reimported.

Are message properties, such as when the message was sent or received, the list of recipients and other properties, preserved when PST files are imported to a Microsoft 365 mailbox?

Yes. The original message metadata isn't changed during the import process

Is there a limit to the number of levels in a folder hierarchy for a PST file that I want to import to a mailbox?

Yes. You can't import a PST file that has 300 or more levels of nested folders.

Can I use drive shipping to import PST files to an inactive mailbox in Microsoft 365?

Yes, this capability is now available.

Can I use drive shipping to import PST files to an online archive mailbox in an Exchange hybrid deployment?

Yes, this capability is now available.

Can I use drive shipping to import PST files to public folders in Exchange Online?

No, you can't import PST files to public folders.

Can Microsoft wipe my hard drive before they ship it back to me?

No, Microsoft can't wipe hard drives before shipping them back to customers. Hard drives are returned to you in the same state they were in when they were received by Microsoft.

Can Microsoft shred my hard drive instead of shipping it back to me?

No, Microsoft can't destroy your hard drive. Hard drives are returned to you in the same state they were in when they were received by Microsoft.

What courier services are supported for return shipping?

If you're a customer in the United States or Europe, Microsoft uses FedEx to return your hard drive. For all other regions, Microsoft uses DHL.

What are the return shipping costs?

Return shipping costs vary, depending on your proximity to the Microsoft data center that you shipped your hard drive to. Microsoft will bill your FedEx or DHL account to return your hard drive. The cost of return shipping is your responsibility.

Can I use a custom courier shipping service, such as FedEx Custom Shipping, to ship my hard drive to Microsoft?

Yes.

If I have to ship my hard drive to another country, is there anything I need to do?

The hard drive that you ship to Microsoft might have to cross international borders. If this is the case, you're responsible for ensuring that the hard drive and the data it contains are imported and/or exported in accordance with the applicable laws. Before shipping a hard drive, check with your advisors to verify that your drive and data can legally be shipped to the specified Microsoft data center. This will help to ensure that it reaches Microsoft in a timely manner.

Use network upload to import your organization's PST files to Microsoft 365

11/2/2020 • 29 minutes to read • Edit Online

> **NOTE**
> This article is for administrators. Are you trying to import PST files to your own mailbox? See Import email, contacts, and calendar from an Outlook .pst file

Here are the step-by-step instructions required to use network upload to bulk-import multiple PST files to Microsoft 365 mailboxes. For frequently asked questions about using network upload to bulk-import PST files to Microsoft 365 mailboxes, see FAQs for using network upload to import PST files.

You have to perform Step 1 only once to import PST files to Microsoft 365 mailboxes. After you perform these steps, follow Step 2 through Step 6 each time you want to upload and import a batch of PST files.

Before you import PST files

- You have to be assigned the Mailbox Import Export role in Exchange Online to import PST files to Microsoft 365 mailboxes. By default, this role isn't assigned to any role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group. Or you can create a role group, assign the Mailbox Import Export role, and then add yourself as a member. For more information, see the "Add a role to a role group" or the "Create a role group" sections in Manage role groups.

  Also, to create import jobs in the Security & Compliance Center, one of the following must be true:

  - You have to be assigned the Mail Recipients role in Exchange Online. By default, this role is assigned to the Organization Management and Recipient Management roles groups.

    Or

  - You have to be a global administrator in your organization.

  > **TIP**
  > Consider creating a new role group in Exchange Online that's specifically intended for importing PST files. For the minimum level of privileges required to import PST files, assign the Mailbox Import Export and Mail Recipients roles to the new role group, and then add members.

- The only supported method for importing PST files to Microsoft 365 is to use the AzCopy tool, as described in this topic. You can't use the Azure Storage Explorer to upload PST files directly to the Azure Storage area.
- You need to store the PST files that you want to import to Microsoft 365 on a file server or shared folder in your organization. In Step 2, you run the AzCopy tool to upload the PST files that are stored on a file server or shared folder to Microsoft 365.
- Large PST files may impact the performance of the PST import process. So we recommend that each PST file you upload to the Azure Storage location in Step 2 should be no larger than 20 GB.
- This procedure involves copying and saving a copy of a URL that contains an access key. This information will be used in Step 2 to upload your PST files, and in Step 3 if you want to view a list of the PST files uploaded to Office 365. Be sure to take precautions to protect this URL like you would protect passwords or other security-related information. For example, you might save it to a password-protected Microsoft Word document or to an encrypted USB drive. See the More information section for an example of this combined URL and key.
- You can import PST files to an inactive mailbox in Office 365. You do this by specifying the GUID of the inactive mailbox in the `Mailbox` parameter in the PST Import mapping file. See Step 4 on the **Instructions** tab in this topic for information.
- In an Exchange hybrid deployment, you can import PST files to a cloud-based archive mailbox for a user whose primary mailbox is on-premises. You do this by doing the following in the PST Import mapping file:

  - Specify the email address for the user's on-premises mailbox in the `Mailbox` parameter.
  - Specify the **TRUE** value in the `IsArchive` parameter.

    See Step 4 for more information.

- After PST files are imported, the retention hold setting for the mailbox is turned on for an indefinite duration. This means that the retention policy assigned to the mailbox won't be processed until you turn off the retention hold or set a date to turn off the hold. Why do we do this? If messages imported to a mailbox are old, they might be permanently deleted (purged) because their retention period has expired

based on the retention settings configured for the mailbox. Placing the mailbox on retention hold gives the mailbox owner time to manage these newly imported messages or give you time to change the retention settings for the mailbox. See the More information section in this topic for suggestions about managing the retention hold.

- By default, the maximum message size that can be received by a Microsoft 365 mailbox is 35 MB. That's because the default value for the *MaxReceiveSize* property for a mailbox is set to 35 MB. However, the limit for the maximum message receive size in Microsoft 365 is 150 MB. So if you import a PST file that contains an item larger than 35 MB, the Office 365 Import service we will automatically change the value of the *MaxReceiveSize* property on the target mailbox to 150 MB. This allows messages up to 150 MB to be imported to user mailboxes.

  TIP
  To identify the message receive size for a mailbox, you can run this command in Exchange Online PowerShell: `Get-Mailbox <user mailbox> | FL MaxReceiveSize`.

- For a high-level overview of the PST Import process, see How the import process works section in this article.

Step 1: Copy the SAS URL and install AzCopy

The first step is to download and install the AzCopy tool, which is the tool that you run in Step 2 to upload PST files to Office 365. You also copy the SAS URL for your organization. This URL is a combination of the network URL for the Azure Storage location in the Microsoft cloud for your organization and a Shared Access Signature (SAS) key. This key provides you with the necessary permissions to upload PST files to your Azure Storage location. Be sure to take precautions to protect the SAS URL. It's unique to your organization and will be used in Step 2.

IMPORTANT
To import PST files using the network upload method and command syntax documented in this article, you must use the version of AzCopy that can be downloaded in step 6b in the following procedure. You can also download that same version of AzCopy from here. Using a different version of AzCopy isn't supported.

1. Go to https://protection.office.com and sign in using the credentials for an administrator account in your organization.
2. In the left pane of the Security & Compliance Center, click **Information governance** > **Import** > **Import PST files**.

   NOTE
   You have to be assigned the appropriate permissions to access the **Import** page in the Security & Compliance Center. See the **Before you begin** section for more information.

3. On the **Import PST files** page, click **+ New import job**.

   The import job wizard is displayed.

4. Type a name for the PST import job, and then click **Next**. Use lowercase letters, numbers, hyphens, and underscores. You can't use uppercase letters or include spaces in the name.
5. On the **Do you want to upload or ship data?** page, click **Upload your data** and then click **Next**.

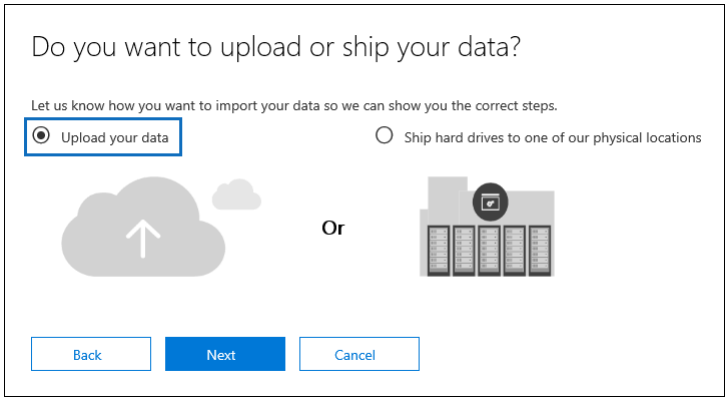

6. On the **Import data** page, do the following two things:

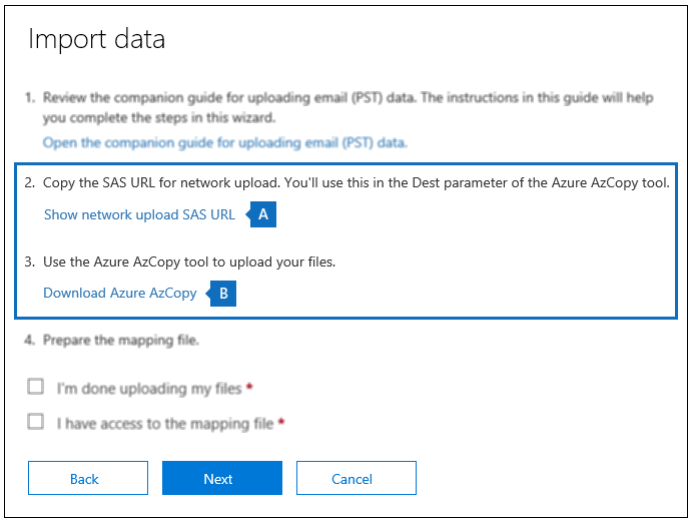

a. In step 2, click **Show network upload SAS URL**. After the SAS URL is displayed, click **Copy to clipboard** and then paste it and save it to a file so you can access it later.

b. In step 3, click **Download Azure AzCopy** to download and install the AzCopy tool. In the pop-up window, click **Run** to install AzCopy.

> **NOTE**
> You can leave the **Import data** page open (in case you need to copy the SAS URL again) or click **Cancel** to close it.

Step 2: Upload your PST files to Office 365

Now you're ready to use the AzCopy.exe tool to upload PST files to Office 365. This tool uploads and stores them in an Azure Storage location in the Microsoft cloud. As previously explained, the Azure Storage location that you upload your PST files is located in the same regional Microsoft datacenter where your organization is located. To complete this step, the PST files have to be located in a file share or file server in your organization. This is known as the source directory in this procedure. Each time you run the AzCopy tool, you can specify a different source directory.

> **NOTE**
> As previously stated, each PST file that you upload to the Azure Storage location should be no larger than 20 GB. PST files larger than 20 GB may impact the performance of the PST import process that you start in Step 6.

1. Open a Command Prompt on your local computer.
2. Go to the directory where you installed the AzCopy.exe tool in Step 1. If you installed the tool in the default location, go to `%ProgramFiles(x86)%\Microsoft SDKs\Azure\AzCopy`.
3. Run the following command to upload the PST files to Office 365.

```
AzCopy.exe /Source:<Location of PST files> /Dest:<SAS URL> /V:<Log file location> /Y
```

> **IMPORTANT**
> You must specify a directory as the source location in the previous command; you can't specify an individual PST file. All PST files in the source directory will be uploaded.

The following table describes the AzCopy.exe parameters and their required values. The information you obtained in the previous step is used in the values for these parameters.

PARAMETER	DESCRIPTION	EXAMPLE
`/Source:`	Specifies the source directory in your organization that contains the PST files that will be uploaded to Office 365. Be sure to surround the value of this parameter with double-quotation marks (" ").	`/Source:"\\FILESERVER01\PSTs"`

PARAMETER	DESCRIPTION	EXAMPLE

`/Dest:`	Specifies the SAS URL that you obtained in Step 1.
	Be sure to surround the value of this parameter with double-quotation marks (" ").
	
	Note: If you use the SAS URL in a script or batch file, you need to watch out for certain characters that need to be escaped. For example, you have to change `%` to `%%` and change `&` to `^&`.
	
	Tip: (Optional) You can specify a subfolder in the Azure Storage location to upload the PST files to. You do this by adding a subfolder location (after "ingestiondata") in the SAS URL. The first example doesn't specify a subfolder. That means the PSTs are uploaded to the root (named *ingestiondata*) of the Azure Storage location. The second example uploads the PST files to a subfolder (named *PSTFiles*) in the root of the Azure Storage location.
		`/Dest:"https://3c3e5952a2764023ad14984.blob.core.windows.net/ingestiondata?sv=201`
		`31T23%3A59%3A59Z&sr=c&si=IngestionSasForAzCopy20160112192049811 7&sig=\`
		Or
		`/Dest:"https://3c3e5952a2764023ad14984.blob.core.windows.net/ingestiondata/PSTFil`
		`31T23%3A59%3A59Z&sr=c&si=IngestionSasForAzCopy20160112192049811 7&sig=\`

`/V:`	Outputs verbose status messages into a log file. By default, the verbose log file is named AzCopyVerbose.log in %LocalAppData%\Microsoft\Azure\AzCopy. If you specify an existing file location for this option, the verbose log will be appended to that file. Be sure to surround the value of this parameter with double-quotation marks (" ").	`/V:"c:\Users\Admin\Desktop\Uploadlog.log"`
`/S`	This optional switch specifies the recursive mode so that the AzCopy tool copies PSTs files that are located in subfolders in the source directory that is specified by the `/Source:` parameter. **Note:** If you include this switch, PST files in subfolders will have a different file pathname in the Azure Storage location after they're uploaded. You'll have to specify the exact file pathname in the CSV file that you create in Step 4.	`/S`
`/Y`	This required switch allows the use of write-only SAS tokens when you upload the PST files to the Azure Storage location. The SAS URL you obtained in step 1 (and specified in `/Dest:` parameter) is a write-only SAS URL, which is why you must include this switch. A write-only SAS URL won't prevent you from using the Azure Storage Explorer to view a list of the PST files uploaded to the Azure Storage location.	`/Y`

Here's an example of the syntax for the AzCopy.exe tool using actual values for each parameter:

```
  AzCopy.exe /Source:"\\FILESERVER1\PSTs"
/Dest:"https://3c3e5952a2764023ad14984.blob.core.windows.net/ingestiondata?sv=2012-02-12&amp;se=9999-12-
31T23%3A59%3A59Z&amp;sr=c&amp;si=IngestionSasForAzCopy20160112192049811 7&amp;sig=Vt5S4hVzlzMcBkuH8bH711atBff
drOS72TlV1mNdORg%3D" /V:"c:\Users\Admin\Desktop\AzCopy1.log" /Y
```

After you run the command, status messages are displayed that show the progress of uploading the PST files. A final status message shows the total number of files that were successfully uploaded.

TIP

After you successfully run the AzCopy.exe command and verify that all the parameters are correct, save a copy of the command line syntax to the same (secured) file where you copied the information you obtained in Step 1. Then you can copy and paste this command in a Command Prompt each time that you want to run the AzCopy.exe tool to upload PST files to Office 365. The only value you might have to change are the ones for the `/Source:` parameter. This depends on the source directory where the PST files are located.

(Optional) Step 3: View a list of the PST files uploaded to Office 365

As an optional step, you can install and use the Microsoft Azure Storage Explorer (which is a free, open-source tool) to view the list of the PST files that you've uploaded to the Azure blob. There are two good reasons to do this:

- Verify that PST files from the shared folder or file server in your organization were successfully uploaded to the Azure blob.
- Verify the filename (and the subfolder pathname if you included one) for each PST file uploaded to the

Azure blob. This is helpful when you're creating the PST mapping file in the next step because you have to specify both the folder pathname and filename for each PST file. Verifying these names can help reduce potential errors in your PST mapping file.

The Microsoft Azure Storage Explorer is in Preview.

> **IMPORTANT**
> You can't use the Azure Storage Explorer to upload or modify PST files. The only supported method for importing PST files is to use AzCopy. Also, you can't delete PST files that you've uploaded to the Azure blob. If you try to delete a PST file, you'll receive an error about not having the required permissions. Note that all PST files are automatically deleted from your Azure storage area. If there are no import jobs in progress, then all PST files in the **ingestiondata** container are deleted 30 days after the most recent import job was created.

To install the Azure Storage Explorer and connect to your Azure Storage area:

1. Download and install the Microsoft Azure Storage Explorer tool.
2. Start the Microsoft Azure Storage Explorer, right-click **Storage Accounts** in the left pane, and then click **Connect to Azure Storage**.

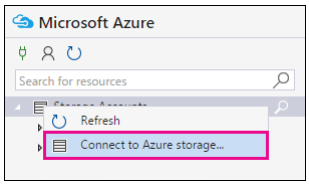

3. Click **Use a shared access signature (SAS) URI or connection string** and click **Next**.
4. Click **Use a SAS URI**, paste the SAS URL that you obtained in Step 1 into the box under **URI**, and then click **Next**.
5. On the **Connection summary** page, you can review the connection information, and then click **Connect**.

   The **ingestiondata** container is opened. It contains the PST files that you uploaded in Step 2. The **ingestiondata** container is located under **Storage Accounts** > **(SAS-Attached Services)** > **Blob Containers**.

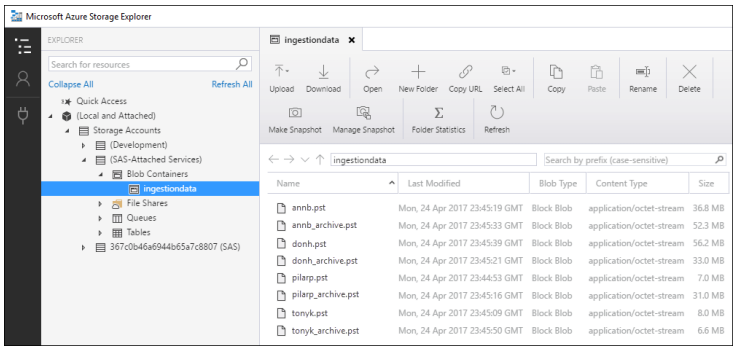

6. When you're finished using the Microsoft Azure Storage Explorer, right-click **ingestiondata**, and then click **Detach** to disconnect from your Azure Storage area. Otherwise, you'll receive an error the next time you try to attach.

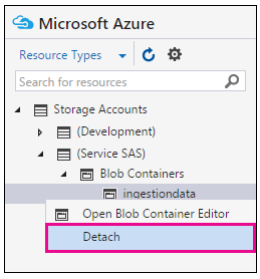

Step 4: Create the PST Import mapping file

After the PST files have been uploaded to the Azure Storage location for your organization, the next step is to create a comma-separated value (CSV) file that specifies which user mailboxes the PST files will be imported to. You'll submit this CSV file in the next step when you create a PST Import job.

1. Download a copy of the PST Import mapping file.
2. Open or save the CSV file to your local computer. The following example shows a completed PST Import mapping file (opened in NotePad). It's much easier to use Microsoft Excel to edit the CSV file.

```
Workload,FilePath,Name,Mailbox,IsArchive,TargetRootFolder,ContentCodePage,SPFileContainer,SPManifestC
ontainer,SPSiteUrl
Exchange,,annb.pst,annb@contoso.onmicrosoft.com,FALSE,/,,,,
Exchange,,annb_archive.pst,annb@contoso.onmicrosoft.com,TRUE,,,,,
Exchange,,donh.pst,donh@contoso.onmicrosoft.com,FALSE,/,,,,
Exchange,,donh_archive.pst,donh@contoso.onmicrosoft.com,TRUE,,,,,
Exchange,PSTFiles,pilarp.pst,pilarp@contoso.onmicrosoft.com,FALSE,/,,,,
Exchange,PSTFiles,pilarp_archive.pst,pilarp@contoso.onmicrosoft.com,TRUE,/ImportedPst,,,,
Exchange,PSTFiles,tonyk.pst,tonyk@contoso.onmicrosoft.com,FALSE,,,,,
Exchange,PSTFiles,tonyk_archive.pst,tonyk@contoso.onmicrosoft.com,TRUE,/ImportedPst,,,,
Exchange,PSTFiles,zrinkam.pst,zrinkam@contoso.onmicrosoft.com,FALSE,,,,,
Exchange,PSTFiles,zrinkam_archive.pst,zrinkam@contoso.onmicrosoft.com,TRUE,/ImportedPst,,,,
```

The first row, or header row, of the CSV file lists the parameters that will be used by the PST Import service to import the PST files to user mailboxes. Each parameter name is separated by a comma. Each row under the header row represents the parameter values for importing a PST file to a specific mailbox. You need a row for each PST file that you want to import to a user mailbox. You can have a maximum of 500 rows in the CSV mapping file. To import more than 500 PST files, you'll have to create multiple mapping files and create multiple import jobs in Step 5.

> **NOTE**
> Don't change anything in the header row, including the SharePoint parameters; they will be ignored during the PST Import process. Also, be sure to replace the placeholder data in the mapping file with your actual data.

3. Use the information in the following table to populate the CSV file with the required information.

PARAMETER	DESCRIPTION	EXAMPLE
`Workload`	Specifies the service that data will be imported to. To import PST files to user mailboxes, use `Exchange`.	`Exchange`
`FilePath`	Specifies the folder location in the Azure Storage location that you uploaded the PST files to in Step 2. If you didn't include an optional subfolder name in the SAS URL in the `/Dest:` parameter in Step 2, leave this parameter blank in the CSV file. If you included a subfolder name, specify it in this parameter (see the second example). The value for this parameter is case-sensitive. Either way, *don't* include "ingestiondata" in the value for the `FilePath` parameter. **Important:** The case for the file path name must be the same as the case you used if you included an optional subfolder name in the SAS URL in the `/Dest:` parameter in Step 2. For example, if you used `PSTFiles` for the subfolder name in Step 2 and then use `pstfiles` in the `FilePath` parameter in CSV file, the import for the PST file will fail. Be sure to use the same case in both instances.	(leave blank) Or `PSTFiles`
`Name`	Specifies the name of the PST file that will be imported to the user mailbox. The value for this parameter is case-sensitive. **Important:** The case for the PST file name in the CSV file must be the same as the PST file that was uploaded to the Azure Storage location in Step 2. For example, if you use `annb.pst` in the `Name` parameter in the CSV file, but the name of the actual PST file is `AnnB.pst`, the import for that PST file will fail. Be sure that the name of the PST in the CSV file uses the same case as the actual PST file.	`annb.pst`

PARAMETER	DESCRIPTION	EXAMPLE

`Mailbox`	Specifies the email address of the mailbox that the PST file will be imported to. You can't specify a public folder because the PST Import Service doesn't support importing PST files to public folders. To import a PST file to an inactive mailbox, you have to specify the mailbox GUID for this parameter. To obtain this GUID, run the following PowerShell command in Exchange Online:
	`Get-Mailbox <identity of inactive mailbox> -InactiveMailboxOnly | FL Guid`
	
	Note: Sometimes you might have multiple mailboxes with the same email address, where one mailbox is an active mailbox and the other mailbox is in a soft-deleted (or inactive) state. In these situations, you have to specify the mailbox GUID to uniquely identify the mailbox to import the PST file to. To obtain this GUID for active mailboxes, run the following PowerShell command:
	`Get-Mailbox <identity of active mailbox> | FL Guid`
	. To obtain the GUID for soft-deleted (or inactive) mailboxes, run this command
	`Get-Mailbox <identity of soft-deleted or inactive mailbox> -SoftDeletedMailbox | FL Guid`
	.
		`annb@contoso.onmicrosoft.com`
		Or
		`2d7a87fe-d6a2-40cc-8aff-1ebea80d4ae7`

`IsArchive`	Specifies whether to import the PST file to the user's archive mailbox. There are two options:
	
	FALSE: Imports the PST file to the user's primary mailbox.
	TRUE: Imports the PST file to the user's archive mailbox. This assumes that the user's archive mailbox is enabled.
	
	If you set this parameter to `TRUE` and the user's archive mailbox isn't enabled, the import for that user will fail. If an import fails for one user (because their archive isn't enabled and this property is set to `TRUE`), the other users in the import job won't be affected.
	If you leave this parameter blank, the PST file is imported to the user's primary mailbox.
	
	Note: To import a PST file to a cloud-based archive mailbox for a user whose primary mailbox is on-premises, just specify `TRUE` for this parameter and specify the email address for the user's on-premises mailbox for the `Mailbox` parameter.
		`FALSE`
		Or
		`TRUE`

PARAMETER	DESCRIPTION	EXAMPLE

`TargetRootFolder`	Specifies the mailbox folder that the PST file is imported to.
	
	If you leave this parameter blank, the PST file will be imported to a new folder named **Imported** at the root level of the mailbox (the same level as the Inbox folder and the other default mailbox folders).
	
	If you specify `/`, the folders and items in the PST file are imported to the top of the folder structure in the target mailbox or archive. If a folder exists in the target mailbox (for example, default folders such as Inbox, Sent Items, and Deleted Items), the items in that folder in the PST are merged into the existing folder in the target mailbox. For example, if the PST file contains an Inbox folder, items in that folder are imported to the Inbox folder in the target mailbox. New folders are created if they don't exist in the folder structure for the target mailbox.
	
	If you specify `/<foldername>`, items and folders in the PST file are imported to a folder named *<foldername>*. For example, if you use `/ImportedPst`, items would be imported to a folder named **ImportedPst**. This folder will be located in the user's mailbox at the same level as the Inbox folder.
	
	Tip: Consider running a few test batches to experiment with this parameter so you can determine the best folder location to import PSTs files to.
		(leave blank)
		Or
		`/`
		Or
		`/ImportedPst`

`ContentCodePage`	This optional parameter specifies a numeric value for the code page to use for importing PST files in the ANSI file format. This parameter is used for importing PST files from Chinese, Japanese, and Korean (CJK) organizations because these languages typically use a double byte character set (DBCS) for character encoding. If this parameter isn't used to import PST files for languages that use DBCS for mailbox folder names, the folder names are often garbled after they're imported. For a list of supported values to use for this parameter, see Code Page Identifiers. **Note:** As previously stated, this is an optional parameter and you don't have to include it in the CSV file. Or you can include it and leave the value blank for one or more rows.	(leave blank) Or `932` (which is the code page identifier for ANSI/OEM Japanese)
`SPFileContainer`	For PST Import, leave this parameter blank.	Not applicable
`SPManifestContainer`	For PST Import, leave this parameter blank.	Not applicable
`SPSiteUrl`	For PST Import, leave this parameter blank.	Not applicable

Step 5: Create a PST Import job

The next step is to create the PST Import job in the Import service in Microsoft 365. As previously explained, you submit the PST Import mapping file that you created in Step 4. After you create the job, Microsoft 365 analyzes the data in the PST files and then gives you an opportunity to filter the data that actually gets imported to the mailboxes specified in the PST import mapping file (see Step 6).

1. Go to https://protection.office.com and sign in using the credentials for an administrator account in your organization.
2. In the left pane of the Security & Compliance Center, click **Information governance > Import > Import PST files**.
3. On the **Import PST files** page, click **+New import job**.

> **NOTE**
> You have to be assigned the appropriate permissions to access the **Import** page in the Security & Compliance Center to create an import job. See the **Before you begin** section for more information.

4. Type a name for the PST import job, and then click **Next**. Use lowercase letters, numbers, hyphens, and underscores. You can't use uppercase letters or include spaces in the name.

5. On the **Do you want to upload or ship data?** page, click **Upload your data** and then click **Next**.

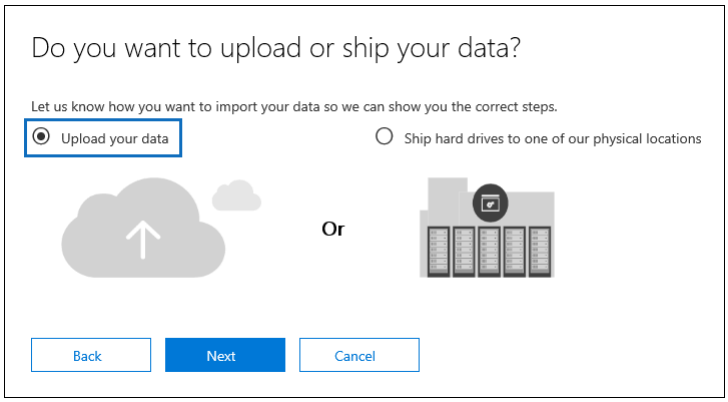

6. In step 4 on the **Import data** page, click the **I'm done uploading my files** and **I have access to the mapping file** check boxes, and then click **Next**.

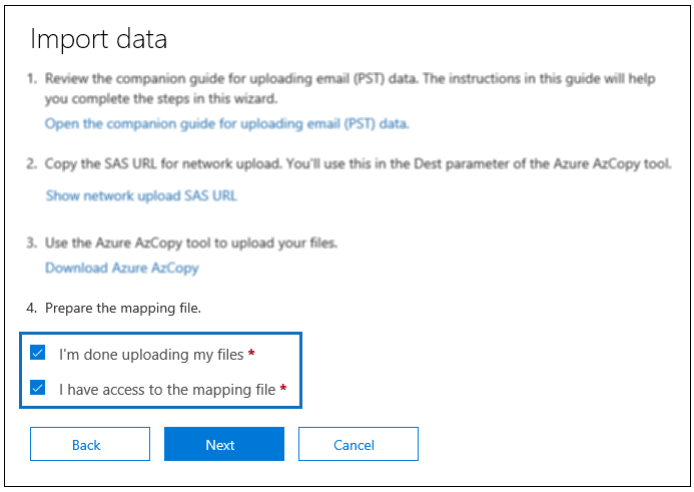

7. On the **Select the mapping file** page, click **Select mapping file** to submit the CSV mapping file that you created in Step 4.

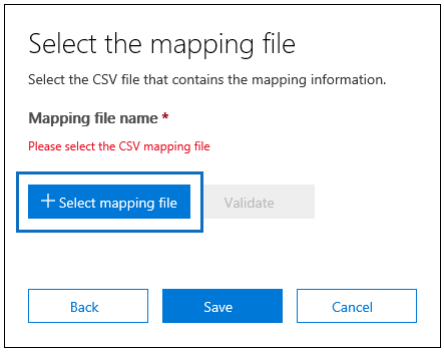

8. After the name of the CSV file appears under **Mapping file name**, click **Validate** to check your CSV file for errors.

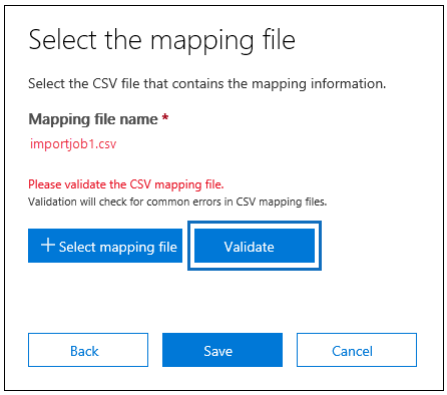

The CSV file has to be successfully validated to create a PST Import job. The file name is changed to green after it's successfully validated. If the validation fails, click the **View log** link. A validation error report is opened, with an error message for each row in the file that failed.

> **NOTE**
> As previously explained, a mapping file can have a maximum of 500 rows. Validation will fail if the mapping file contains more than 500 rows. To import more than 500 PST files, you'll have to create multiple mapping files and multiple import jobs.

9. After the mapping file is successfully validated, read the terms and conditions document, and then click the checkbox.
10. Click **Save** to submit the job, and then click **Close** after the job is successfully created.

    A status flyout page is displayed, with a status of **Analysis in progress** and the new import job is displayed in the list on the **Import PST files** page.
11. Click **Refresh** to update the status information that's displayed in the **Status** column. When the analysis is complete and the data is ready to be imported, the status is changed to **Analysis completed**.

    You can click the import job to display the status flyout page, which contains more detailed information about the import job such as the status of each PST file listed in the mapping file.

Step 6: Filter data and start the PST Import job

After you create the import job in Step 5, Microsoft 365 analyzes the data in the PST files (in a safe and secure manner) by identifying the age of the items and the different message types included in the PST files. When the analysis is completed and the data is ready to import, you have the option to import all the data contained in the PST files or you can trim the data that's imported by setting filters that control what data gets imported.

1. On the **Import PST files** page in the Security & Compliance Center, click **Ready to import to Office 365** for the import job that you created in Step 5.

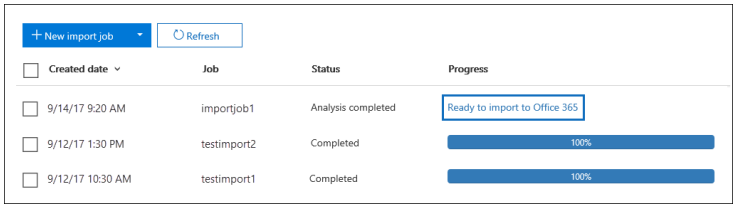

   A fly out page is displayed with information about the PST files and other information about the import job.
2. On the flyout page, click **Import to Office 365**.

   The **Filter your data** page is displayed. It contains the data insights resulting from the analysis performed on the PST files by Office 365, including information about the age of the data. At this point, you have the option to filter the data that will be imported or import all the data as is.

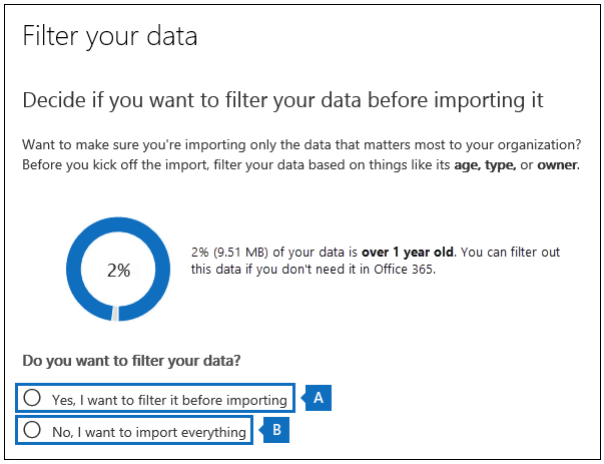

3. Do one of the following:

a. To trim the data that you import, click **Yes, I want to filter it before importing**.

   For detailed step-by-step instructions about filtering the data in the PST files and then starting the import job, see Filter data when importing PST files to Office 365.

   Or

b. To import all data in the PST files, click **No, I want to import everything**, and click **Next**.

4. If you chose to import all the data, click **Import data** to start the import job.

   The status of the import job is display on the **Import PST files** page. Click **Refresh** to update the status information that's displayed in the **Status** column. Click the import job to display the status flyout page, which displays status information about each PST file being imported.

More information

- Why import PST files to Microsoft 365?
  - It's a good way to import your organization's archival messaging data to Microsoft 365.
  - The data is available to the user from all devices because it's stored in the cloud.
  - It helps address compliance needs of your organization by letting you apply Microsoft 365 compliance features to the data from the PST files that you imported. This includes:
  - Enabling archive mailboxes and auto-expanding archiving to give users additional mailbox storage space to store the data that you imported.
  - Placing mailboxes on Litigation Hold to retain the data that you imported.
  - Using Microsoft eDiscovery tools to search the data that you imported.
  - Using Microsoft 365 retention policies to control how long the data that you imported will be retained, and what action to take after the retention period expires.
  - Searching the audit log for mailbox-related events that affect the data that you imported.
  - Importing data to inactive mailboxes to archive data for compliance purposes.
  - Using data loss prevention policies to prevent sensitive data from leaking outside your organization.
- Here's an example of the Shared Access Signature (SAS) URL that's obtained in Step 1. This example also contains the syntax for the command that you run in the AzCopy.exe tool to upload PST files. Be sure to take precautions to protect the SAS URL just like you would protect passwords or other security-related information.

```
SAS URL: https://3c3e5952a2764023ad14984.blob.core.windows.net/ingestiondata?sv=2012-02-
12&amp;se=9999-12-
31T23%3A59%3A59Z&amp;sr=c&amp;si=IngestionSasForAzCopy20160112192049811 7&amp;sig=Vt5S4hVzlzMcBkuH8bH7
11atBffdrOS72TlV1mNdORg%3D

AzCopy.exe /Source:<Location of PST files> /Dest:<SAS URL> /V:<Log file location> /Y

EXAMPLES

This example uploads PST files to the root of the Azure storage location:

AzCopy.exe /Source:"\\FILESERVER1\PSTs"
/Dest:"https://3c3e5952a2764023ad14984.blob.core.windows.net/ingestiondata?sv=2012-02-12&amp;se=9999-
12-
31T23%3A59%3A59Z&amp;sr=c&amp;si=IngestionSasForAzCopy20160112192049811 7&amp;sig=Vt5S4hVzlzMcBkuH8bH7
11atBffdrOS72TlV1mNdORg%3D" /V:"c:\Users\Admin\Desktop\AzCopy1.log" /Y

This example uploads PST files to a subfolder named PSTFiles  in the Azure storage location:

AzCopy.exe /Source:"\\FILESERVER1\PSTs"
/Dest:"https://3c3e5952a2764023ad14984.blob.core.windows.net/ingestiondata/PSTFiles?sv=2012-02-
12&amp;se=9999-12-
31T23%3A59%3A59Z&amp;sr=c&amp;si=IngestionSasForAzCopy20160112192049811 7&amp;sig=Vt5S4hVzlzMcBkuH8bH7
11atBffdrOS72TlV1mNdORg%3D" /V:"c:\Users\Admin\Desktop\AzCopy1.log" /Y
```

- As previously explained, the Office 365 Import service turns on the retention hold setting (for an indefinite duration) after PST files are imported to a mailbox. This means the *RetentionHoldEnabled* property is set to **True** so that the retention policy assigned to the mailbox won't be processed. This gives the mailbox owner time to manage the newly imported messages by preventing a deletion or archive policy from deleting or archiving older messages. Here are some steps you can take to manage this retention hold:
  - After a certain time, you can turn off the retention hold by running the **Set-Mailbox -RetentionHoldEnabled $false** command. For instructions, see Place a mailbox on retention hold.
  - You can configure the retention hold so that it's turned off on some date in the future. You do this by running the **Set-Mailbox -EndDateForRetentionHold *date*** command. For example, assuming that today's date is June 1, 2016 and you want the retention hold turned off in 30 days, you would run the following command: **Set-Mailbox -EndDateForRetentionHold 7/1/2016**. In this scenario, you would leave the **RetentionHoldEnabled** property set to *True*. For more information, see Set-Mailbox.
  - You can change the settings for the retention policy that's assigned to the mailbox so that older items that were imported won't be immediately deleted or moved to the user's archive mailbox. For example, you could lengthen the retention age for a deletion or archive policy that's assigned to the mailbox. In this scenario, you would turn off the retention hold on the mailbox after you changed the settings of the retention policy. For more information, see Set up an archive and

deletion policy for mailboxes in your organization.

How the import process works

You can use the network upload option and the Office 365 Import service to bulk-import PST files to user mailboxes. Network upload means that you upload the PST files a temporary storage area in the Microsoft cloud. Then the Office 365 Import service copies the PST files from the storage area to the target user mailboxes.

Here's an illustration and description of the network upload process to import PST files to mailboxes in Office 365.

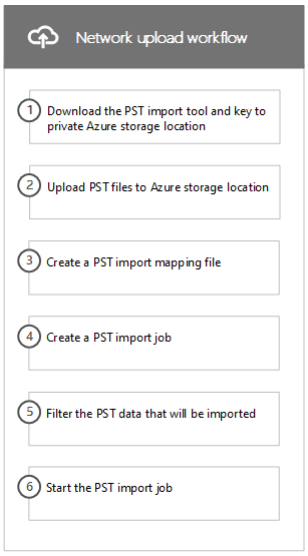

1. **Download the PST import tool and key to private Azure Storage location:** The first step is to download the AzCopy command-line tool and an access key used to upload the PST files to an Azure Storage location in the Microsoft cloud. You obtain these from the **Import** page in the Security & Compliance Center. The key (called a secure access signature (SAS) key, provides you with the necessary permissions to upload PST files to a private and secure Azure Storage location. This access key is unique to your organization and helps prevent unauthorized access to your PST files after they're uploaded to the Microsoft cloud. Importing PST files doesn't require your organization to have a separate Azure subscription.

2. **Upload the PST files to the Azure Storage location:** The next step is to use the AzCopy.exe tool (downloaded in step 1) to upload and store your PST files in an Azure Storage location that resides in the same regional Microsoft datacenter where your organization is located. To upload them, the PST files that you want to import have to be located in a file share or file server in your organization.

   There's an optional step that you can perform to view the list of PST files after they're uploaded to the Azure Storage location.

3. **Create a PST import mapping file:** After the PST files have been uploaded to the Azure Storage location, the next step is to create a comma-separated value (CSV) file that specifies which user mailboxes the PST files will be imported to, note that a PST file can be imported to a user's primary mailbox or their archive mailbox. The Office 365 Import service uses the information in the CSV file to import the PST files.

4. **Create a PST import job:** The next step is to create a PST import job on the **Import PST files** page in the Security & Compliance Center and submit the PST import mapping file created in the previous step. After you create the import job, Microsoft 365 analyzes the data in the PST files and then gives you an opportunity to set filters that control what data actually gets imported to the mailboxes specified in the PST import mapping file.

5. **Filter the PST data that will be imported to mailboxes:** After the import job is created and started, Microsoft 365 analyzes the data in the PST files (safely and securely) by identifying the age of the items and the different message types included in the PST files. When the analysis is completed and the data is ready to import, you have the option to import all the data contained in the PST files or you can trim the data that's imported by setting filters that control what data gets imported.

6. **Start the PST import job:** After the import job is started, Microsoft 365 uses the information in the PST import mapping file to import the PSTs files from the Azure Storage location to user mailboxes. Status information about the import job (including information about each PST file being imported) is displayed on the **Import PST files** page in the Security & Compliance Center. When the import job is finished, the status for the job is set to **Complete**.

Use drive shipping to import your organization's PST files

11/2/2020 • 34 minutes to read • Edit Online

This article is for administrators. Are you trying to import PST files to your own mailbox? See Import email, contacts, and calendar from an Outlook .pst file

Use the Office 365 Import service and drive shipping to bulk-import PST files to user mailboxes. Drive shipping means that you copy the PST files to a hard disk drive and then physically ship the drive to Microsoft. When Microsoft receives your hard drive, data center personnel copies the data from the hard drive to a storage area in the Microsoft cloud. Then you have the opportunity to trim the PST data that's imported to the target mailboxes by setting filters that control what data gets imported. After you start the import job, the Import service imports the PST data from the storage area to user mailboxes. Using drive shipping to import PST files to user mailboxes is one way to migrate your organization's email to Office 365.

Here are the steps required to use drive shipping to import PST files to Microsoft 365 mailboxes:

Step 1: Download the secure storage key and PST Import tool

Step 2: Copy the PST files to the hard drive

Step 3: Create the PST Import mapping file

Step 4: Create a PST Import job in Office 365

Step 5: Ship the hard drive to Microsoft

Step 6: Filter data and start the PST Import job

> **IMPORTANT**
> You have to perform Step 1 once to down load the secure storage key and the import tool. After you perform these steps, follow Step 2 through Step 6 each time you want to ship a hard drive to Microsoft.

For frequently asked questions about using drive shipping to import PST files to Office 365, see FAQs for using drive shipping to import PST files.

Before you import PST files

- You have to be assigned the Mailbox Import Export role in Exchange Online to import PST files to Microsoft 365 mailboxes. By default, this role isn't assigned to any role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group. Or you can create a role group, assign the Mailbox Import Export role, and then add yourself as a member. For more information, see the "Add a role to a role group" or the "Create a role group" sections in Manage role groups.

  Additionally, to create import jobs in the Security & Compliance Center, one of the following must be true:

  - You have to be assigned the Mail Recipients role in Exchange Online. By default, this role is assigned to the Organization Management and Recipient Management roles groups.

    Or

  - You have to be a global administrator in your organization.

    > **TIP**
    > Consider creating a new role group in Exchange Online that's specifically intended for importing PST files to Office 365. For the minimum level of privileges required to import PST files, assign the Mailbox Import Export and Mail Recipients roles to the new role group, and then add members.

- You need to store the PST files that you want to copy to a hard drive on a file server or shared folder in your organization. In Step 2, you run the Azure Import Export tool (WAImportExport.exe) that copies the PST files that are stored on this file server or shared folder to the hard drive.
- Large PST files may impact the performance of the PST import process. So we recommend that each PST file you copy to the hard drive in Step 2 should be no larger than 20 GB.
- Only 2.5-inch solid-state drives (SSDs) or 2.5-inch or 3.5-inch SATA II/III internal hard drives are supported for use with the Office 365 Import service. You can use hard drives up to 10 TB. For import jobs, only the first data volume on the hard drive will be processed. The data volume must be formatted with NTFS. When copying data to a hard drive, you can attach it directly using a 2.5-inch SSD or 2.5-inch or 3.5-inch SATA II/III connector or you can attach it externally using an external 2.5-inch SSD or 2.5-inch

or 3.5-inch SATA II/III USB adaptor.

> **IMPORTANT**
> External hard drives that come with an built-in USB adaptor aren't supported by the Office 365 Import service. Additionally, the disk inside the casing of an external hard drive can't be used. Please don't ship external hard drives.

- The hard drive that you copy the PST files to must be encrypted with BitLocker. The WAImportExport.exe tool that you run in Step 2 will help you set up BitLocker. It also generates a BitLocker encryption key that Microsoft data center personnel use to access the drive to upload the PST files to the Azure Storage area in the Microsoft cloud.
- Drive shipping is available through a Microsoft Enterprise Agreement (EA). Drive shipping isn't available through a Microsoft Products and Services Agreement (MPSA).
- The cost to import PST files to Microsoft 365 mailboxes using drive shipping is $2 USD per GB of data. For example, if you ship a hard drive that contains 1,000 GB (1 TB) of PST files, the cost is $2,000 USD. You can work with a partner to pay the import fee. For information about finding a partner, see Find your Microsoft partner or reseller.
- You or your organization must have an account with FedEx or DHL.
  - Organizations in the United States, Brazil, and Europe must have FedEx accounts.
  - Organizations in East Asia, Southeast Asia, Japan, Republic of Korea, and Australia must have DHL accounts.

    Microsoft uses (and charges) this account to return the hard drive back to you.
- The hard drive that you ship to Microsoft may cross international borders. In this case, you're responsible for ensuring that the hard drive and the data it contains are imported and/or exported in accordance with the applicable laws. Before shipping a hard drive, check with your advisors to verify that your drive and data can legally be shipped to the identified Microsoft data center. This helps to ensure that it reaches Microsoft in a timely manner.
- This procedure involves copying and saving a secure storage key and a BitLocker encryption key. Be sure to take precautions to protect these keys like you would protect passwords or other security-related information. For example, you might save them to a password-protected Microsoft Word document or save them to an encrypted USB drive. See the More information section for an example of these keys.
- After PST files are imported to a Microsoft 365 mailbox, the retention hold setting for the mailbox is turned on for an indefinite duration. This means that the retention policy assigned to the mailbox won't be processed until you turn off the retention hold or set a date to turn off the hold. Why do we do this? If messages imported to a mailbox are old, they might be permanently deleted (purged) because their retention period has expired based on the retention settings configured for the mailbox. Placing the mailbox on retention hold gives the mailbox owner time to manage these newly imported messages or give you time to change the retention settings for the mailbox. See the More information section for suggestions about managing the retention hold.
- By default, the maximum message size that can be received by a Microsoft 365 mailbox is 35 MB. That's because the default value for the *MaxReceiveSize* property for a mailbox is set to 35 MB. However, the limit for the maximum message receive size in Microsoft 365 is 150 MB. So if you import a PST file that contains an item larger than 35 MB, the Office 365 Import service we will automatically change the value of the *MaxReceiveSize* property on the target mailbox to 150 MB. This allows messages up to 150 MB to be imported to user mailboxes.

  > **TIP**
  > To identify the message receive size for a mailbox, you can run this command in Exchange Online PowerShell: `Get-Mailbox <user mailbox> | FL MaxReceiveSize`.
- You can import PST files to an inactive mailbox in Office 365. You do this by specifying the GUID of the inactive mailbox in the `Mailbox` parameter in the PST Import mapping file. See Step 3: Create the PST Import mapping file for more information.
- In an Exchange hybrid deployment, you can import PST files to a cloud-based archive mailbox for a user whose primary mailbox is on-premises. You do this by doing the following in the PST Import mapping file:
  - Specify the email address for the user's on-premises mailbox in the `Mailbox` parameter.
  - Specify the **TRUE** value in the `IsArchive` parameter.

    See Step 3: Create the PST Import mapping file for more information.

Step 1: Download the secure storage key and PST Import tool

The first step is to download the secure storage key and the tool and that you use in Step 2 to copy PST files to the hard drive.

> **IMPORTANT**
> You have to use Azure Import/Export tool version 1 (WAimportExportV1) to successfully import PST files by using the drive shipping method. Version 2 of the Azure Import/Export tool isn't supported and using it will result in incorrectly preparing the hard drive for the import job. Be sure to download the Azure Import/Export tool from the Security & Compliance Center by following the procedures in this step.

1. Go to https://protection.office.com/ and sign in using the credentials for an administrator account in your organization.
2. In the left pane of the Security & Compliance Center, click **Information governance** > **Import** > **Import PST files**.

   > **NOTE**
   > As previously stated, you have to be assigned the appropriate permissions to access the **Import** page in the Security & Compliance Center.

3. On the **Import PST files** page, click **+ New import job**.
4. In the import job wizard, type a name for the PST import job, and then click **Next**. Use lowercase letters, numbers, hyphens, and underscores. You can't use uppercase letters or include spaces in the name.
5. On the **Choose import job type** page, click **Ship hard drives to one of our physical locations** and then click **Next**.

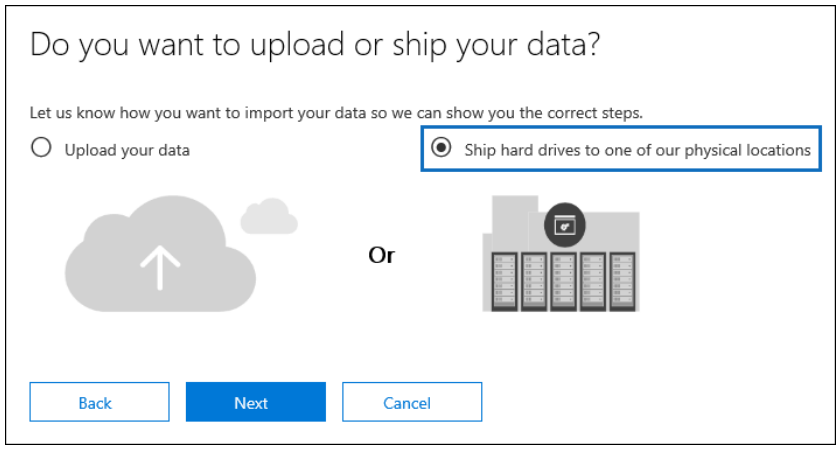

6. On the **Import data** page, do the following two things:

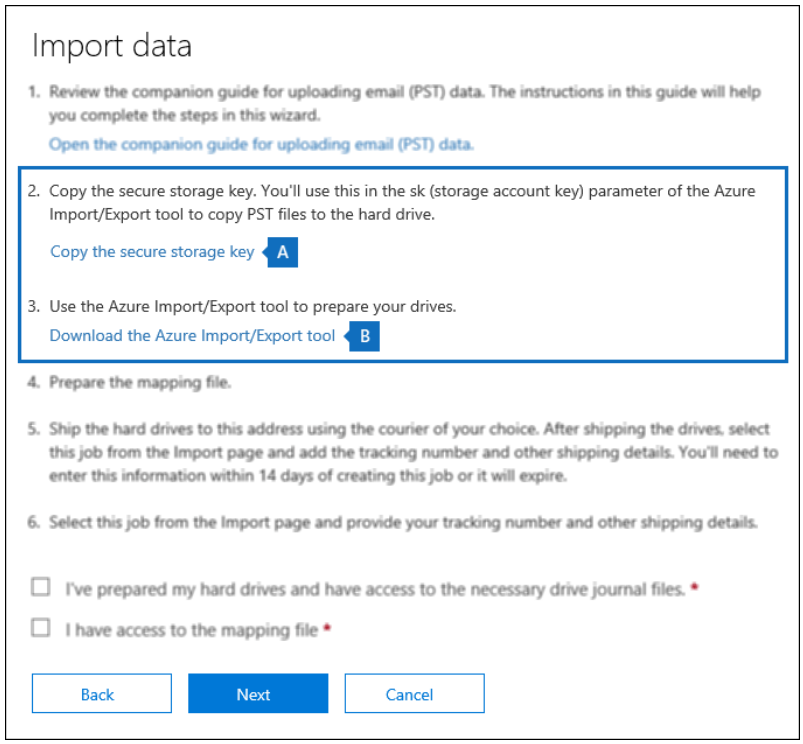

   a. In step 2, click **Copy the secure storage key**. After the storage key is displayed, click **Copy to**

clipboard and then paste it and save it to a file so you can access it later.

b. In step 3, **Download the Azure Import/Export tool** to download and install the Azure Import/Export (version 1) tool.

- In the pop-up window, click **Save** > **Save as** to save the WaImportExportV1.zip file to a folder on your local computer.
- Extract the WaImportExportV1.zip file.

7. Click **Cancel** to close the wizard.

   You come back to the **Import** page in the Security & Compliance Center when you create the import job in Step 4.

Step 2: Copy the PST files to the hard drive

The next step is to use the WAImportExport.exe tool to copy PST files to the hard drive. This tool encrypts the hard drive with BitLocker, copies the PSTs to the hard drive, and creates a journal file that stores information about the copy process. To complete this step, the PST files have to be located in a file share or file server in your organization. This is known as the source directory in the following procedure.

As previously stated, each PST file that you copy to the hard drive should be no larger than 20 GB. PST files larger than 20 GB may impact the performance of the PST import process that you start in Step 6.

> **IMPORTANT**
> After you run the WAImportExport.exe tool the first time for a hard drive, you have to use a different syntax each time after that. This syntax is explained in step 4 of this procedure to copy PST files to the hard drive.

1. Open a Command Prompt on your local computer.

   > **TIP**
   > If you run the command prompt as an administrator (by selecting "Run as administrator" when you open it) error messages will be displayed in the command prompt window. This can help you troubleshoot problems running the WAImportExport.exe tool.

2. Go to the directory where you installed the WAImportExport.exe tool in Step 1.

3. Run the following command the first time that you use the WAImportExport.exe to copy PST files to a hard drive.

   ```
   WAImportExport.exe PrepImport /j:<Name of journal file> /t:<Drive letter> /id:<Name of session>
   /srcdir:<Location of PST files> /dstdir:<PST file path> /sk:<Storage account key> /blobtype:BlockBlob
   /encrypt /logdir:<Log file location>
   ```

   The following table describes the parameters and their required values.

PARAMETER	DESCRIPTION	EXAMPLE
`/j:`	Specifies the name of the journal file. This file is saved to the same folder where the WAImportExport.exe tool is located. Each hard drive you ship to Microsoft must have one journal file. Every time you run the WAImportTool.exe to copy PST files to a hard drive, information will be appended to the journal file for that drive. Microsoft data center personnel use the information in the journal file to associate the hard drive with the import job that you create in Step 4, and to upload the PST files to the Azure Storage area in the Microsoft cloud.	`/j:PSTHDD1.jrn`
`/t:`	Specifies the drive letter of the hard drive when it's connected to your local computer.	`/t:h`

PARAMETER	DESCRIPTION	EXAMPLE
`/id:`	Specifies the name of the copy session. A session is defined as each time you run the WAImportExport.exe tool to copy files to the hard drive. The PST files are copied to a folder named with the session name specified by this parameter.	`/id:driveship1`
`/srcdir:`	Specifies the source directory in your organization that contains the PST files that will be copied during the session. Be sure to surround the value of this parameter with double-quotation marks (" ").	`/srcdir:"\\FILESERVER01\PSTs"`

`/dstdir:`	Specifies the destination directory in the Azure Storage area in the Microsoft cloud where the PSTs will be uploaded. You must use the value `ingestiondata/`. Be sure to surround the value of this parameter with double-quotation marks (" ").
	Optionally, you can also add an extra file path to the value of this parameter. For example, you can use the file path of the source directory on the hard drive (converted to a URL format), which is specified in the `/srcdir:` parameter. For example, `\\FILESERVER01\PSTs` is changed to `FILESERVER01/PSTs`. In this case, you still must include `ingestiondata` in the file path. So in this example, the value for the `/dstdir:` parameter would be `"ingestiondata/FILESERVER01/PSTs"`.
	One reason to add the additional file path is if you have PSTs files with the same filename.
	> [!NOTE]> If you include the optional pathname, the namespace for a PST file after it's uploaded to the Azure Storage area includes the pathname and the name of the PST file; for example, `FILESERVER01/PSTs/annb.pst`. If you don't include a pathname, the namespace is only the PST filename; for example `annb.pst`.
		`/dstdir:"ingestiondata/"`
		Or
		`/dstdir:"ingestiondata/FILESERVER01/PSTs"`

`/sk:`	Specifies the storage account key that you obtained in Step 1. Be sure to surround the value of this parameter with double-quotation marks (" ").	`"yaNIIs9Uy5g25Yoak+LlSHfqVBGOeNwjqtBEBGqRMoidq6/e5k/VPkjOXdDIXJHxHv`
`/blobtype:`	Specifies the type of blobs in the Azure Storage area to import the PST files to. For importing PST files, use the value **BlockBlob**. This parameter is required.	`/blobtype:BlockBlob`
`/encrypt`	This switch turns on BitLocker for the hard drive. This parameter is required the first time you run the WAImportExport.exe tool. The BitLocker encryption key is copied to the journal file and the log file that is created if you use the `/logfile:` parameter. As previously explained, the journal file is saved to the same folder where the WAImportExport.exe tool is located.	`/encrypt`

PARAMETER	DESCRIPTION	EXAMPLE
`/logdir:`	This optional parameter specifies a folder to save log files to. If not specified, the log files are saved to the same folder where the WAImportExport.exe tool is located. Be sure to surround the value of this parameter with double-quotation marks (" ").	`/logdir:"c:\users\admin\desktop\PstImportLogs"`

Here's an example of the syntax for the WAImportExport.exe tool using actual values for each parameter:

```
WAImportExport.exe PrepImport /j:PSTHDD1.jrn /t:f /id:driveship1 /srcdir:"\\FILESERVER01\PSTs"
/dstdir:"ingestiondata/"
/sk:"yaNIIs9Uy5g25Yoak+LlSHfqVBGOeNwjqtBEBGqRMoidq6/e5k/VPkjOXdDIXJHxHvNoNoFH5NcVUJXHwu9ZxQ=="
blobtype:BlockBlob /encrypt /logdir:"c:\users\admin\desktop\PstImportLogs"
```

After you run the command, status messages are displayed that show the progress of copying the PST files to the hard drive. A final status message shows the total number of files that were successfully copied.

4. Run this command each subsequent time you run the WAImportExport.ext tool to copy PST files to the same hard drive.

```
WAImportExport.exe PrepImport /j:<Name of journal file> /id:<Name of new session> /srcdir:<Location
of PST files> /dstdir:<PST file path> /blobtype:BlockBlob
```

Here's an example of the syntax for running subsequent sessions to copy PST files to the same hard drive.

```
WAImportExport.exe PrepImport /j:PSTHDD1.jrn /id:driveship2 /srcdir:"\\FILESERVER01\PSTs\SecondBatch"
/dstdir:"ingestiondata/" /blobtype:BlockBlob
```

Step 3: Create the PST Import mapping file

After Microsoft data center personnel upload the PST files from the hard drive to the Azure Storage area, the Import service will use the information in the PST Import mapping file, which is a comma-separated value (CSV) file, that specifies which user mailboxes the PST files are imported to. You will submit this CSV file in the next step when you create a PST Import job.

1. Download a copy of the PST Import mapping file.

2. Open or save the CSV file to your local computer. The following example shows a completed PST Import mapping file (opened in NotePad). It's much easier to use Microsoft Excel to edit the CSV file.

```
Workload,FilePath,Name,Mailbox,IsArchive,TargetRootFolder,ContentCodePage,SPFileContainer,SPManifestC
ontainer,SPSiteUrl
Exchange,FILESERVER01/PSTs,annb.pst,annb@contoso.onmicrosoft.com,FALSE,/,,,,
Exchange,FILESERVER01/PSTs,annb_archive.pst,annb@contoso.onmicrosoft.com,TRUE,/ImportedPst,,,,
Exchange,FILESERVER01/PSTs,donh.pst,donh@contoso.onmicrosoft.com,FALSE,/,,,,
Exchange,FILESERVER01/PSTs,donh_archive.pst,donh@contoso.onmicrosoft.com,TRUE,/ImportedPst,,,,
Exchange,FILESERVER01/PSTs,pilarp.pst,pilarp@contoso.onmicrosoft.com,FALSE,/,,,,
Exchange,FILESERVER01/PSTs,pilarp_archive.pst,pilarp@contoso.onmicrosoft.com,TRUE,/ImportedPst,,,,
Exchange,,tonyk.pst,tonyk@contoso.onmicrosoft.com,FALSE,/,,,,
Exchange,,tonyk_archive.pst,tonyk@contoso.onmicrosoft.com,TRUE,,,,,
Exchange,,zrinkam.pst,zrinkam@contoso.onmicrosoft.com,FALSE,/,,,,
Exchange,,zrinkam_archive.pst,zrinkam@contoso.onmicrosoft.com,TRUE,,,,,
```

The first row, or header row, of the CSV file lists the parameters that will be used by the PST Import service to import the PST files to user mailboxes. Each parameter name is separated by a comma. Each row under the header row represents the parameter values for importing a PST file to a specific mailbox. You need a row for each PST file that was copied to the hard drive. Be sure to replace the placeholder data in the mapping file with your actual data.

> **NOTE**
> Don't change anything in the header row, including the SharePoint parameters; they will be ignored during the PST Import process.

3. Use the information in the following table to populate the CSV file with the required information.

PARAMETER	DESCRIPTION	EXAMPLE
`Workload`	Specifies the service that data will be imported to. To import PST files to user mailboxes, use `Exchange`.	`Exchange`

PARAMETER	DESCRIPTION	EXAMPLE
`FilePath`	Specifies the folder location in the Azure Storage area that PST files will be copied to when the hard drive is shipped to Microsoft. What you add in this column in the CSV file depends on what you specified in for the `/dstdir:` parameter in the previous step. If you have subfolders on the source location, then the value in the `FilePath` parameter must contain the relative path for the subfolder; for example, /folder1/user1/. If you used `/dstdir:"ingestiondata/"`, then leave this parameter blank in the CSV file. If you included an optional pathname for the value of the `/dstdir:` parameter (for example, `/dstdir:"ingestiondata/FILESERVER01/PSTs"`, then use that pathname (not including "ingestiondata") for this parameter in the CSV file. The value for this parameter is case-sensitive. Either way, *don't* include "ingestiondata" in the value for the `FilePath` parameter. Leave this parameter blank or specify only the optional pathname. > [!IMPORTANT]> The case for the file path name must be the same case that you specified in the `/dstdir:` parameter in the previous step. For example, if you used `"ingestiondata/FILESERVER01/PSTs"` for the subfolder name in the previous step, but then used `fileserver01/psts` in the `FilePath` parameter in CSV file, the import for the PST file will fail. Be sure to use the same case in both instances.	(leave blank) Or `FILESERVER01/PSTs`
`Name`	Specifies the name of the PST file that will be imported to the user mailbox. The value for this parameter is case-sensitive. > [!IMPORTANT]> The case for the PST file name in the CSV file must be the same as the PST file that was uploaded to the Azure Storage location in Step 2. For example, if you use `annb.pst` in the `Name` parameter in the CSV file, but the name of the actual PST file is `AnnB.pst`, the import for that PST file will fail. Be sure that the name of the PST in the CSV file uses the same case as the actual PST file.	`annb.pst`

PARAMETER	DESCRIPTION	EXAMPLE
`Mailbox`	Specifies the email address of the mailbox that the PST file will be imported to. You can't specify a public folder because the PST Import Service doesn't support importing PST files to public folders. To import a PST file to an inactive mailbox, you have to specify the mailbox GUID for this parameter. To obtain this GUID, run the following PowerShell command in Exchange Online: `Get-Mailbox <identity of inactive mailbox> -InactiveMailboxOnly | FL Guid` > [!NOTE]> Sometimes, you may have multiple mailboxes with the same email address, where one mailbox is an active mailbox and the other mailbox is in a soft-deleted (or inactive) state. In these situations, you have to specify the mailbox GUID to uniquely identify the mailbox to import the PST file to. To obtain this GUID for active mailboxes, run the following PowerShell command: `Get-Mailbox <identity of active mailbox> | FL Guid` . To obtain the GUID for soft-deleted (or inactive) mailboxes, run this command: `Get-Mailbox <identity of soft-deleted or inactive mailbox> -SoftDeletedMailbox | FL Guid` .	`annb@contoso.onmicrosoft.com` Or `2d7a87fe-d6a2-40cc-8aff-1ebea80d4ae7`
`IsArchive`	Specifies whether to import the PST file to the user's archive mailbox. There are two options: **FALSE** Imports the PST file to the user's primary mailbox. **TRUE** Imports the PST file to the user's archive mailbox. This assumes that the user's archive mailbox is enabled. If you set this parameter to `TRUE` and the user's archive mailbox isn't enabled, the import for that user will fail. If an import fails for one user (because their archive isn't enabled and this property is set to `TRUE`), the other users in the import job won't be affected. If you leave this parameter blank, the PST file is imported to the user's primary mailbox. **Note:** To import a PST file to a cloud-based archive mailbox for a user whose primary mailbox is on-premises, just specify `TRUE` for this parameter and specify the email address for the user's on-premises mailbox for the `Mailbox` parameter.	`FALSE` Or `TRUE`
`TargetRootFolder`	Specifies the mailbox folder that the PST file is imported to. If you leave this parameter blank, the PST will be imported to a new folder named **Imported** located at the root level of the mailbox (the same level as the Inbox folder and the other default mailbox folders). If you specify `/`, items in the PST file will be imported directly in to the user's Inbox folder. If you specify `/<foldername>`, items in the PST file will be imported to a folder named *<foldername>*. For example, if you use `/ImportedPst`, items would be imported to a folder named **ImportedPst**. This folder will be located in the user's mailbox at the same level as the Inbox folder.	(leave blank) Or `/` Or `/ImportedPst`

PARAMETER	DESCRIPTION	EXAMPLE
`ContentCodePage`	This optional parameter specifies a numeric value for the code page to use for importing PST files in the ANSI file format. This parameter is used for importing PST files from Chinese, Japanese, and Korean (CJK) organizations because these languages typically use a double byte character set (DBCS) for character encoding. If this parameter isn't used to import PST files for languages that use DBCS for mailbox folder names, the folder names are often garbled after they're imported. For a list of supported values to use for this parameter, see Code Page Identifiers. > [!NOTE]> As previously stated, this is an optional parameter and you don't have to include it in the CSV file. Or you can include it and leave the value blank for one or more rows.	(leave blank) Or `932` (which is the code page identifier for ANSI/OEM Japanese)
`SPFileContainer`	For PST Import, leave this parameter blank.	Not applicable
`SPManifestContainer`	For PST Import, leave this parameter blank.	Not applicable
`SPSiteUrl`	For PST Import, leave this parameter blank.	Not applicable

Step 4: Create a PST Import job in Office 365

The next step is to create the PST Import job in the Import service in Office 365. As previously explained, you submit the PST Import mapping file that you created in Step 3. After you create the job, the Import service will use the information in the mapping file to import the PST files to the specified user mailbox after the PST files are copied from the hard drive to the Azure Storage area and you create and start the import job.

1. Go to https://protection.office.com and sign in using the credentials for an administrator account in your organization.
2. In the left pane of the Security & Compliance Center, click **Information governance** > **Import** > **Import PST files**.
3. On the **Import PST files** page, click **+ New import job**.

> **NOTE**
> As previously stated, you have to be assigned the appropriate permissions to access the **Import** page in the Security & Compliance Center.

4. Type a name for the PST import job, and then click **Next**. Use lowercase letters, numbers, hyphens, and underscores. You can't use uppercase letters or include spaces in the name.
5. On the **Choose import job type** page, click **Ship hard drives to one of our physical locations** and then click **Next**.

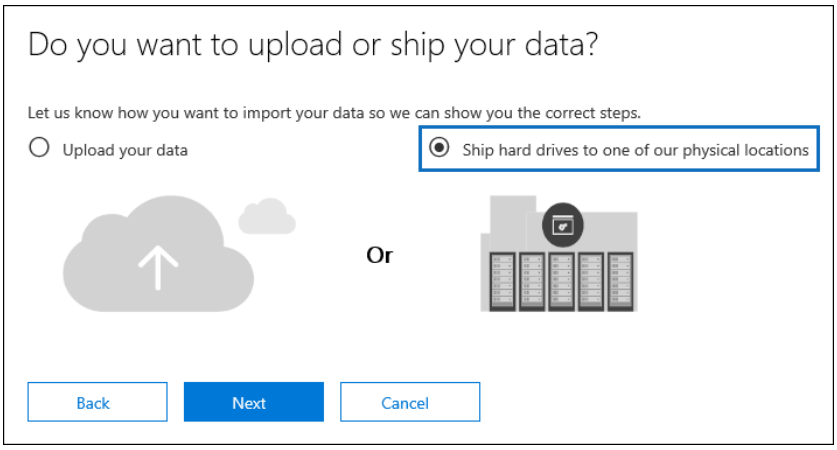

6. In step 6, click the **I've prepared my hard drives and have access to the necessary drive journal**

files and **I have access to the mapping file** check boxes, and then click **Next**.

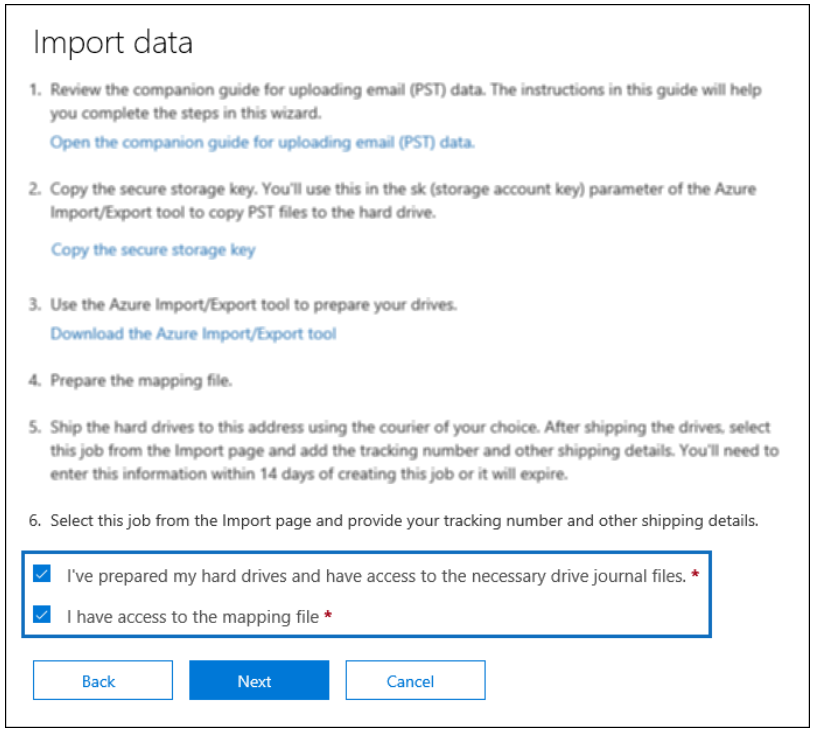

7. On the **Select the drive file** page, click **Select drive file**, and then go to the same folder where the WAImportExport.exe tool is located. The journal file that was created in Step 2 was copied to this folder.

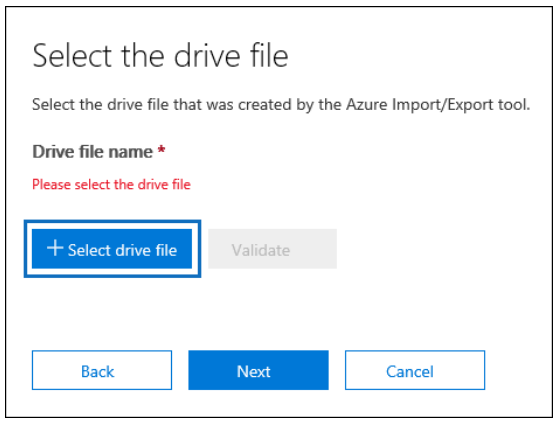

8. Select the journal file; for example, `PSTHDD1.jrn`.

> **TIP**
> When you ran the WAImportExport.exe tool in Step 2, the name of the journal file was specified by the `/j:` parameter.

9. After the name of the drive file appears under **Drive file name**, click **Validate** to check your drive file for errors.

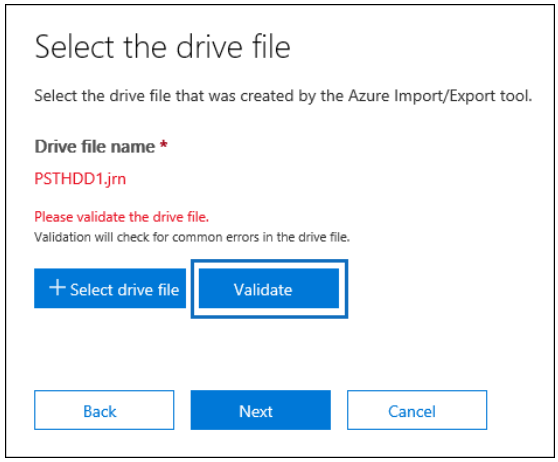

The drive file has to be successfully validated to create a PST Import job. Note that the file name is

changed to green after it's successfully validated. If the validation fails, click the **View log** link. A validation error report is opened, with an error message with information about why the file failed.

> **NOTE**
> You must add and validate a journal file for each hard drive you ship to Microsoft.

10. After adding and validating a journal file for each hard drive that you ship to Microsoft, click **Next**.
11. Click **+ Select mapping file** to submit the PST Import mapping file that you created in Step 3.

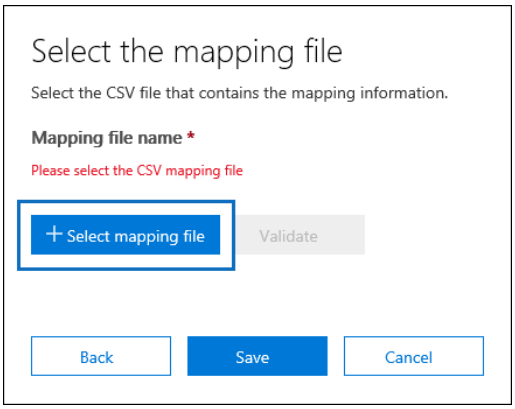

12. After the name of the CSV file appears under **Mapping file name**, click **Validate** to check your CSV file for errors.

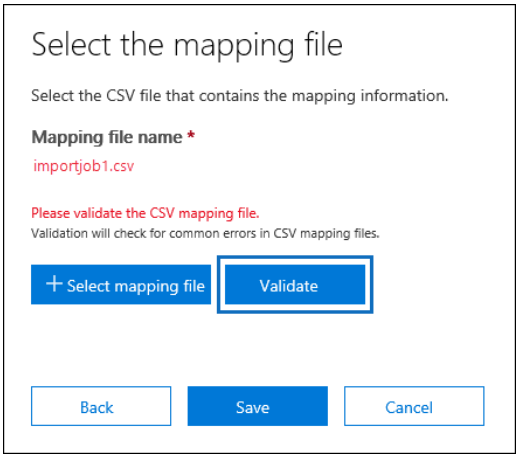

    The CSV file has to be successfully validated to create a PST Import job. Note that the file name is changed to green after it's successfully validated. If the validation fails, click the **View log** link. A validation error report is opened, with an error message for each row in the file that failed.

13. After the PST mapping file is successfully validated, click **Next**.
14. On the **Provide contact information** page, type your contact information in the applicable boxes.

    The address for the Microsoft location that you ship your hard drives to is displayed. This address is auto-generated based on your Microsoft datacenter location. Copy this address to a file or take a screenshot.

15. Read the terms and conditions document, click the checkbox, and then click **Save** to submit the import job.

    When the import job is successfully created, a status page is displayed that explains the next steps of the drive shipping process.

16. On the **Import PST files** page, click **Refresh** to displayed the new drive shipping import job in the list of import jobs. The status is set to **Waiting for tracking number**. You can also click the import job to display the status flyout page, which contains more detailed information about the import job.

Step 5: Ship the hard drive to Microsoft

The next step is to ship the hard drive to Microsoft, and then provide the tracking number for the shipment and return shipment information for the drive shipping job. After the drive is received by Microsoft, it will take between 7 and 10 business days for data center personnel to upload your PST files to the Azure Storage area for your organization.

> **NOTE**
> If you don't provide the tracking number and return shipment information within 14 days of creating the import job, the import job will be expired. If this happens, you'll have to create a new drive shipping import job (see Step 4: Create a PST Import job in Office 365) and re-submit the drive file and the PST import mapping file.

Ship the hard drive

Keep the following things in mind when you ship hard drives to Microsoft:

- Don't ship the SATA-to-USB adapter; you only have to ship the hard drive.
- Package the hard drive properly; for example, use an anti-static bag or bubble wrap.
- Use a delivery carrier of your choice to ship the hard drive to Microsoft.
- Ship the hard drive to the address for the Microsoft location that was displayed when you created the import job in Step 4. Be sure to include "Office 365 Import Service" in the ship-to address.
- After you ship the hard drive, be sure to write down the name of the delivery carrier and the tracking number. You'll provide these in the next step.

Enter the tracking number and other shipping information

After you've shipped the hard drive to Microsoft, complete the following procedure on the Import service page.

1. Go to https://protection.office.com and sign in using the credentials for an administrator account in your organization.
2. In the left pane, click **Information governance** > **Import** > **Import PST files**.
3. On the **Import PST files** page, click the job for the drive shipment that you want to enter the tracking number for.
4. On the status flyout page, click **Enter tracking number**.
5. Provide the following shipping information:
6. **Delivery carrier** Type the name of the delivery carrier that you used to ship the hard drive to Microsoft.
7. **Tracking number** Type the tracking number for the hard drive shipment.
8. **Return carrier account number** Type your organization's account number for the carrier that listed under **Return carrier**. Microsoft uses (and charges) this account to ship your hard drive back to you. Organizations in the USA and Europe, must have an account with FedEx. Organizations in Asia and the rest of the world, must have an account with DHL.
9. Click **Save** to save this information for the import job.

   On the **Import PST files** page, click **Refresh** to update the information for your drive shipping import job. Notice that status is now set to **Drives in transit**.

Step 6: Filter data and start the PST Import job

After your hard drive is received by Microsoft, the status for the import job on the **Import PST files** page will change to **Drives received**. Data center personnel use the information in the journal file to upload your PST files to the Azure Storage area for your organization. At this point, the status changes to **Import in-progress**. As previously stated, it will take between 7 and 10 business days after receiving your hard drive to upload the PST files.

After PST files are uploaded to Azure, the status is changed to **Analysis in progress**. This indicates that Microsoft 365 is analyzing the data in the PST files (in a safe and secure manner) to identify the age of the items and the different message types included in the PST files. When the analysis is completed and the data is ready to import, the status for the import job is changed to **Analysis completed**. At this point, you have the option to import all the data contained in the PST files or you can trim the data that's imported by setting filters that control what data gets imported.

1. Go to https://protection.office.com and sign in using the credentials for an administrator account in your organization.
2. In the left pane, click **Information governance** > **Import** > **Import PST files**.
3. On the **Import PST files** page, click **Ready to import to Office 365** for the import job that you created in Step 4.

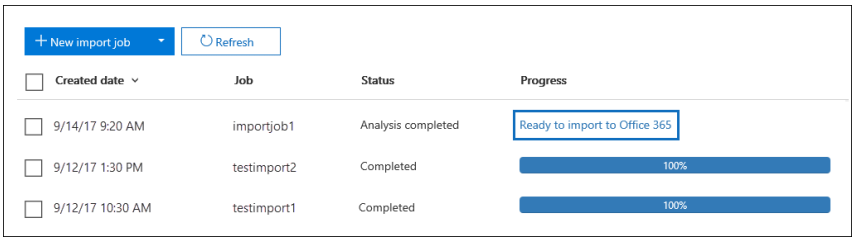

A fly out page is displayed with information about the PST files and other information about the import job.

4. Click **Import to Office 365**.

5. The **Filter your data** page is displayed. It contains the data insights resulting from the analysis performed on the PST files by Office 365, including information about the age of the data. At this point, you have the option to filter the data that will be imported or import all the data as is.

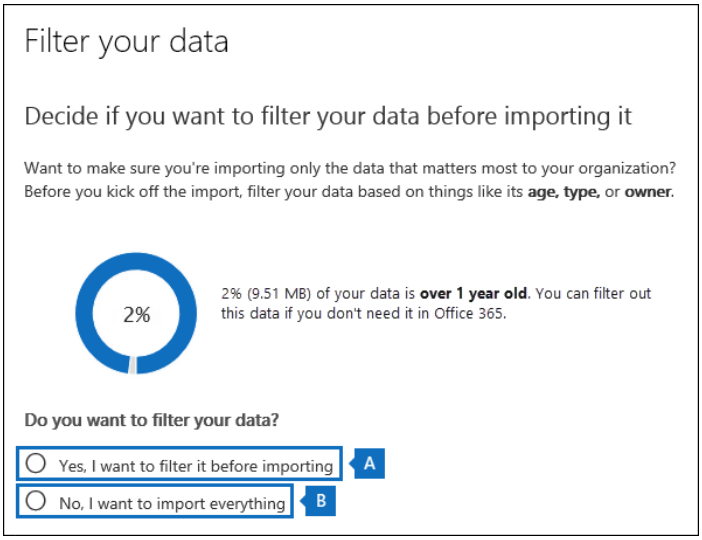

6. Do one of the following:

   a. To trim the data that you import, click **Yes, I want to filter it before importing**.

   For detailed step-by-step instructions about filtering the data in the PST files and then starting the import job, see Filter data when importing PST files to Office 365.

   Or

   b. To import all data in the PST files, click **No, I want to import everything,** and click **Next**.

7. If you chose to import all the data, click **Import data** to start the import job.

   The status of the import job is displayed on the **Import PST files** page. Click **Refresh** to update the status information that's displayed in the **Status** column. Click the import job to display the status flyout page, which displays status information about each PST file being imported. When the import is complete and PST files have been imported to user mailboxes, the status will be changed to **Completed**.

View a list of the PST files uploaded to Microsoft 365

You can install and use the Microsoft Azure Storage Explorer (which is a free, open-source tool) to view the list of the PST files that we're uploaded (by Microsoft data center personnel) to the Azure Storage area for your organization. You can do this to verify that PST files from the hard drives that you sent to Microsoft were successfully uploaded to the Azure Storage area.

The Microsoft Azure Storage Explorer is in Preview.

Important: You can't use the Azure Storage Explorer to upload or modify PST files. The only supported method for importing PST files to Microsoft 365 is to use AzCopy. Also, you can't delete PST files that you've uploaded to the Azure blob. If you try to delete a PST file, you receive an error about not having the required permissions. All PST files are automatically deleted from your Azure Storage area. If there are no import jobs in progress, then all PST files in the ** ingestiondata ** container are deleted 30 days after the most recent import job was created.

To install the Azure Storage Explorer and connect to your Azure Storage area:

1. Perform the following steps to get the Shared Access Signature (SAS) URL for your organization. This URL is a combination of the network URL for the Azure Storage location in the Microsoft cloud for your organization and an SAS key. This key provides you with the necessary permissions to access your organization's Azure Storage location.

2. Go to https://protection.office.com/ and sign in using the credentials for an administrator account in your

organization.

3. In the left pane of the Security & Compliance Center, click **Information governance** > **Import** > **Import PST files**.
4. On the **Import PST files** page, click **+ New import job**.
5. In the import job wizard, type a name for the PST import job, and then click **Next**. Use lowercase letters, numbers, hyphens, and underscores. You can't use uppercase letters or include spaces in the name.
6. On the **Choose import job type** page, click **Upload your data**, and then click **Next**.
7. In step 2, click **Show network upload SAS URL**.
8. After the URL is displayed, copy it and save it to a file. Be sure to copy the entire URL.

   > **IMPORTANT**
   > Be sure to take precautions to protect the SAS URL. This can be used by anyone to access the Azure storage area for your organization.

9. Click **Cancel** to close the import job wizard.
10. Download and install the Microsoft Azure Storage Explorer tool.
11. Start the Microsoft Azure Storage Explorer, right-click **Storage Accounts** in the left pane, and then click **Connect to Azure Storage**.

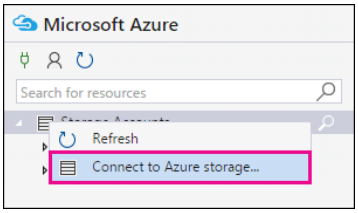

12. Click **Use a shared access signature (SAS) URI or connection string** and click **Next**.
13. Click **Use a SAS URI**, paste the SAS URL that you obtained in step 1 in to in the box under **URI**, and then click **Next**.
14. On the **Connection summary** page, you can review the connection information, and then click **Connect**.

    The **ingestiondata** container is opened. It contains the PST files from your hard drive. The **ingestiondata** container is located under **Storage Accounts** > **(SAS-Attached Services)** > **Blob Containers**.

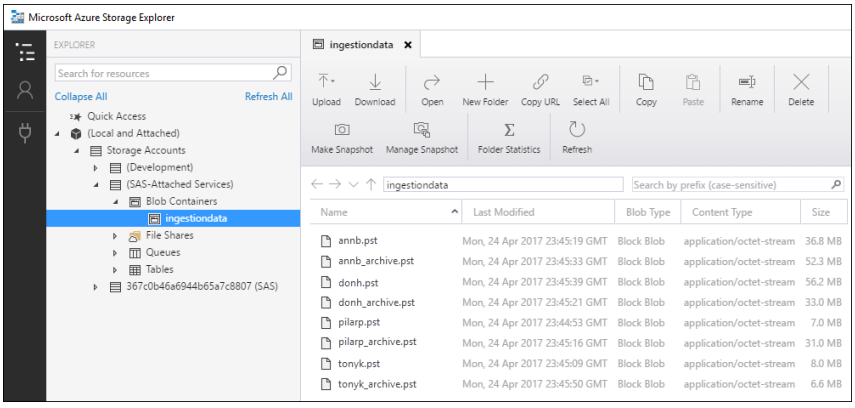

15. When you're finished using the Microsoft Azure Storage Explorer, right-click **ingestiondata**, and then click **Detach** to disconnect from your Azure Storage area. Otherwise, you'll receive an error the next time you try to attach.

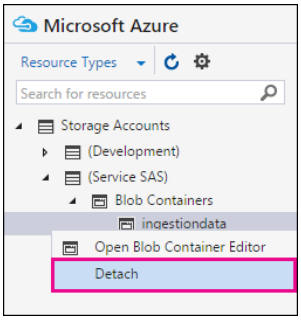

Troubleshooting tips

- **What happens if the import job fails because of errors in the PST Import CSV mapping file?** If an import job fails because of errors in the mapping file, you don't have to reship the hard drive to Microsoft to create an import job. That's because the PST files from the hard drive that you submitted for the drive shipping import job have already been uploaded to the Azure Storage area for your organization. In this case, you only have to fix the errors in the PST Import CSV mapping file, and then create a new "network upload" import job and submit the revised CSV mapping file. To create and start a new network upload import job, see Step 5: Create a PST Import job in Microsoft 365 and Step 6: Filter data and start the PST Import job in the topic "Use network upload to import PST files to Office 365."

  > **NOTE**
  > To help you troubleshoot the PST Import CSV mapping file, use the Azure Storage Explorer tool to view the folder structure in the **ingestiondata** container for the PST files from your hard drive that were uploaded to the Azure storage area. Mapping file errors are typically caused by an incorrect value in the FilePath parameter. This parameter specifies the location of a PST file in the Azure storage area. See the description of the FilePath parameter in table in Step 3. As previously explained, the location of PST files in the Azure storage area was specified by the `/dstdir:` parameter when you ran the WAImportExport.exe tool in Step 2.

More information

- Drive shipping is an effective way to import large amounts of archival messaging data to Microsoft 365 to take advantage of the compliance features that are available to your organization. After archival data is imported to user mailboxes, you can:
  - Enable archive mailboxes and auto-expanding archiving to give users more mailbox storage space for the data.
  - Place mailboxes on Litigation Hold to retain the data.
  - Use Microsoft eDiscovery tools to search the data.
  - Apply Microsoft 365 retention policies to control how long the data is retained, and what action to take after the retention period expires.
  - Search the audit log for events related to this data.
  - Import data to inactive mailboxes to archive data for compliance purposes.
  - Protect your organization against data loss of sensitive information.
- Here's an example of the secure storage account key and a BitLocker encryption key. This example also contains the syntax for the WAImportExport.exe command that you run to copy PST files to a hard drive. Be sure to take precautions to protect these just like you would protect passwords or other security-related information.

```
Secure storage account key:

yaNIIs9Uy5g25Yoak+LlSHfqVBGOeNwjqtBEBGqRMoidq6/e5k/VPkjOXdDIXJHxHvNoNoFH5NcVUJXHwu9ZxQ==

BitLocker encryption key:

397386-221353-718905-535249-156728-127017-683716-083391

COMMAND SYNTAX

First time:

WAImportExport.exe PrepImport /j:<Name of journal file> /t:<Drive letter> /id:<Name of session>
/srcdir:<Location of PST files> /dstdir:<PST file path> /sk:<Storage account key> /blobtype:BlockBlob
/encrypt /logdir:<Log file location>

Subsequent times:

WAImportExport.exe PrepImport /j:<Name of journal file> /id:<Name of new session> /srcdir:<Location
of PST files> /dstdir:<PST file path> /blobtype:BlockBlob

EXAMPLES

First time:

WAImportExport.exe PrepImport /j:PSTHDD1.jrn /t:f /id:driveship1 /srcdir:"\\FILESERVER1\PSTs"
/dstdir:"ingestiondata/"
/sk:"yaNIIs9Uy5g25Yoak+LlSHfqVBGOeNwjqtBEBGqRMoidq6/e5k/VPkjOXdDIXJHxHvNoNoFH5NcVUJXHwu9ZxQ=="
/blobtype:BlockBlob /encrypt /logdir:"c:\users\admin\desktop\PstImportLogs"

Subsequent times:

WAImportExport.exe PrepImport /j:PSTHDD1.jrn /id:driveship2 /srcdir:"\\FILESERVER1\PSTs\SecondBatch"
/dstdir:"ingestiondata/" /blobtype:BlockBlob
```

- As previously explained, the Office 365 Import service turns on the retention hold setting (for an

indefinite duration) after PST files are imported to a mailbox. This means the *RentionHoldEnabled* property is set to `True` so that the retention policy assigned to the mailbox won't be processed. This gives the mailbox owner time to manage the newly imported messages by preventing a deletion or archive policy from deleting or archiving older messages. Here are some steps you can take to manage this retention hold:

- After a certain period of time, you can turn off the retention hold by running the `Set-Mailbox -RetentionHoldEnabled $false` command. For instructions, see Place a mailbox on retention hold.
- You can configure the retention hold so that it's turned off on some date in the future. You do this by running the `Set-Mailbox -EndDateForRetentionHold <date>` command. For example, assuming that today's date is June 1, 2016 and you want the retention hold turned off in 30 days, you would run the following command: `Set-Mailbox -EndDateForRetentionHold 7/1/2016`. In this scenario, you would leave the *RentionHoldEnabled* property set to *True*. For more information, see Set-Mailbox.
- You can change the settings for the retention policy that's assigned to the mailbox so that older items that were imported won't be immediately deleted or moved to the user's archive mailbox. For example, you could lengthen the retention age for a deletion or archive policy that's assigned to the mailbox. In this scenario, you would turn off the retention hold on the mailbox after you changed the settings of the retention policy. For more information, see Set up an archive and deletion policy for mailboxes in your organization.

Filter data when importing PST files

11/2/2020 • 6 minutes to read • Edit Online

Use the new Intelligent Import feature in the Office 365 Import service to filter the items in PST files that actually get imported to the target mailboxes. Here's how it works:

- After you create and submit a PST import job, PST files are uploaded to an Azure storage area in the Microsoft cloud.
- Microsoft 365 analyzes the data in the PST files, in a safe and secure manner, by identifying the age of the mailbox items and the different message types included in the PST files.
- When the analysis is complete and the data is ready to import, you have the option to import all data in the PST files as is or trim the data that's imported by setting filters that control what data gets imported. For example, you can choose to:
  - Import only items of a certain age.
  - Import selected message types.
  - Exclude messages sent or received by specific people.
- After you configure the filter settings, Microsoft 365 imports only the data that meets the filtering criteria to the target mailboxes specified in the import job.

The following graphic shows the Intelligent Import process, and highlights the tasks you perform and the tasks performed by Office 365.

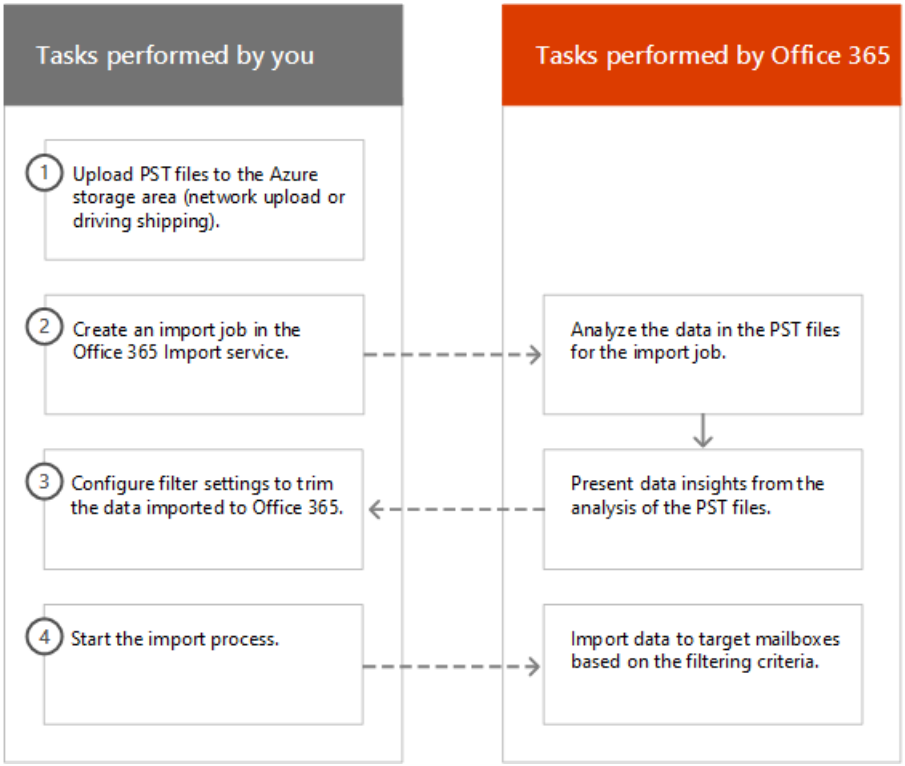

Create a PST import job

- The steps in this topic assume that you've created a PST import job in the Office 365 Import service by using network upload or drive shipping. For step-by-step instructions, see one of the following topics:

- Use network upload to import PST files to Office 365
- Use drive shipping to import PST files to Office 365

- After you create an import job by using network upload, the status for the import job on the Import page in the Security & Compliance Center is set to **Analysis in progress**, which means that Microsoft 365 is analyzing the data in the PST files that you uploaded. Click **Refresh** to update the status for the import job.
- For drive shipping import jobs, the data will be analyzed by Microsoft 365 after Microsoft datacenter personnel receive your hard drive and upload the PST files to the Azure storage area for your organization.

Filter data that gets imported to mailboxes

After you've created a PST import job, follow these steps to filter the data before you import it to Office 365.

1. Go to https://protection.office.com/ and sign in using the credentials for an administrator account in your organization.
2. Click **Information governance** > **Import** > **Import PST files**.

   The import jobs for your organization are listed on the **Import PST files** page. Note that the **Analysis completed** value in the **Status** column indicates the import jobs that have been analyzed by Microsoft 365 and are ready for you to import.

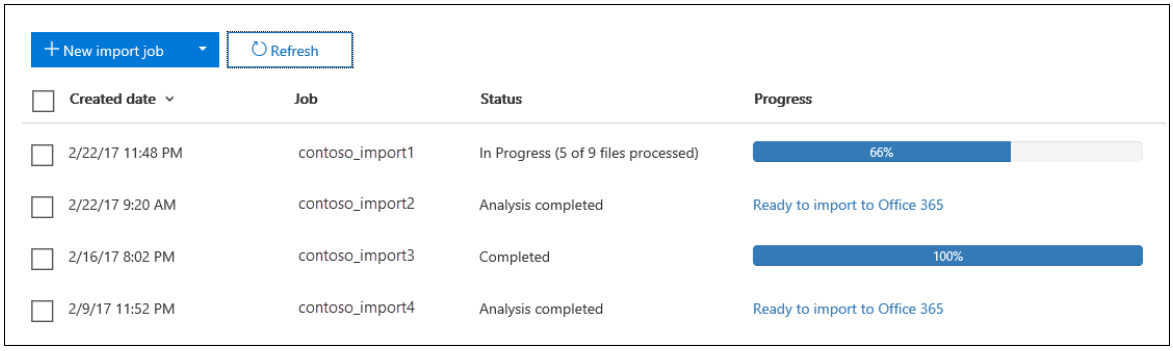

3. Click **Ready to import to Office 365** for the import job that you want to complete.

   A fly out page is displayed with information about the PST files and other information about the import job.

4. Click **Import to Office 365**.

   The **Filter your data** page is displayed. It contains data insights about the data in the PST files for the import job, including information about the age of the data.

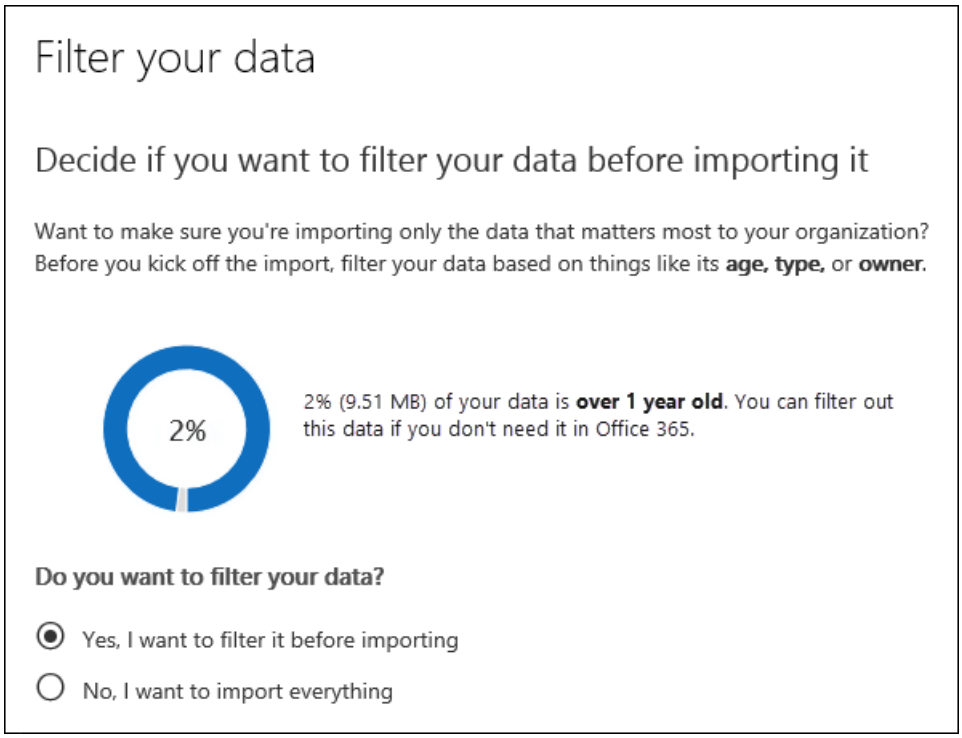

5. Based on whether or not you want to trim the data that's imported to Microsoft 365, under **Do you want to filter your data?**, do one of the following:

   a. Click **Yes, I want to filter it before importing** to trim the data that you import, and then click **Next**.

   The **Import data to Office 365 page** page is displayed with detailed data insights from the analysis that Microsoft 365 performed.

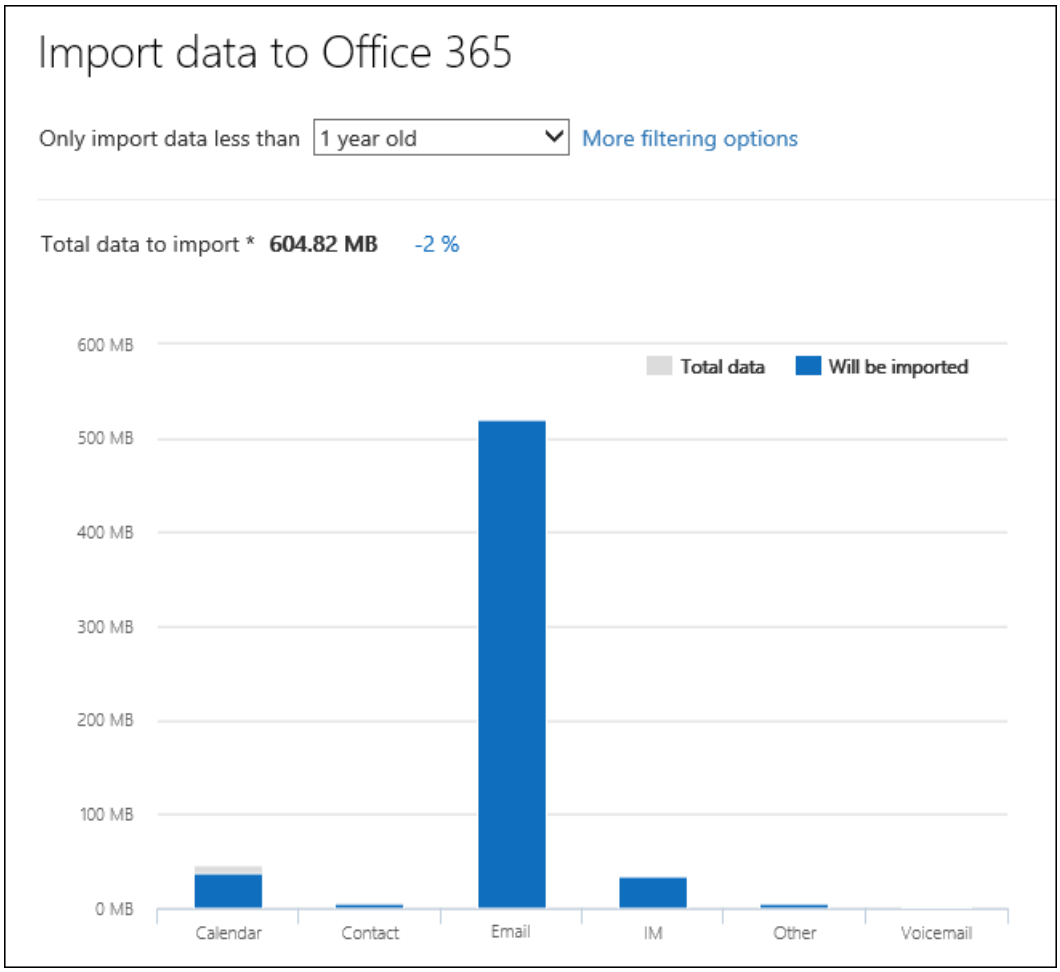

The graph on this page shows the amount of data that will be imported. Information about each message

type found in the PST files is displayed in the graph. You can hover the cursor over each bar to display specific information about that message type. There is also a drop-down list with different age values based on the analysis of the PST files. When you select an age in the drop-down list, the graph is updated to show how much data will be imported for the selected age.

b. To configure addition filters to reduce the amount of data that's imported, click **More filtering options**.

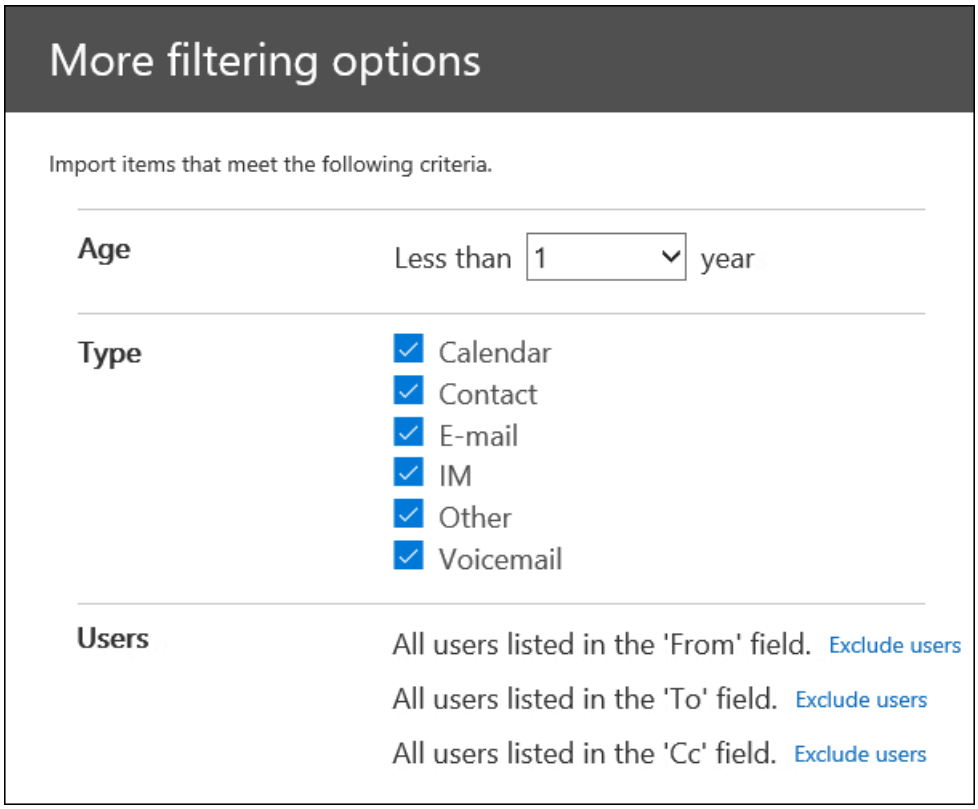

You can configure these filters:

- **Age** - Select an age so only items that are newer than the specified age will be imported. See the More information section for a description about how Microsoft 365 determines the age buckets for the **Age** filter.
- **Type** - This section shows all the message types that were found in the PST files for the import job. You can uncheck a box next to a message type that you want to exclude. Note that you can't exclude the Other message type. See the More information section for a list of mailbox items that are included in the Other category.
- **Users** - You can exclude messages that are sent or received by specific people. To exclude people who appear in the From: field, To: field, or the Cc: field of messages, click **Exclude users** next to that recipient type. Type the email address (SMTP address) of the person, click **Add+** to add them to the list of excluded users for that recipient type, and then click **Save** to save the list of excluded users.

  > **NOTE**
  > Microsoft 365 doesn't show data insights that result from setting the **People** filter. However, if you set this filter to exclude messages sent or received by specific people, those messages will be excluded during the actual import process.

c. Click **Apply** in the **More filtering options** fly out page to save your filter settings.

The data insights on the **Import data to Office 365** page are updated based on your filter settings,

including the total amount of data that will be imported based on the filter settings. Note that a summary of the filter settings is also shown. You can click **Edit** next to a filter to change the setting if necessary.

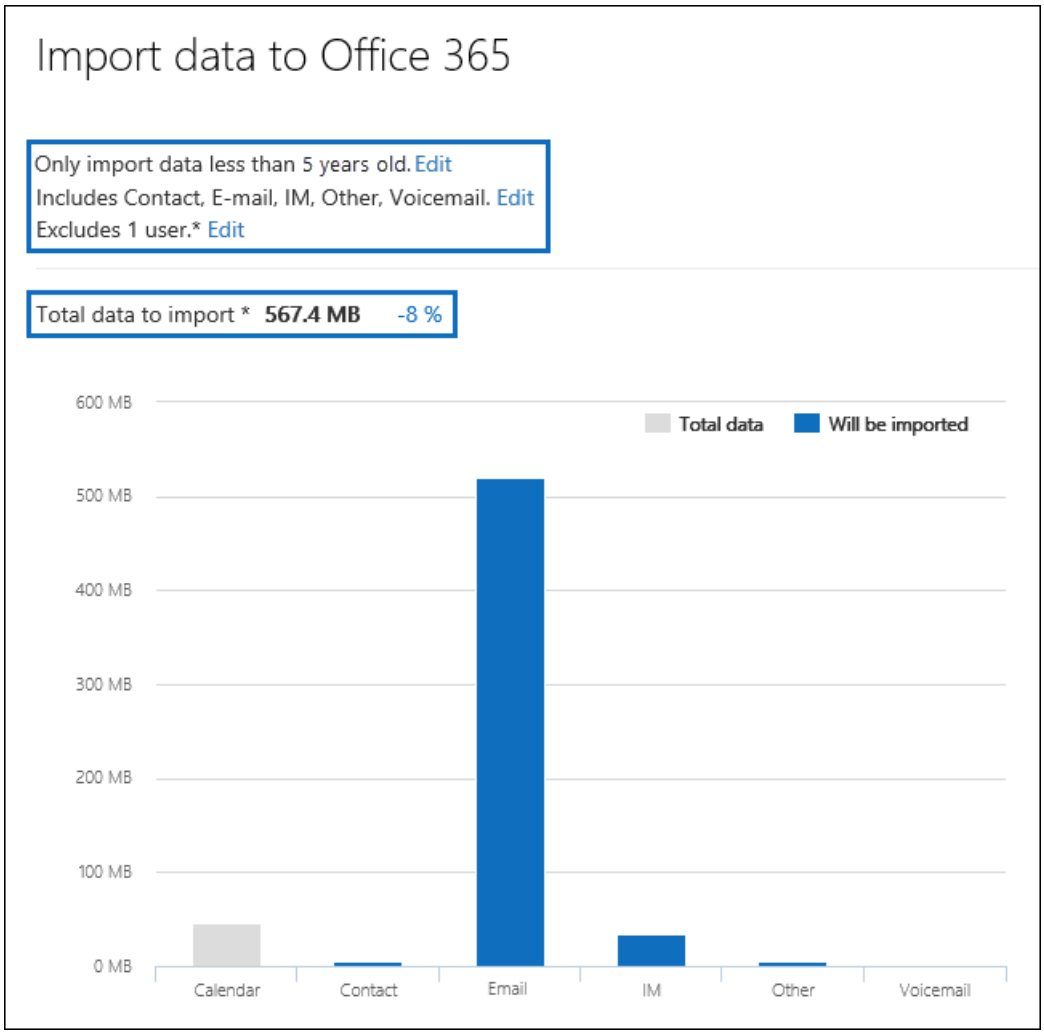

d. Click **Next**.

A status page is displayed showing your filter settings. Again, you can edit any of the filter settings.

e. Click **Import data** to start the import . Note that the total amount of data that will be imported is displayed.

Or

a. Click **No, I want to import everything** to import all data in the PST files to Office 365, and then click **Next**.

b. On the **Import data to Office 365** page, click **Import data** to start the import. Note that the total amount of data that will be imported is displayed.

6. On the **Import PST files** page, click **Refresh** . The status for the import job is displayed in the **Status** column.
7. Click the import the job to display more detailed information, such as the status for each PST file and the filter settings that you configured.

More information

- How does Microsoft 365 determine the increments for the age filter? When Microsoft 365 analyzes a PST file, it looks at the sent or received time stamp of each item (if an item has both a sent and received timestamp, the oldest date is selected). Then Microsoft 365 looks at the year value for that timestamp and

compares it to the current date to determine the age of the item. These ages are then used as the values in the drop-down list for the **Age** filter. For example, if a PST file has messages from 2016, 2015, and 2014, then values in the **Age** filter would be **1 year**, **2 years**, and **3 years**.

- The following table lists the message types that are included in the **Other** category in the **Type** filter on the **More options** fly out page (see Step 5b in the previous procedure). Currently, you can't exclude items in the "Other" category when you import PSTs to Office 365.

MESSAGE CLASS ID	MAILBOX ITEMS THAT USE THIS MESSAGE CLASS
IPM.Activity	Journal entries
IPM.Document	Documents and files (not attached to an email message)
IPM.File	(same as IPM.Document)
IPM.Note.IMC.Notification	Reports sent by Internet Mail Connect, which is the Exchange Server gateway to the Internet
IPM.Note.Microsoft.Fax	Fax messages
IPM.Note.Rules.Oof.Template.Microsoft	Out-of-office auto-reply messages
IPM.Note.Rules.ReplyTemplate.Microsoft	Replies sent by an inbox rule
IPM.OLE.Class	Exceptions for a recurring series
IPM.Recall.Report	Message recall reports
IPM.Remote	Remote mail messages
IPM.Report	Item status reports

FAQ about importing PST files

11/2/2020 • 14 minutes to read • Edit Online

This article is for administrators. Do you want to import PST files to your own mailbox? See Import email, contacts, and calendar from an Outlook .pst file

Here are some frequently asked questions about using the Office 365 Import Service to bulk-import PST files to Microsoft 365 mailboxes. For more information about how to import PST files, see Overview of importing PST files to Office 365.

Using network upload to import PST files

For step-by-step instructions, see Use network upload to import PST files to Office 365.

What permissions are required to create import jobs in the Office 365 Import Service?

You have to be assigned the Mailbox Import Export role in Exchange Online to import PST files to Microsoft 365 mailboxes. By default, this role isn't assigned to any role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group. Or you can create a new role group, assign the Mailbox Import Export role, and then add yourself or other users as a member. For more information, see the "Add a role to a role group" or the "Create a role group" sections in Manage role groups in Exchange Online.

Additionally, to create import jobs in the Security & Compliance Center, one of the following must be true:

- You have to be assigned the Mail Recipients role in Exchange Online. By default, this role is assigned to the Organization Management and Recipient Management roles groups.

  Or

- You have to be a global administrator in your organization.

> **TIP**
> Consider creating a new role group in Exchange Online that's specifically intended for importing PST files to Office 365. For the minimum level of privileges required to import PST files, assign the Mailbox Import Export and Mail Recipients roles to the new role group, and then add members.

Where is network upload available?

Network upload is currently available in these regions: United States, Canada, Brazil, the United Kingdom, France, Germany, Switzerland, Norway, Europe, India, East Asia, Southeast Asia, Japan, Republic of Korea, Australia, and United Arab Emirates (UAE). Network upload will be available in more regions soon.

What is the pricing for importing PST files by using network upload?

Using network upload to import PST files is free.

This also means that after PST files are deleted from the Azure Storage area, they're no longer displayed in the list of files for a completed import job in the Microsoft 365 admin center. Although an import job might still be listed on the **Import data to Office 365** page, the list of PST files might be empty when you view the details of older import jobs.

What version of the PST file format is supported for importing to Office 365?

There are two versions of the PST file format: ANSI and Unicode. We recommend importing files that use the Unicode PST file format. However, files that use the ANSI PST file format, such as those for languages that use a double-byte character set (DBCS), can also be imported to Office 365. For more information about importing ANSI PST files, see Step 4 in Use network upload to import your organization's PST files to Office 365.

Additionally, PST files from Outlook 2007 and later versions can be imported to Office 365.

After I upload my PST files to the Azure Storage area, how long are they kept in Azure before they're deleted?

When you use the network upload method to import PST files, you upload them to an Azure blob container named `ingestiondata`. If there are no import jobs in progress on the **Import PST files** page in the Security & Compliance Center), then all PST files in the `ingestiondata` container in Azure are deleted 30 days after the most recent import job was created in the Security & Compliance Center. That also means you have to create a new import job in the Security & Compliance Center (described in Step 5 in the network upload instructions) within 30 days of uploading PST files to Azure.

This also means that after PST files are deleted from the Azure Storage area, they're no longer displayed in the list of files for a completed import job in the Security & Compliance Center. Although an import job might still be listed on the **Import PST files** page in the Security & Compliance Center, the list of PST files might be empty when you view the details of older import jobs.

How long does it take to import a PST file to a mailbox?

It depends on the capacity of your network, but it typically takes several hours for each terabyte (TB) of data to be uploaded to the Azure Storage area for your organization. After the PST files are copied to the Azure Storage area, a PST file is imported to a Microsoft 365 mailbox at a rate of at least 24 GB per day. If this rate doesn't meet your needs, you might consider other methods for migrating email data to Office 365. For more information, see Ways to migrate multiple email accounts to Office 365.

If different PST files are imported to different target mailboxes, the import process occurs in parallel; in other words, each PST/mailbox pair is imported simultaneously. Likewise, if multiple PST files are imported to the same mailbox, they will be simultaneously imported.

How does the PST import process handle duplicate email items?

The PST import process checks for duplicate items and doesn't copy the items from a PST file to the mailbox or archive if a matching item exists in the target folder in the target mailbox or target archive. If you reimport the same PST file and specify a different target folder (using the TargetRootFolder property in the PST import mapping file) than the one you specified in a previous import job, all items in the PST file will be reimported.

Is there a message size limit when importing PST files?

Yes. If a PST file contains a mailbox item that is larger than 150 MB, the item will be skipped during the import process.

Are message properties, such as when the message was sent or received, the list of recipients and other properties, preserved when PST files are imported to a Microsoft 365 mailbox?

Yes. The original message metadata isn't changed during the import process.

Is there a limit to the number of levels in a folder hierarchy for a PST file that I want to import to a mailbox?

Yes. You can't import a PST file that has 300 or more levels of nested folders.

Can I use network upload to import PST files to an inactive mailbox in Office 365?

Yes, this capability is now available.

Can I use network upload to import PST files to an online archive mailbox in an Exchange hybrid deployment?

Yes, this capability is now available.

Can I use network upload to import PST files to public folders in Exchange Online?

No, you can't import PST files to public folders.

Using drive shipping to import PST files

For step-by-step instructions, see Use drive shipping to import PST files to Office 365.

What permissions are required to create import jobs in the Office 365 Import Service?

You have to be assigned the Mailbox Import Export role to import PST files to Microsoft 365 mailboxes. By default, this role isn't assigned to any role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group. Or you can create a new role group, assign the Mailbox Import Export role, and then add yourself or other users as a member. For more information, see the "Add a role to a role group" or the "Create a role group" sections in Manage role groups in Exchange Online.

Additionally, to create import jobs in the Security & Compliance Center, one of the following must be true:

- You have to be assigned the Mail Recipients role in Exchange Online. By default, this role is assigned to the Organization Management and Recipient Management roles groups.

  Or

- You have to be a global administrator in your organization.

> **TIP**
> Consider creating a new role group in Exchange Online that's specifically intended for importing PST files to Office 365. For the minimum level of privileges required to import PST files, assign the Mailbox Import Export and Mail Recipients roles to the new role group, and then add members.

Where is drive shipping available?

Drive shipping is currently available in the United States, Canada, Brazil, the United Kingdom, Europe, India, East Asia, Southeast Asia, Japan, Republic of Korea, and Australia. Drive shipping will be available in more regions soon.

> **NOTE**
> At this time, drive shipping to import PST files is not available in Germany and Switzerland. This FAQ will be updated when drive shipping is available in these countries.

What commercial licensing agreements support drive shipping?

Drive shipping to import PST files to Microsoft 365 is available through a Microsoft Enterprise Agreement (EA). Drive shipping isn't available through a Microsoft Products and Services Agreement (MPSA).

What is the pricing for using drive shipping to import PST files to Microsoft 365?

The cost to use drive shipping to import PST files to Microsoft 365 mailboxes is $2 USD per GB of data. For example, if you ship a hard drive that contains 1,000 GB (1 TB) of PST files, the cost is $2,000 USD. You can work with a partner to pay the import fee. For information about finding a partner, see Find your Microsoft partner or reseller.

What kind of hard drives are supported for drive shipping?

Only 2.5-inch solid-state drives (SSDs) or 2.5-inch or 3.5-inch SATA II/III internal hard drives are supported for use with the Office 365 Import service. You can use hard drives up to 10 TB. For import jobs, only the first data volume on the hard drive will be processed. The data volume must be formatted with NTFS. When copying data to a hard drive, you can attach it directly using a 2.5-inch SSD or 2.5-inch or 3.5-inch SATA II/III connector or you can attach it externally using an external 2.5-inch SSD or 2.5-inch or 3.5-inch SATA II/III USB adaptor.

> **IMPORTANT**
> External hard drives that come with an built-in USB adaptor aren't supported by the Office 365 Import service. Additionally, the disk inside the casing of an external hard drive can't be used. Please don't ship external hard drives.

How many hard drives can I ship for a single import job?

You can ship a maximum of 10 hard drives for a single import job.

After I ship my hard drive, how long does it take to get to the Microsoft data center?

That depends on a few things, such as your proximity to the Microsoft data center and what kind of shipping option you used to ship your hard drive (such as, next-day delivery, two-day delivery, or ground-delivery). With most shippers, you can use the tracking number to track the status of your delivery.

After my hard drive arrives at the Microsoft data center, how long does it take to upload my PST files to Azure?

After your hard drive is received at the Microsoft data center, it will take between 7 to 10 business days to upload the PST files to the Azure Storage area for your organization. The PST files will be uploaded to an Azure blob container named `ingestiondata`.

How long does it take to import a PST file to a mailbox?

After the PST files are uploaded to the Azure Storage area, Microsoft 365 analyzes the data in the PST files (in a safe and secure manner) to identify the age of the items and the different message types included in the PST files. When this analysis is complete, you'll have the option to import all the data in the PST files or set filters to that control what data gets imported. After you start the import job, a PST file is imported to a Microsoft 365 mailbox at a rate of at least 24 GB per day. If this rate doesn't meet your needs, you might consider other methods for importing email data to Office 365. For more information, see Ways to migrate multiple email accounts to Office 365.

If different PST files are imported to different target mailboxes, the import process occurs in parallel; in other words, each PST/mailbox pair is imported simultaneously. Likewise, if multiple PST files are imported to the same mailbox, they will be simultaneously imported.

After Microsoft uploads my PST files to Azure, how long are they kept in Azure before they're deleted?

All PST files in the Azure Storage location for your organization (in blob container named `ingestiondata`), are deleted 30 days after the most recent import job was created on the **Import PST files** page in the Security & Compliance Center.

This also means that after PST files are deleted from the Azure Storage area, they're no longer displayed in the list of files for a completed import job in the Security & Compliance Center. Although an import job might still be listed on the **Import PST files** page in the Security & Compliance Center, the list of PST files might be empty when you view the details of older import jobs.

What version of the PST file format is supported for importing to Microsoft 365?

There are two versions of the PST file format: ANSI and Unicode. We recommend importing files that use the Unicode PST file format. However, files that use the ANSI PST file format, such as those for languages that use a double-byte character set (DBCS), can also be imported to Microsoft 365. For more information about importing ANSI PST files, see Step 3 in Use drive shipping to import PST files to Office 365.

Additionally, PST files from Outlook 2007 and later versions can be imported to Office 365.

Is there a message size limit when importing PST files?

Yes. If a PST file contains a mailbox item that is larger than 150 MB, the item will be skipped during the import process.

How does the PST import process handle duplicate email items?

The PST import process checks for duplicate items and doesn't copy the items from a PST file to the mailbox or archive if a matching item exists in the target folder in the target mailbox or target archive. If you reimport the same PST file and specify a different target folder (using the TargetRootFolder property in the PST import mapping file) than the one you specified in a previous import job, all items in the PST file will be reimported.

Are message properties, such as when the message was sent or received, the list of recipients and other properties, preserved when PST files are imported to a Microsoft 365 mailbox?

Yes. The original message metadata isn't changed during the import process

Is there a limit to the number of levels in a folder hierarchy for a PST file that I want to import to a mailbox?

Yes. You can't import a PST file that has 300 or more levels of nested folders.

Can I use drive shipping to import PST files to an inactive mailbox in Microsoft 365?

Yes, this capability is now available.

Can I use drive shipping to import PST files to an online archive mailbox in an Exchange hybrid deployment?

Yes, this capability is now available.

Can I use drive shipping to import PST files to public folders in Exchange Online?

No, you can't import PST files to public folders.

Can Microsoft wipe my hard drive before they ship it back to me?

No, Microsoft can't wipe hard drives before shipping them back to customers. Hard drives are returned to you in the same state they were in when they were received by Microsoft.

Can Microsoft shred my hard drive instead of shipping it back to me?

No, Microsoft can't destroy your hard drive. Hard drives are returned to you in the same state they were in when they were received by Microsoft.

What courier services are supported for return shipping?

If you're a customer in the United States or Europe, Microsoft uses FedEx to return your hard drive. For all other regions, Microsoft uses DHL.

What are the return shipping costs?

Return shipping costs vary, depending on your proximity to the Microsoft data center that you shipped your hard drive to. Microsoft will bill your FedEx or DHL account to return your hard drive. The cost of return shipping is your responsibility.

Can I use a custom courier shipping service, such as FedEx Custom Shipping, to ship my hard drive to Microsoft?

Yes.

If I have to ship my hard drive to another country, is there anything I need to do?

The hard drive that you ship to Microsoft might have to cross international borders. If this is the case, you're responsible for ensuring that the hard drive and the data it contains are imported and/or exported in accordance with the applicable laws. Before shipping a hard drive, check with your advisors to verify that your drive and data can legally be shipped to the specified Microsoft data center. This will help to ensure that it reaches Microsoft in a timely manner.

Archive third-party data

2/18/2021 • 5 minutes to read • Edit Online

Microsoft 365 lets administrators use data connectors to import and archive third-party data from social media platforms, instant messaging platforms, and document collaboration platforms, to mailboxes in your Microsoft 365 organization. One primary benefit of using data connectors to import and archive third-party data in Microsoft 365 is that you can apply various Microsoft 365 compliance solutions to that after it's been imported. This helps you ensure that your organization's non-Microsoft data is in compliance with the regulations and standards that affect your organization.

Third-party data connectors

The following table lists the third-party data connectors available in the Microsoft 365 compliance center. The table also summarizes the compliance solutions that you can apply to third-party data after you import and archive in Microsoft 365. See the next section for a more detailed description of each compliance solution and how it can benefit third-party data.

TIP

Click the link in the **Third-party data** column to go the step-by-step instructions for creating a connector for that data type.

THIRD-PARTY DATA	LITIGATION HOLD	EDISCOVERY	RETENTION SETTINGS	RECORDS MANAGEMENT	COMMUNICATION COMPLIANCE	INSIDER RISK MANAGEMENT
Android [1]	✓	✓	✓	✓	✓	
AT&T Network [1]	✓	✓	✓	✓	✓	
Bell Network [1]	✓	✓	✓	✓	✓	
Bloomberg Message	✓	✓	✓	✓	✓	
CellTrust [2]	✓	✓	✓	✓	✓	
Cisco Jabber [2]	✓	✓	✓	✓	✓	
EML [2]	✓	✓	✓	✓		
Enterprise Number [1]	✓	✓	✓	✓	✓	
Facebook	✓	✓	✓	✓		
FX Connect [2]	✓	✓	✓	✓	✓	

THIRD-PARTY DATA	LITIGATION HOLD	EDISCOVERY	RETENTION SETTINGS	RECORDS MANAGEMENT	COMMUNICATION COMPLIANCE	INSIDER RISK MANAGEMENT
Human resources (HR)						✓
ICE Chat	✓	✓	✓	✓	✓	
Instant Bloomberg	✓	✓	✓	✓	✓	
Jive [2]	✓	✓	✓	✓	✓	
LinkedIn	✓	✓	✓	✓		
MS SQL Database [2]	✓	✓	✓	✓		
O2 Network [1]	✓	✓	✓	✓	✓	
Physical badging						✓
Pivot [2]	✓	✓	✓	✓	✓	
Redtail Speak [2]	✓	✓	✓	✓	✓	
Reuters Dealing [2]	✓	✓	✓	✓	✓	
Reuters Eikon [2]	✓	✓	✓	✓	✓	
Reuters FX [2]	✓	✓	✓	✓	✓	
Salesforce Chatter [2]	✓	✓	✓	✓		
ServiceNow [2]	✓	✓	✓	✓		
Slack eDiscovery [2]	✓	✓	✓	✓	✓	
Symphony [2]	✓	✓	✓	✓	✓	
TELUS Network [1]	✓	✓	✓	✓	✓	
Text-delimited [2]	✓	✓	✓	✓		

THIRD-PARTY DATA	LITIGATION HOLD	EDISCOVERY	RETENTION SETTINGS	RECORDS MANAGEMENT	COMMUNICATION COMPLIANCE	INSIDER RISK MANAGEMENT
Twitter	✓	✓	✓	✓		
Verizon Network [1]	✓	✓	✓	✓	✓	
Webex Teams [2]	✓	✓	✓	✓	✓	
Webpages [2]	✓	✓	✓	✓		
WhatsApp [1]	✓	✓	✓	✓	✓	
Workplace from Facebook [2]	✓	✓	✓	✓	✓	
XIP [2]	✓	✓	✓	✓	✓	
XSLT/XML [2]	✓	✓	✓	✓		
Yieldbroker [2]	✓	✓	✓	✓	✓	
Zoom Meetings [2]	✓	✓	✓	✓	✓	

NOTE

[1] Data connector provided by TeleMessage. Before you can archive data in Microsoft 365, you have to work with TeleMessage to set up their archiving service for your organization. For more information, see the prerequisite section in the step-by-step instructions for this data type.

[2] Data connector provided by Globanet. Before you can archive data in Microsoft 365, you have to work with Globanet to set up their archiving service for your organization. For more information, see the prerequisite section in the step-by-step instructions for this data type.

The third-party data listed in the previous table (except for HR data and physical badging data) is imported into user mailboxes. The corresponding compliance solutions that support third-party data are applied to the user mailbox where the data is stored.

Overview of compliance solutions that support third-party data

The following sections describe some of the things that the Microsoft 365 compliance solutions can help you to manage the third-party data listed in the previous table.

Litigation hold

You place a Litigation hold on a user mailbox to retain third-party data. When you create a hold, you can specify a hold duration (also called a *time-based hold*) so that deleted and modified third-party data is retained for a specified period and then permanently deleted from the mailbox. Or you can just retain content indefinitely

(called an *infinite hold*) or until the Litigation hold is removed.

eDiscovery

The three primary eDiscovery tools in Microsoft 365 are Content search, Core eDiscovery, and Advanced eDiscovery.

- **Content search**. You can use the content search tool to search mailboxes for third-party data that you imported. You can use search queries and conditions to narrow your search results, and the export the search results.
- **Core eDiscovery**. This tool builds on the basic search and export functionality by enabling you to create cases that let you control who can access case data, place a hold on user mailboxes or mailbox content that matches search criteria. That means you can place an eDiscovery hold on the third-party data that was imported to user mailboxes.
- **Advanced eDiscovery**. This powerful tool expands the case functionality of Core eDiscovery by letting you add custodians to a case, placing custodian's data on hold, and then loading a custodian's third-party data into a review for further analysis such as themes and duplicate detection. After you load third-party data into a review set, you can query and filter it to a narrow result set.

  Both Core eDiscovery and Advanced eDiscovery let you manage third-party data that may be relevant to your organization's legal or internal investigations.

Retention settings

You can apply a retention policy to user mailboxes to retain and then delete third-party data (and other mailbox content) after retention period expires. You can also use retention policies to delete third-party data of a certain age or use retention labels to trigger a disposition review when the retention period for third-party data expires.

Records management

The records management feature in Microsoft 365 lets you declare third-party data as a record. This can be done manually by users who apply a retention label that marks third-party data in their mailbox as record. Or you can auto-apply retention labels by identifying sensitive information, keywords, or content types in third-party data.

Communication compliance

You can use Communication compliance to examine third-party data to make sure it is compliant with your organization's data standards. You can do this by you detecting, capturing, and taking remediation actions for inappropriate messages in your organization. For example, you can monitor the third-party data that you import for offensive language, sensitive information, and regulatory compliance.

Insider risk management

Signals from third-party data, like selective HR data, can be used by the Insider risk management solution to minimize internal risks by letting you to detect, investigate, and act on risky activities in your organization. For example, data imported by the HR data connector is used as risk indicators to help detect departing employee data theft.

Working with a Microsoft partner to archive third-party data

Another option for importing and archiving third-party data is for your organization to work with a Microsoft Partner. If a third-party data type isn't supported by the data connectors available in the Microsoft compliance center, you can work with a partner who can provide a custom connector that will be configured to extract items from the third-party data source on a regular basis and then connect to the Microsoft cloud by a third-party API and import those items to Microsoft 365. The partner connector also converts the content of an item from the third-party data source to an email message and then imports it to a mailbox in Microsoft 365.

For a list of partners that you can work with and the step-by-step process for this method, see Work with a

partner to archive third-party data in Microsoft 365.

Set up a connector to archive Bloomberg Message data

2/18/2021 • 8 minutes to read • Edit Online

Use a data connector in the Microsoft 365 compliance center to import and archive financial services email data from the Bloomberg Message collaboration tool. After you set up and configure a connector, it connects to your organization's Bloomberg secure FTP (SFTP) site once every day, and imports email items to mailboxes in Microsoft 365.

After Bloomberg Message data is stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation hold, content search, In-place archiving, auditing, Communication compliance, and Microsoft 365 retention policies to Bloomberg Message data. For example, you can search Bloomberg Message emails using the content search tool or associate the mailbox that contains the Bloomberg Message data with a custodian in an Advanced eDiscovery case. Using a Bloomberg Message connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving Bloomberg Message data

The following overview explains the process of using a connector to archive Bloomberg Message data in Microsoft 365.

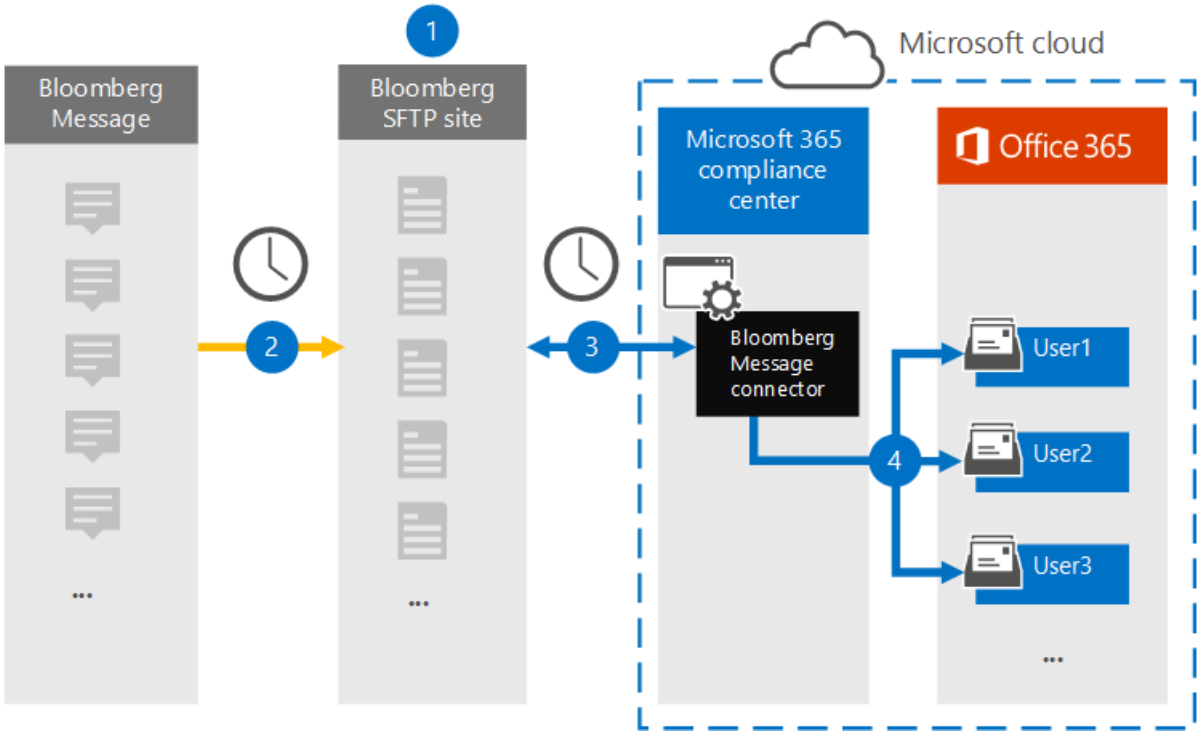

1. Your organization works with Bloomberg to set up a Bloomberg SFTP site. You'll also work with Bloomberg to configure Bloomberg Message to copy email messages to the Bloomberg SFTP site.
2. Once every 24 hours, email messages from Bloomberg Message are copied to the Bloomberg SFTP site.
3. The Bloomberg Message connector that you create in the Microsoft 365 compliance center connects to the Bloomberg SFTP site every day and transfers the email messages from the previous 24 hours to a secure Azure Storage area in the Microsoft Cloud.
4. The connector imports the email message items to the mailbox of a specific user. A new folder named

BloombergMessage is created in the specific user's mailbox and the items will be imported to it.

The connector does this by using the value of the CorporateEmailAddress property. Every email message contains this property, which is populated with the email address of every participant of the email message. In addition to automatic user mapping using the value of the *CorporateEmailAddress* property, you can also define a custom mapping by uploading a CSV mapping file. This mapping file contains a Bloomberg UUID and the corresponding Microsoft 365 mailbox address for each user in your organization. If you enable automatic user mapping and provide a custom mapping, for every email item the connector will first look at the custom-mapping file. If it doesn't find a valid Microsoft 365 user that corresponds to a user's Bloomberg UUID, the connector uses the *CorporateEmailAddress* property of the email item. If the connector doesn't find a valid Microsoft 365 user in either the custom-mapping file or the *CorporateEmailAddress* property of the email item, the item won't be imported.

Before you begin

Some of the implementation steps required to archive Bloomberg Message data are external to Microsoft 365 and must be completed before you can create the connector in the compliance center.

- Subscribe to Bloomberg Anywhere. This is required so that you can log in to Bloomberg Anywhere to access the Bloomberg SFTP site that you have to set up and configure.
- Set up a Bloomberg SFTP (Secure file transfer protocol) site. After working with Bloomberg to set up the SFTP site, data from Bloomberg Message is uploaded to the SFTP site every day. The connector you create in Step 2 connects to this SFTP site and transfers the email data to Microsoft 365 mailboxes. SFTP also encrypts the Bloomberg Message data that is sent to mailboxes during the transfer process.

  For information about Bloomberg SFTP (also called *BB-SFTP*):

  - See the "SFTP Connectivity Standards" document at Bloomberg Support.
  - Contact Bloomberg customer support.

  > **NOTE**
  > If your organization already deployed a connector to archive Instant Bloomberg data, you don't need to set up another SFTP site. You can use the same SFTP site for the Bloomberg Message connector.

- After you work with Bloomberg to set up an SFTP site, Bloomberg will provide some information to you after you respond to the Bloomberg implementation email message. Save a copy of the following information. You use it to set up a connector in Step 3.
  - Firm code, which is an ID for your organization and is used to log in to the Bloomberg SFTP site.
  - Password for your Bloomberg SFTP site
  - URL for Bloomberg SFTP site (for example, sftp.bloomberg.com). In addition, Bloomberg may also provide a corresponding IP address for the Bloomberg SFTP site, which also can be used to set up the connector.
  - Port number for Bloomberg SFTP site
- The Bloomberg Message connector can import a total of 200,000 items in a single day. If there are more than 200,000 items on the SFTP site, none of those items will be imported to Microsoft 365.
- The user who creates a Bloomberg Message connector in Step 3 (and who downloads the public keys and IP address in Step 1) must be assigned the Mailbox Import Export role in Exchange Online. This is required to add connectors in the **Data connectors** page in the Microsoft 365 compliance center. By

default, this role isn't assigned to any role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Step 1: Obtain SSH and PGP public keys

The first step is to obtain a copy of the public keys for Secure Shell (SSH) and Pretty Good Privacy (PGP). You use these keys in Step 2 to configure the Bloomberg SFTP site to allow the connector (that you create in Step 3) to connect to the SFTP site and transfer the Bloomberg Message email data to Microsoft 365 mailboxes. You also obtain an IP address in this step, which you use when configuring the Bloomberg SFTP site.

1. Go to https://compliance.microsoft.com and click **Data connectors** in the left nav.
2. On the **Data connectors** page under **Bloomberg Message**, click **View**.
3. On the **Bloomberg Message** product description page, click **Add connector**
4. On the **Terms of service** page, click **Accept**.
5. On the **Add credentials for Bloomberg SFTP site** under step 1, click the **Download SSH key**, **Download PGP key**, and **Download IP address** links to save a copy of each file to your local computer. These files contain the following items that are used to configure the Bloomberg SFTP site in Step 2:
   - SSH public key: This key is used to configure Secure Shell (SSH) to enable a secure remote login when the connector connects to the Bloomberg SFTP site.
   - PGP public key: This key is used to configure the encryption of data that's transferred from the Bloomberg SFTP site to Microsoft 365.
   - IP address: The Bloomberg SFTP site is configured to accept a connection request only from this IP address, which is used by the Bloomberg Message connector that you create in Step 3.
6. Click **Cancel** to close the wizard. You come back to this wizard in Step 3 to create the connector.

Step 2: Configure the Bloomberg SFTP site

> **NOTE**
> As previously stated, if you're organization has previously set up a Bloomberg SFTP site to archive Instant Bloomberg data, you don't have to set up another one. You can specify the same SFTP site when you create the connector in Step 3.

The next step is to use the SSH and PGP public keys and the IP address that you obtained in Step 1 to configure SSH authentication and PGP encryption for the Bloomberg SFTP site. This lets the Bloomberg Message connector that you create in Step 3 connect to the Bloomberg SFTP site and transfer Bloomberg Message data to Microsoft 365. You need to work with Bloomberg customer support to set up your Bloomberg SFTP site. Contact Bloomberg customer support for assistance.

> **IMPORTANT**
> Bloomberg recommends that you attach the three files that you downloaded in Step 1 to an email message and send it to their customer support team when working with them to set up your Bloomberg SFTP site.

Step 3: Create a Bloomberg Message connector

The last step is to create a Bloomberg Message connector in the Microsoft 365 compliance center. The connector uses the information you provide to connect to the Bloomberg SFTP site and transfer email messages to the corresponding user mailbox boxes in Microsoft 365.

1. Go to https://compliance.microsoft.com and click **Data connectors** in the left nav.
2. On the **Data connectors** page under **Bloomberg Message**, click **View**.
3. On the **Bloomberg Message** product description page, click **Add connector**
4. On the **Terms of service** page, click **Accept**.
5. On the **Add credentials for Bloomberg SFTP site** page, under Step 3, enter the required information in the following boxes and then click **Next**.
   - **Firm code:** The ID for your organization that is used as the username for the Bloomberg SFTP site.
   - **Password:** The password for your organization's Bloomberg SFTP site.
   - **SFTP URL:** The URL for the Bloomberg SFTP site (for example, sftp.bloomberg.com).
   - **SFTP port:** The port number for the Bloomberg SFTP site. The connector uses this port to connect to the SFTP site.
6. On the **User-mapping** page, enable automatic user mapping and provide custom user mapping as required
7. Click **Next**, review your settings, and then click prepare to create the connector.
8. Go to the **Data connectors** page to see the progress of the import process for the new connector.

Known issues

- Threading of Bloomberg Message email imported to Microsoft 365 isn't supported. Individual messages sent to a person are imported, but they aren't presented in a threaded conversation. Microsoft is working to support threading in later versions of the Bloomberg Message data connector.

Set up a connector to archive Facebook data (preview)

2/18/2021 • 5 minutes to read • Edit Online

Use a connector in the Microsoft 365 compliance center to import and archive data from Facebook Business pages to Microsoft 365. After you set up and configure the connector, it connects to the Facebook Business page (on a scheduled basis), converts the content of Facebook items to an email message format, and then imports those items to a mailbox in Microsoft 365.

After the Facebook data is imported, you can apply Microsoft 365 compliance features such as Litigation Hold, Content Search, In-Place Archiving, Auditing, Communication compliance, and Microsoft 365 retention policies to the Facebook data. For example, when a mailbox is placed on Litigation Hold or assigned to a retention policy, the Facebook data is preserved. You can search third-party data using Content Search or associate the mailbox where the Facebook data is stored with a custodian in an Advanced eDiscovery case. Using a connector to import and archive Facebook data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Prerequisites for setting up a connector for Facebook Business pages

Complete the following prerequisites before you can set up and configure a connector in the Microsoft 365 compliance center to import and archive data from your organization's Facebook Business pages.

- You need a Facebook account for your organization's business pages (you need to sign in to this account when setting up the connector). Currently, you can only archive data from Facebook Business pages; you can't archive data from individual Facebook profiles.
- Your organization must have a valid Azure subscription. If you don't have an existing Azure subscription, you can sign up for one of these options:
  - Sign up for a free one year Azure subscription
  - Sign up for a Pay-As-You-Go Azure subscription

  > **NOTE**
  > The free Azure Active Directory subscription that's included with your Microsoft 365 subscription doesn't support the connectors in the Security & Compliance Center.

- The connector for Facebook Business pages can import a total of 200,000 items in a single day. If there are more than 200,000 Facebook Business items in a day, none of those items will be imported to Microsoft 365.
- The user who sets up the custom connector in the Microsoft 365 compliance center (in Step 5) must be assigned the Mailbox Import Export role in Exchange Online. By default, this role isn't assigned to any role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Step 1: Create an app in Azure Active Directory

The first step is to register a new app in Azure Active Directory (AAD). This app corresponds to the web app resource that you implement in Step 4 and Step 5 for the Facebook connector.

For step-by-step instructions, see Create an app in Azure Active Directory.

During the completion of this step (by using the previous step-by-step instructions), you'll save the following information to a text file. These values are used in later steps in the deployment process.

- AAD application ID
- AAD application secret
- Tenant Id

Step 2: Deploy the connector web service from GitHub to your Azure account

The next step is to deploy the source code for the Facebook Business pages connector app that will use the Facebook API to connect to your Facebook account and extract data so you can import it to Microsoft 365. The Facebook connector that you deploy for your organization will upload the items from your Facebook Business pages to the Azure Storage location that is created in this step. After you create a Facebook business pages connector in the Microsoft 365 compliance center (in Step 5), the Import service will copy the Facebook business pages data from the Azure Storage location to a mailbox in your Microsoft 365 organization. As previous explained in the Prerequisites section, you must have a valid Azure subscription to create an Azure Storage account.

For step-by-step instructions, see Deploy the connector web service from GitHub to your Azure account.

In the step-by-step instructions to complete this step, you'll provide the following information:

- APISecretKey: You create this secret during the completion of this step. It's used in Step 5.
- TenantId: The tenant ID of your Microsoft 365 organization that you copied after creating the Facebook connector app in Azure Active Directory in Step 1.

After completing this step, be sure to copy the Azure app service URL (for example, https://fbconnector.azurewebsites.net). You need to use this URL to complete Step 3, Step 4, and Step 5).

Step 3: Register the web app on Facebook

The next step is to create and configure a new app on Facebook. The Facebook business pages connector that you create in Step 5 uses the Facebook web app to interact with the Facebook API to obtain data from your organization's Facebook Business pages.

For step-by-step instructions, see Register the Facebook app.

During the completion of this step (by following the step-by-step instructions), you save the following information to a text file. These values are used to configure the Facebook connector app in Step 4.

- Facebook application ID
- Facebook application secret
- Facebook webhooks verify token

Step 4: Configure the Facebook connector app

The next step is to add configuration settings to the Facebook connector app that you uploaded when you created the Azure web app resource in Step 1. You do this by going to the home page of your connector app and

configuring it.

For step-by-step instructions, see Configure the Facebook connector app.

During the completion of this step (by following the step-by-step instructions), you provide the following information (that you've copied to a text file after completing the previous steps):

- Facebook application ID (obtained in Step 3)
- Facebook application secret (obtained in Step 3)
- Facebook webhooks verify token (obtained in Step 3)
- Azure Active Directory application ID (the AAD application ID obtained in Step 1)
- Azure Active Directory application secret (the AAD application secret obtained in Step 1)

Step 5: Set up a Facebook Business pages connector in the Microsoft 365 compliance center

The final step is to set up the connector in the Microsoft 365 compliance center that will import data from your Facebook Business pages to a specified mailbox in Microsoft 365. After you complete this step, the Microsoft 365 Import service will start importing data from your Facebook Business pages to Microsoft 365.

For step-by-step instructions, see Step 5: Set up a Facebook connector in the Microsoft 365 compliance center.

During the completion of this step (by following the step-by-step instructions), you provide the following information (that you've copied to a text file after completing the steps).

- AAD application ID (obtained in Step 1)
- Azure app service URL (obtained in Step 1; for example, https://fbconnector.azurewebsites.net)
- APISecretKey (that you created in Step 2)

Set up a connector to import HR data

2/18/2021 • 24 minutes to read • Edit Online

You can set up a data connector in the Microsoft 365 compliance center to import human resources (HR) data related to events such as a user's resignation or a change in a user's job level. The HR data can then be used by the insider risk management solution to generate risk indicators that can help you identity possible malicious activity or data theft by users inside your organization.

Setting up a connector for HR data that insider risk management policies can use to generate risk indicators consists of creating a CSV file that contains that contains the HR data, creating an app in Azure Active Directory that's used for authentication, creating an HR data connector in the Microsoft 365 compliance center, and then running a script (on a scheduled basis) that ingests the HR data in CSV files to the Microsoft cloud so it's available to the insider risk management solution.

Before you begin

- Determine which HR scenarios and data to import to Microsoft 365. This will help you determine how many CSV files and HR connectors you'll need to create, and how to generate and structure the CSV files. The HR data that you import is determined by the insider risk management policies that you want to implement. For more information, see Step 1.
- Determine how to retrieve or export the data from your organization's HR system (and on a regular basis) and add it to the CSV files that you create in Step 1. The script that you run in Step 4 will upload the HR data in the CSV files to the Microsoft cloud.
- The user who creates the HR connector in Step 3 must be assigned the Mailbox Import Export role in Exchange Online. By default, this role isn't assigned to any role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a new role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".
- The sample script that you run in Step 4 will upload your HR data to the Microsoft cloud so that it can be used by the insider risk management solution. This sample script isn't supported under any Microsoft standard support program or service. The sample script is provided AS IS without warranty of any kind. Microsoft further disclaims all implied warranties including, without limitation, any implied warranties of merchantability or of fitness for a particular purpose. The entire risk arising out of the use or performance of the sample script and documentation remains with you. In no event shall Microsoft, its authors, or anyone else involved in the creation, production, or delivery of the scripts be liable for any damages whatsoever (including, without limitation, damages for loss of business profits, business interruption, loss of business information, or other pecuniary loss) arising out of the use of or inability to use the sample scripts or documentation, even if Microsoft has been advised of the possibility of such damages.

Step 1: Prepare a CSV file with your HR data

The first step is to create a CSV file that contains the HR data that the connector will import to Microsoft 365. This data will be used by the insider risk solution to generate potential risk indicators. Data for the following HR scenarios can be imported to Microsoft 365:

- Employee resignation. Information about users who have left your organization.

- Job level changes. Information about job level changes for users, such as promotions and demotions.
- Performance reviews. Information about user performance.
- Performance improvement plans. Information about performance improvement plans for users.

The type of HR data to import depends on the insider risk management policy and corresponding policy template that you want to implement. The following table shows which HR data type is required for each policy template:

POLICY TEMPLATE	HR DATA TYPE
Data theft by departing users	Employee resignations
General data leaks	Not applicable
Data leaks by priority users	Not applicable
Data leaks by disgruntled users	Job level changes, Performance reviews, Performance improvement plans
General security policy violations	Not applicable
Security policy violations by departing users	Employee resignations
Security policy violations by priority users	Not applicable
Security policy violations by disgruntled users	Job level changes, Performance reviews, Performance improvement plans
Offensive language in email	Not applicable

For more information about policy templates for insider risk management, see Insider risk management policies.

For each HR scenario, you'll need to provide the corresponding HR data in one or more CSV files. The number of CSV files to use for your insider risk management implementation is discussed later in this section.

After you create the CSV file with the required HR data, store it on the local computer that you run the script on in Step 4. You should also implement an update strategy to make sure the CSV file always contains the most current information so that whatever you run the script, the most current HR data will be uploaded to the Microsoft cloud and accessible to the insider risk management solution.

> **IMPORTANT**
> The column names described in the following sections are not required parameters, but only examples. You can use any column name in your CSV files. However, the column names you use in a CSV file *must* be mapped to the data type when you create the HR connector in Step 3. Also note that the sample CSV files in the following sections are show in NotePad view. It's much easier to view and edit CSV files in Microsoft Excel.

The follow sections describe the required CSV data for each HR scenario.

CSV file for employee resignation data

Here's an example of a CSV file for employee resignation data.

```
EmailAddress,ResignationDate,LastWorkingDate
sarad@contoso.com,2019-04-23T15:18:02.4675041+05:30,2019-04-29T15:18:02.4675041+05:30
pilarp@contoso.com,2019-04-24T09:15:49Z,2019-04-29T15:18:02.7117540
```

The following table describes each column in the CSV file for employee resignation data.

COLUMN	DESCRIPTION
EmailAddress	Specifies the email address (UPN) of the terminated user.
ResignationDate	Specifies the date the user's employment was officially terminated in your organization. For example, this may be the date when the user gave their notice about leaving your organization. This date may be the different than the date of the person's last day of work. Use the following date format: `yyyy-mm-ddThh:mm:ss.nnnnnn+|-hh:mm`, which is the ISO 8601 date and time format.
LastWorkingDate	Specifies the last day of work for the terminated user. Use the following date format: `yyyy-mm-ddThh:mm:ss.nnnnnn+|-hh:mm`, which is the ISO 8601 date and time format.

CSV file for job level changes data

Here's an example of a CSV file for job level changes data.

```
EmailAddress,EffectiveDate,OldLevel,NewLevel
sarad@contoso.com,2019-04-23T15:18:02.4675041+05:30,Level 61 – Sr. Manager,Level 60- Manager
pillar@contoso.com,2019-04-23T15:18:02.4675041+05:30,Level 62 – Director,Level 60- Sr. Manager
```

The following table describes each column in the CSV file for job level changes data.

COLUMN	DESCRIPTION
EmailAddress	Specifies the user's email address (UPN).
EffectiveDate	Specifies the date that the user's job level was officially changed. Use the following date format: `yyyy-mm-ddThh:mm:ss.nnnnnn+|-hh:mm`, which is the ISO 8601 date and time format.
Remarks	Specifies the remarks that evaluator has provided about the change of job level. You can enter a limit of 200 characters. This parameter is optional. You don't have to include it in the CSV file.
OldLevel	Specifies the user's job level before it was changed. This is a free-text parameter and can contain hierarchical taxonomy for your organization. This parameter is optional. You don't have to include it in the CSV file.

COLUMN	DESCRIPTION
NewLevel	Specifies the user's job level after it was changed. This is a free-text parameter and can contain hierarchical taxonomy for your organization. This parameter is optional. You don't have to include it in the CSV file.

CSV file for performance review data

Here's an example of a CSV file for performance data.

```
EmailAddress,EffectiveDate,Remarks,Rating
sarad@contoso.com,2019-04-23T15:18:02.4675041+05:30,Met expectations but bad attitude,2-Below expectation
pillar@contoso.com,2019-04-23T15:18:02.4675041+05:30, Multiple conflicts with the team
```

The following table describes each column in the CSV file for performance review data.

COLUMN	DESCRIPTION
EmailAddress	Specifies the user's email address (UPN).
EffectiveDate	Specifies the date that the user was officially informed about the result of their performance review. This can be the date when the performance review cycle ended. Use the following date format: `yyyy-mm-ddThh:mm:ss.nnnnnn+|-hh:mm`, which is the ISO 8601 date and time format.
Remarks	Specifies any remarks that evaluator has provided to the user for the performance review. This is a text parameter with a limit of 200 characters. This parameter is optional. You don't have to include it in the CSV file.
Rating	Specifies the rating provided for the performance review. This is a text parameter and can contain any free-form text that your organization uses to recognize the evaluation. For example, "3 Met expectations" or "2 Below average". This is a text parameter with a limit of 25 characters. This parameter is optional. You don't have to include it in the CSV file.

CSV file for performance improvement plan data

Here's an example of a CSV file for the data for the performance improvement plan data.

```
EmailAddress,EffectiveDate,ImprovementRemarks,PerformanceRating
sarad@contoso.com,2019-04-23T15:18:02.4675041+05:30,Met expectation but bad attitude,2-Below expectation
pillar@contoso.com,2019-04-23T15:18:02.4675041+05:30, Multiple conflicts with the team
```

The following table describes each column in the CSV file for performance review data.

COLUMN	DESCRIPTION
EmailAddress	Specifies the user's email address (UPN).

COLUMN	DESCRIPTION
EffectiveDate	Specifies the date when the user was officially informed about their performance improvement plan. You must use the following date format: `yyyy-mm-ddThh:mm:ss.nnnnnn+|-hh:mm`, which is the ISO 8601 date and time format.
Remarks	Specifies any remarks that evaluator has provided about the performance improvement plan. This is a text parameter with a limit of 200 characters. This is an optional parameter. You don't have to include it in the CSV file.
Rating	Specifies any rating or other information related to the performance review. performance improvement plan. This is a text parameter and can contain any free form text that your organization uses to recognize the evaluation. For example, "3 Met expectations" or "2 Below average". This is a text parameter with limit of 25 characters. This is an optional parameter. You don't have to include it in the CSV file.

Determining how many CSV files to use for HR data

In Step 3, you can choose to create separate connectors for each HR data type or you can choose to create single connector for all data types. You can use separate CSV files that contain data for one HR scenario (like the examples of the CSV files described in the previous sections). Alternatively, you can use a single CSV file that contains data for two or more HR scenarios. Here are some guidelines to help you determine how many CSV files to use for HR data.

- If the insider risk management policy that you want to implement requires multiple HR data types, consider using a single CSV file that contains all the required data types.
- The method for generating or collecting the HR data may determine the number of CSV files. For example, if the different types of HR data used to configure an HR connector are located in a single HR system in your organization, then you may be able to export the data to a single CSV file. But if data is distributed across different HR systems, then it might be easier to export data to different CSV files. For example, Employee resignation data may be located in a different HR system than Job level or Performance review data. In this case, it may be easier to have separate CSV files rather than having to manually combine the data into a single CSV file. So, how you retrieve or export data from your HR systems may determine how the number of CSV files you'll need.
- As a general rule, the number of HR connectors that you'll need to create is determined by the data types in a CSV file. For example, if a CSV file contains all the data types required to support your insider risk management implementation, then you only need one HR connector. But if you have two separate CSV files that each contain a single data type, then you'll have to create two HR connectors. An exception to this is that if you add an **HRScenario** column to a CSV file (see the next section), you can configure a single HR connector that can process different CSV files.

Configuring a single CSV file for multiple HR data types

You can add multiple HR data types to a single CSV file. This is useful if the insider risk management solution you're implementing requires multiple HR data types or if the data types are located in a single HR system in your organization. Having fewer CSV files always allows you to have fewer HR connectors to create and manage.

Here are requirements for configuring a CSV file with multiple data types:

- You have to add the required columns (and optional if you use them) for each data type and the corresponding column name in the header row. If a data type doesn't correspond to a column, you can leave the value blank.
- To use a CSV file with multiple types of HR data, the HR connector needs to know which rows in the CSV file contain which type HR data. This is accomplished by adding an additional **HRScenario** column to the CSV file. The values in this column identify the type of HR data in each row. For example, values that correspond to the four HR scenarios could be `Resignation`, `Job level change`, `Performance review`, and `Performance improvement plan`.
- If you have multiple CSV files that contain an HRScenario** column, be sure that each file uses the same column name and the same values that identify the specific HR scenarios.

The following example shows a CSV file that contains the **HRScenario** column. The values in the HRScenario column identify the type of data in the corresponding row.

```
HRScenario,EmailAddress,ResignationDate,LastWorkingDate,EffectiveDate,Remarks,Rating,OldLevel,NewLevel
Resignation,sarad@contoso.com,2019-04-23T15:18:02.4675041+05:30,2019-04-29T15:18:02.4675041+05:30,,,,
Resignation,pilarp@contoso.com,2019-04-24T09:15:49Z,2019-04-29T15:18:02.7117540,,,,
Job level change,sarad@contoso.com,2019-04-23T15:18:02.4675041+05:30,,,,,Level 61 Sr. Manager, Level 60
Manager
Job level change,pillarp@contoso.com,2019-04-23T15:18:02.4675041+05:30,,,,,Level 62 Director,Level 60 Sr
Manager
Performance review,sarad@contoso.com,,,2019-04-23T15:18:02.4675041+05:30,Met expectation but bad attitude,2
Below expectations,,
Performance review,pillarp@contoso.com,,,2019-04-23T15:18:02.4675041+05:30, Multiple conflicts with the
team,,
Performance improvement plan,sarad@contoso.com,,,2019-04-23T15:18:02.4675041+05:30,Met expectations but bad
attitude,2 Below expectations,,
Performance improvement plan,pillarp@contoso.com,,,2019-04-23T15:18:02.4675041+05:30,Multiple conflicts with
the team,,
```

> **NOTE**
>
> You can use any name for the column that identifies HR data type because you will map the name of the column in your CSV file as the column that identifies the HR data type when you set up the connector in Step 3. You will also map the values used for the data type column when you set up the connector.

Adding the HRScenario column to a CSV file that contains a single data type

Based on your organization's HR systems and how you will export HR data to CSV file, you might have to create multiple CSV files that contain a single HR data type. In this case, you can still create a single HR connector to import data from different CSV files. To do this, you'll just have to add an HRScenario column to the CSV file and specify the HR data type. Then you can run the script for each CSV file, but use the same job ID for the connector. See Step 4.

Step 2: Create an app in Azure Active Directory

The next step is to create and register a new app in Azure Active Directory (Azure AD). The app will correspond to the HR connector that you create in Step 3. Creating this app will allow Azure AD to authenticate the HR connector when it runs and attempts to access your organization. This app will also be used to authenticate the script that you run in Step 4 to upload your HR data to the Microsoft cloud. During the creation of this Azure AD app, be sure to save the following information. These values will be used in Step 3 and Step 4.

- Azure AD application ID (also called the *app Id* or *client Id*)
- Azure AD application secret (also called the *client secret*)

- Tenant Id (also called the *directory Id*)

For step-by-step instructions for creating an app in Azure AD, see Register an application with the Microsoft identity platform.

Step 3: Create the HR connector

The next step is to create an HR connector in the Microsoft 365 compliance center. After you run the script in Step 4, the HR connector that you create will ingest the HR data from the CSV file to your Microsoft 365 organization. Before you create a connector, be sure that you have a list of the HR scenarios and the corresponding CSV column names for each one. You have to map the data required for each scenario to the actual column names in your CSV file when configuring the connector. Alternatively, you can upload a sample CSV file when configuring the connector and the wizard will help you map the name of the columns to the required data types.

After you complete this step, be sure to copy the job ID that's generated when you create the connector. You'll use the job ID when you run the script.

1. Go to https://compliance.microsoft.com and then click **Data connectors** in the left nav.
2. On the **Data connectors** page under **HR**, click **View**.
3. On the **HR Custom** page, click **Add connector**.
4. On the **Setup the connection** page, do the following and then click **Next**:
   a. Type or paste the Azure AD application ID for the Azure app that you created in Step 2.
   b. Type a name for the HR connector.
5. On the HR scenarios page, select one or more HR scenarios that you want to import data for, and then click **Next**.
6. On the file mapping method page, select one of the following options and then click **Next**.
   - **Upload a sample file**. If you select this option, click **Upload sample file** to upload the CSV file that you prepared in Step 1. This option allows you to quickly select column names in your CSV file from a drop-down list to map them to the data types for the HR scenarios that you previously selected.

   OR

   - **Manually provide the mapping details**. If you select this option, you have to type the name of the columns in your CSV file to map them to the data types for the HR scenarios that you previously selected.
7. On the File mapping details page, do one of the following, depending on whether you uploaded a sample CSV file and whether you're configuring the connector for a single HR scenario or for multiple scenarios. If you uploaded a sample file, you don't have to type the column names. You pick them from a dropdown list.
   - If you selected a single HR scenario in the previous step, then type the column header names (also called *parameters*) from the CSV file that you created in Step 1 in each of the appropriate boxes. The column names that you type are not case-sensitive, but be sure to include spaces if the column names in your CSV file include spaces. As previously explained, the names you type in these boxes must match the parameter names in your CSV file. For example, the following screenshot shows the parameter names from the sample CSV file for the employee resignation HR scenario shown in Step 1.
   - If you selected multiple data types in step above, then you need to enter identifier column name that will identify the HR data type in your CSV file. After entering the identifier column name, type

the value that identifies this HR data type, and type the column header names for selected data types from the CSV file(s) that you created in Step 1 in each of the appropriate boxes for each selected data type. As previously explained, the names that you type in these boxes must match the column names in your CSV file.

8. On the **Review** page, review your settings and then click **Finish** to create the connector.

   A status page is displayed that confirms the connector was created. This page contains two important things that you need to complete the next step to run the sample script to upload your HR data.

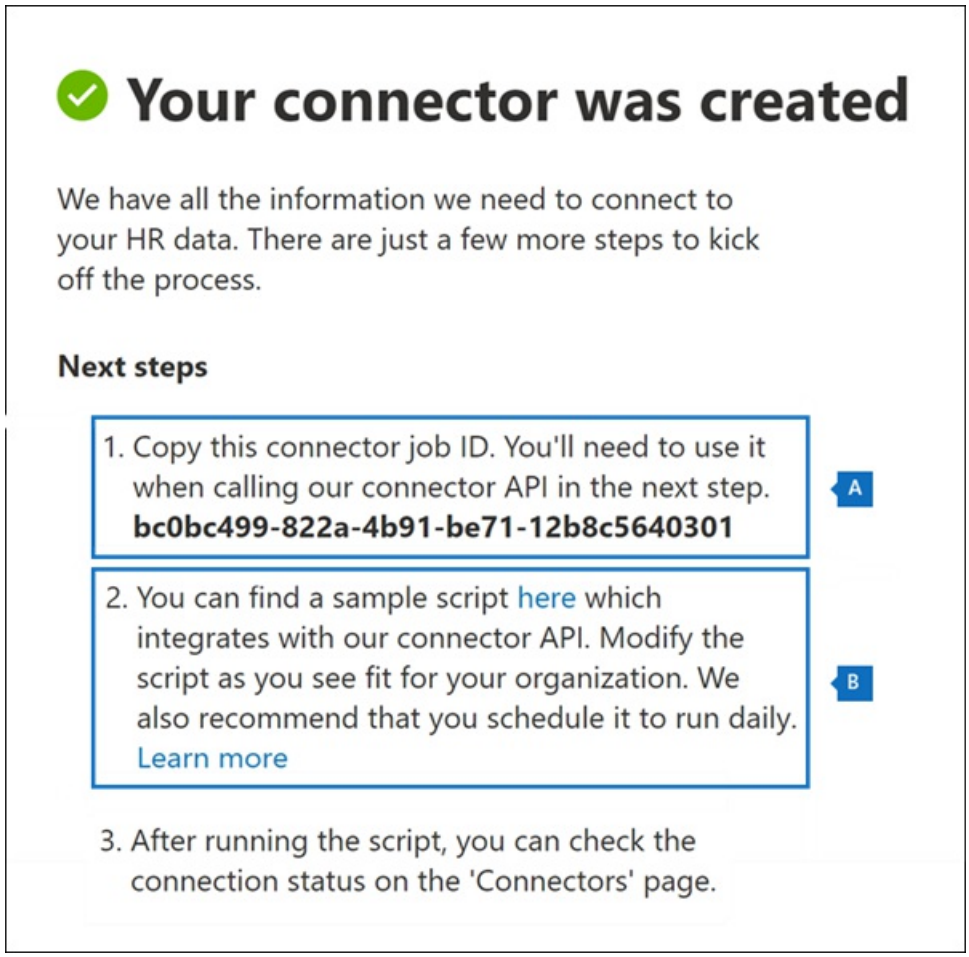

   a. **Job ID.** You'll need this job ID to run the script in the next step. You can copy it from this page or from the connector flyout page.

   b. **Link to sample script.** Click the **here** link to go to the GitHub site to access the sample script (the link opens a new window). Keep this window open so that you can copy the script in Step 4. Alternatively, you can bookmark the destination or copy the URL so you can access it again when you run the script. This link is also available on the connector flyout page.

9. Click **Done**.

   The new connector is displayed in the list on the **Connectors** tab.

10. Click the HR connector that you just created to display the flyout page, which contains properties and other information about the connector.

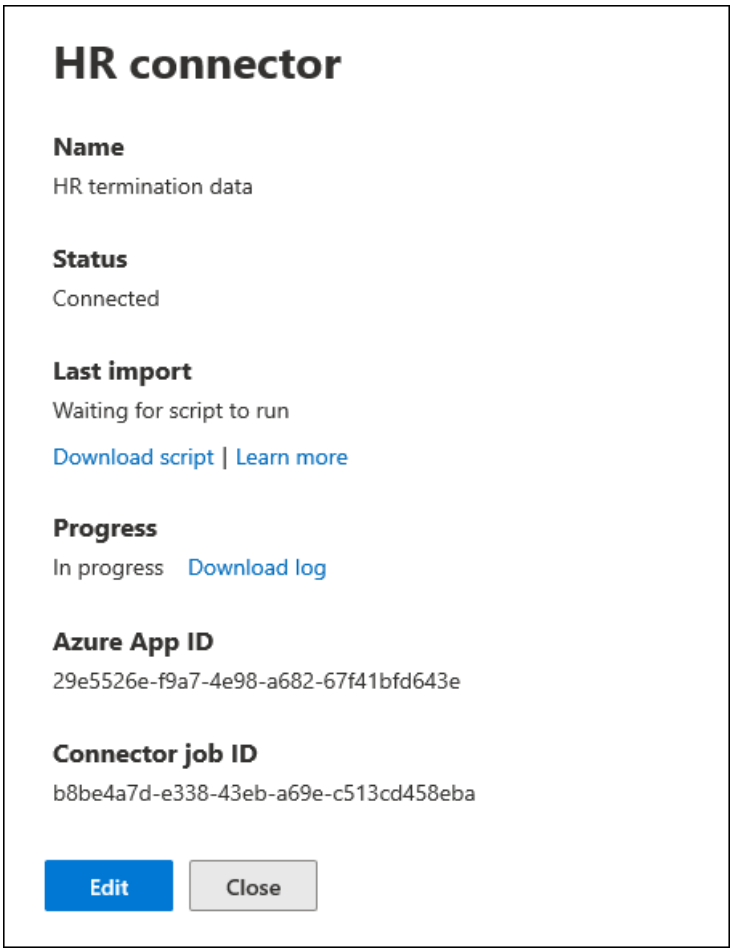

If you haven't already done so, you can copy the values for the **Azure App ID** and **Connector job ID**. You'll need these to run the script in the next step. You can also download the script from the flyout page (or download it using the link in the next step.)

You can also click **Edit** to change the Azure App ID or the column header names that you defined on the **File mapping** page.

Step 4: Run the sample script to upload your HR data

The last step in setting up an HR connector is to run a sample script that will upload the HR data in the CSV file (that you created in Step 1) to the Microsoft cloud. Specifically, the script uploads the data to the HR connector. After you run the script, the HR connector that you created in Step 3 imports the HR data to your Microsoft 365 organization where it can accessed by other compliance tools, such as the Insider risk management solution. After you run the script, consider scheduling a task to run it automatically on a daily basis so the most current employee termination data is uploaded to the Microsoft cloud. See Schedule the script to run automatically.

1. Go to window that you left open from the previous step to access the GitHub site with the sample script. Alternatively, open the bookmarked site or use the URL that you copied.
2. Click the **Raw** button to display the script in text view.
3. Copy all the lines in the sample script and then save them to a text file.
4. Modify the sample script for your organization, if necessary.
5. Save the text file as a Windows PowerShell script file by using a filename suffix of `.ps1`; for example, `HRConnector.ps1`.
6. Open a Command Prompt on your local computer, and go to the directory where you saved the script.
7. Run the following command to upload the HR data in the CSV file to the Microsoft cloud; for example:

```
.\HRConnector.ps1 -tenantId <tenantId> -appId <appId>  -appSecret <appSecret>  -jobId <jobId>  -
csvFilePath '<csvFilePath>'
```

The following table describes the parameters to use with this script and their required values. The information you obtained in the previous steps is used in the values for these parameters.

PARAMETER	DESCRIPTION
`tenantId`	This is the Id for your Microsoft 365 organization that you obtained in Step 2. You can also obtain the tenant Id for your organization on the **Overview** blade in the Azure AD admin center. This is used to identify your organization.
`appId`	This is the Azure AD application Id for the app that you created in Azure AD in Step 2. This is used by Azure AD for authentication when the script attempts to accesses your Microsoft 365 organization.
`appSecret`	This is the Azure AD application secret for the app that you created in Azure AD in Step 2. This also used for authentication.
`jobId`	This is the job ID for the HR connector that you created in Step 3. This is used to associate the HR data that is uploaded to the Microsoft cloud with the HR connector.
`csvFilePath`	This is the file path for the CSV file (stored on the same system as the script) that you created in Step 1. Try to avoid spaces in the file path; otherwise use single quotation marks.

Here's an example of the syntax for the HR connector script using actual values for each parameter:

```
 .\HRConnector.ps1 -tenantId d5723623-11cf-4e2e-b5a5-01d1506273g9 -appId 29ee526e-f9a7-4e98-a682-
67f41bfd643e -appSecret MNubVGbcQDkGCnn -jobId b8be4a7d-e338-43eb-a69e-c513cd458eba -csvFilePath
'C:\Users\contosoadmin\Desktop\Data\employee_termination_data.csv'
```

If the upload is successful, the script displays the **Upload Successful** message.

> **NOTE**
> If you have problems running the previous command because of execution policies, see About Execution Policies and Set-ExecutionPolicy for guidance about setting execution policies.

Step 5: Monitor the HR connector

After you create the HR connector and run the script to upload your HR data, you can view the connector and upload status in the Microsoft 365 compliance center. If you schedule the script to run automatically on a regular basis, you can also view the current status after the last time the script ran.

1. Go to https://compliance.microsoft.com and click **Data connectors** in the left nav.
2. Click the **Connectors** tab and then select the HR connector to display the flyout page. This page contains

the properties and information about the connector.

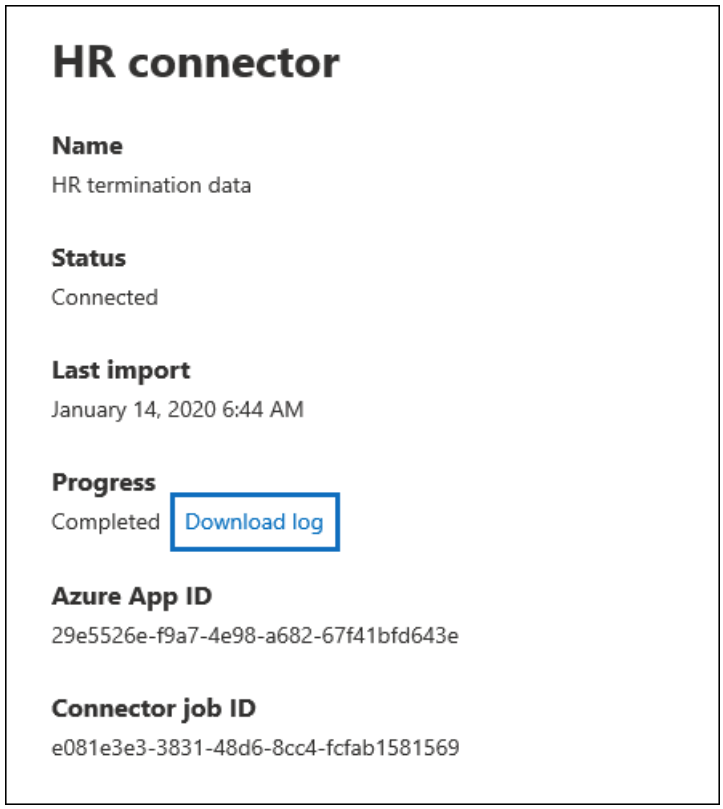

3. Under **Progress**, click the **Download log** link to open (or save) the status log for the connector. This log contains information about each time the script runs and uploads the data from the CSV file to the Microsoft cloud.

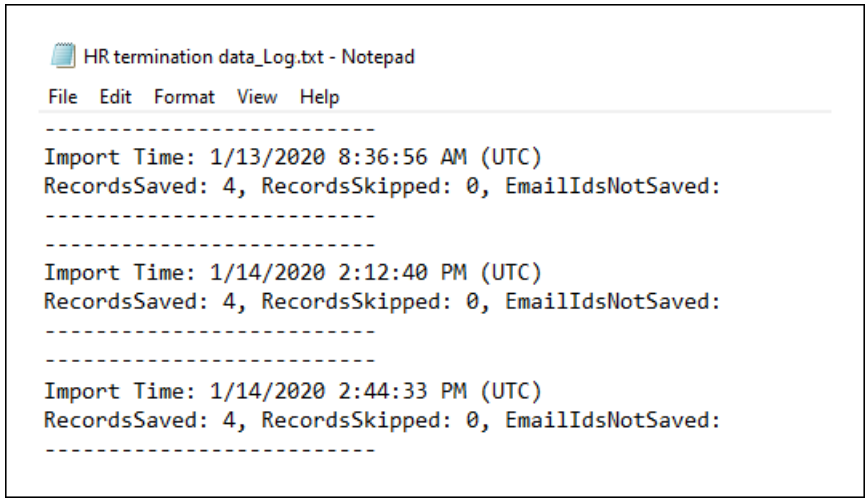

   The `RecordsSaved` field indicates the number of rows in the CSV file that uploaded. For example, if the CSV file contains four rows, then the value of the `RecordsSaved` fields is 4, if the script successfully uploaded all the rows in the CSV file.

If you've haven't run the script in Step 4, a link to download the script is displayed under **Last import**. You can download the script and then follow the steps to run the script.

(Optional) Step 6: Schedule the script to run automatically

To make sure the latest HR data from your organization is available to tools like the insider risk management solution, we recommend that you schedule the script to run automatically on a recurring basis, such as once a day. This also requires that you update the HR data in the CSV file on a similar (if not the same) schedule so that it contains the latest information about employees who leave your organization. The goal is to upload the most

current HR data so that the HR connector can make it available to the insider risk management solution.

You can user the Task Scheduler app in Windows to automatically run the script every day.

1. On your local computer, click the Windows **Start** button and then type **Task Scheduler**.
2. Click the **Task Scheduler** app to open it.
3. In the **Actions** section, click **Create Task**.
4. On the **General** tab, type a descriptive name for the scheduled task; for example, **HR Connector Script**. You can also add an optional description.
5. Under **Security options**, do the following:
   a. Determine whether to run the script only when you're logged on to the computer or run it when you're logged on or not.
   b. Make sure that the **Run with the highest privileges** checkbox is selected.
6. Select the **Triggers** tab, click **New**, and then do the following things:
   a. Under **Settings**, select the **Daily** option, and then choose a date and time to run the script for the first time. The script will every day at the same specified time.
   b. Under **Advanced settings**, make sure the **Enabled** checkbox is selected.
   c. Click **Ok**.
7. Select the **Actions** tab, click **New**, and then do the following things:

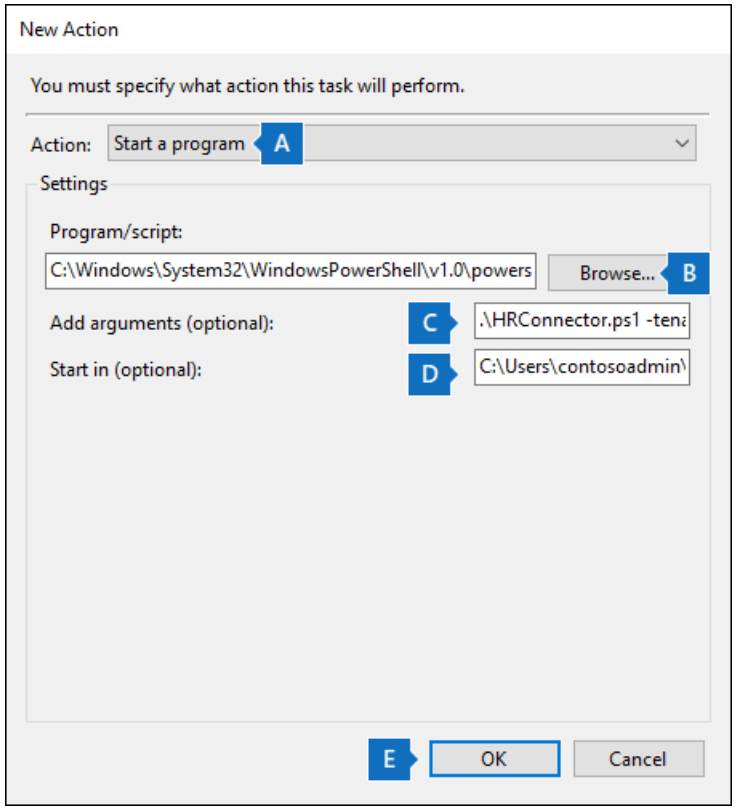

   a. In the **Action** dropdown list, make sure that **Start a program** is selected.
   b. In the **Program/script** box, click **Browse**, and go to the following location and select it so the path is displayed in the box: `C:\Windows\System32\WindowsPowerShell\v1.0\powershell.exe`.
   c. In the **Add arguments (optional)** box, paste the same script command that you ran in Step 4. For example,

```
.\HRConnector.ps1 -tenantId "d5723623-11cf-4e2e-b5a5-01d1506273g9" -appId "c12823b7-b55a-4989-faba-02de41bb97c3" -appSecret "MNubVGbcQDkGCnn" -jobId "e081f4f4-3831-48d6-7bb3-fcfab1581458" -csvFilePath "C:\Users\contosoadmin\Desktop\Data\employee_termination_data.csv"
```

d. In the **Start in (optional)** box, paste the folder location of the script that you ran in Step 4. For example, `C:\Users\contosoadmin\Desktop\Scripts`.

e. Click **Ok** to save the settings for the new action.

8. In the **Create Task** window, click **Ok** to save the scheduled task. You might be prompted to enter your user account credentials.

   The new task is displayed in the Task Scheduler Library.

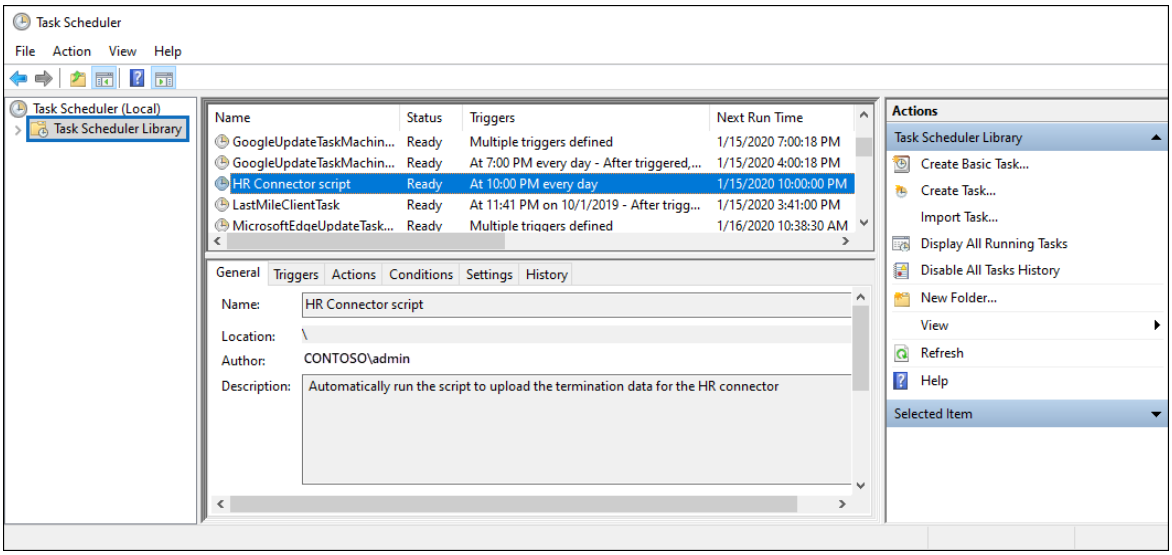

   The last time the script ran and the next time it's scheduled to run is displayed. You can double-click the task to edit it.

   You can also verify the last time the script ran on the flyout page of the corresponding HR connector in the compliance center.

Existing HR connectors

On July 20, 2020, we released additional scenarios that are supported by HR connectors. These are the HR scenarios that were previously described in this article. Any HR connector created before this date only supports the Employee resignation scenario. If you created an HR connector before July 20, 2020, we have migrated it so that it continues to migrate your HR data to the Microsoft cloud. You don't have to do anything to maintain this functionality. You can keep using the connector without any disruption.

If you want to implement additional HR scenarios, please create a new HR connector and configure it for the additional HR scenarios that were released. You'll also need to create one or more new CSV files that contain the data to support the additional HR scenarios. After you create a new HR connector, run the script using the job ID of the new connector and CSV file(s) with the data for your additional HR scenarios.

Set up a connector to archive ICE Chat data

2/18/2021 • 8 minutes to read • Edit Online

Use a native connector in the Microsoft 365 compliance center to import and archive financial services chat data from the ICE Chat collaboration tool. After you set up and configure a connector, it connects to your organization's ICE Chat secure FTP (SFTP) site once every day, converts the content of chat messages to an email message format, and then import those items to mailboxes in Microsoft 365.

After ICE chat data is stored in user mailboxes, you can apply Microsoft 365 compliance features such as litigation hold, eDiscovery, archiving, auditing, communication compliance, and Microsoft 365 retention policies to ICE Chat data. For example, you can search ICE Chat messages using content search or associate the mailbox that contains the ICE Chat data with a custodian in an Advanced eDiscovery case. Using an ICE Chat connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving ICE Chat data

The following overview explains the process of using a connector to archive ICE chat data in Microsoft 365.

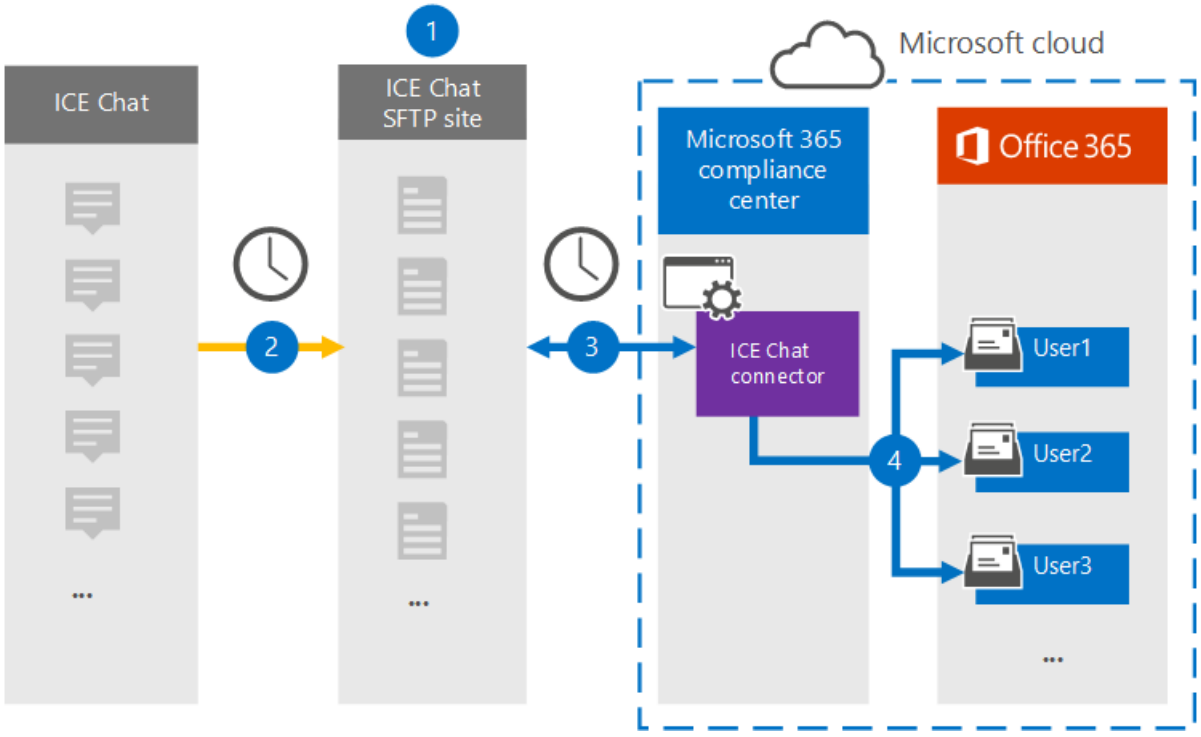

1. Your organization works with ICE Chat to set up an ICE Chat SFTP site. You'll also work with ICE Chat to configure ICE Chat to copy chat messages to your ICE Chat SFTP site.

2. Once every 24 hours, chat messages from ICE Chat are copied to your ICE Chat SFTP site.

3. The ICE Chat connector that you create in the Microsoft 365 compliance center connects to the ICE Chat SFTP site every day and transfers the chat messages from the previous 24 hours to a secure Azure Storage location in the Microsoft Cloud. The connector also converts the content of a chat massage to an email message format.

4. The connector imports chat message items to the mailboxes of specific users. A new folder named **ICE Chat** is created in the user mailboxes and the chat message items are imported to that folder. The connector does by using the value of the *SenderEmail* and *RecipientEmail* properties. Every chat message

contains these properties, which are populated with email address of the sender and every recipient/participant of the chat message.

In addition to automatic user mapping that uses the values of the *SenderEmail* and *RecipientEmail* property (which means that the connector imports a chat message to the sender's mailbox and the mailboxes of every recipient), you can also define custom user mapping by uploading a CSV mapping file. This mapping file contains the ICE Chat *ImId* and the corresponding Microsoft 365 mailbox address for every user in your organization. If you enable automatic user mapping and provide a custom-mapping file, for every chat item the connector will first look at the custom-mapping file. If it doesn't find a valid Microsoft 365 user account that corresponds to a user's ICE Chat ImId, the connector will use the *SenderEmail* and *RecipientEmail* properties of the chat item to import the item to the mailboxes of the chat participants. If the connector doesn't find a valid Microsoft 365 user in either the custom-mapping file or the *SenderEmail* and *RecipientEmail* properties, the item won't be imported.

Before you begin

Some of the implementation steps required to archive ICE Chat data are external to Microsoft 365 and must be completed before you can create the connector in the compliance center.

- ICE Chat charges their customers a fee for external compliance. Your organization should contact the ICE Chat sales group to discuss, and to sign the ICE Chat data services agreement, which you can obtain at https://www.theice.com/publicdocs/agreements/ICE_Data_Services_Agreement.pdf. This agreement is between ICE Chat and your organization and does not involve Microsoft. After you set up an ICE Chat SFTP site in Step 2, ICE Chat provides the FTP credentials directly to your organization. Then you who would provide those credentials to Microsoft when setting up the connector in Step 3.
- You must set up an ICE Chat SFTP site before creating the connector in Step 3. After working with ICE Chat to set up the SFTP site, data from ICE Chat is uploaded to the SFTP site every day. The connector you create in Step 3 connects to this SFTP site and transfers the chat data to Microsoft 365 mailboxes. SFTP also encrypts the ICE Chat data that's sent to mailboxes during the transfer process.
- The ICE Chat connector can import a total of 200,000 items in a single day. If there are more than 200,000 items on the SFTP site, none of those items will be imported to Microsoft 365.
- The admin who creates the ICE Chat connector in Step 3 (and who downloads the public keys and IP address in Step 1) must be assigned the Mailbox Import Export role in Exchange Online. This role is required to add connectors on the **Data connectors** page in the Microsoft 365 compliance center. By default, this role isn't assigned to any role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Step 1: Obtain SSH and PGP public keys

The first step is to obtain a copy of the public keys for Secure Shell (SSH) and Pretty Good Privacy (PGP). You use these keys in Step 2 to configure the ICE Chat SFTP site to allow the connector (that you create in Step 3) to connect to the SFTP site and transfer the ICE Chat data to Microsoft 365 mailboxes. You will also obtain an IP address in this step, which you use when configuring the ICE Chat SFTP site.

1. Go to https://compliance.microsoft.com and click **Data connectors** in the left nav.
2. On the **Data connectors** page under **ICE Chat**, click **View**.
3. On the **ICE Chat** page, click **Add connector**.

4. On the **Terms of service** page, click **Accept**.
5. On the **Add credentials for ICE Chat SFTP site** page under step 1, click the **Download SSH key**, **Download PGP key**, and **Download IP address** links to save a copy of each file to your local computer. These files contain the following items that are used to configure the ICE Chat SFTP site in Step 2:
   - SSH public key: This key is used to configure Secure SSH to enable a secure remote login when the connector connects to the ICE Chat SFTP site.
   - PGP public key: This key is used to configure the encryption of data that's transferred from the ICE Chat SFTP site to Microsoft 365.
   - IP address: The ICE Chat SFTP site is configured to accept a connection request only from this IP address, which is used by the ICE Chat connector that you create in Step 3.
6. Click **Cancel** to close the wizard. You come back to this wizard in Step 3 to create the connector.

Step 2: Configure the ICE Chat SFTP site

The next step is to use the SSH and PGP public keys and the IP address that you obtained in Step 1 to configure SSH authentication and PGP encryption for the ICE Chat SFTP site. This lets the ICE Chat connector that you create in Step 3 connect to the ICE Chat SFTP site and transfer ICE Chat data to Microsoft 365. You need to work with ICE Chat customer support to set up your ICE Chat SFTP site.

Step 3: Create an ICE Chat connector

The last step is to create an ICE Chat connector in the Microsoft 365 compliance center. The connector uses the information you provide to connect to the ICE Chat SFTP site and transfer chat messages to the corresponding user mailbox boxes in Microsoft 365.

1. Go to https://compliance.microsoft.com and click **Data connectors** in the left nav.
2. On the **Data connectors** page under **ICE Chat**, click **View**.
3. On the **ICE Chat** page, click **Add connector**.
4. On the **Terms of service** page, click **Accept**.
5. On the **Add credentials for ICE Chat SFTP site** page, under Step 3, enter the required information in the following boxes and then click **Validate connection**.
   - **Firm code:** The ID for your organization, which is used as the username for the ICE Chat SFTP site.
   - **Password:** The password for your ICE Chat SFTP site.
   - **SFTP URL:** The URL for the ICE Chat SFTP site (for example, sftp.theice.com).
   - **SFTP port:** The port number for the ICE Chat SFTP site. The connector uses this port to connect to the SFTP site.
6. After the connection is validated, click **Next**.
7. On the **Map external users to Microsoft 365 users** page, enable automatic user mapping and provide custom user mapping as required. You can download a copy of the user-mapping CSV file on this page. You can add the user mappings to the file and then upload it.

NOTE

As previously explained, custom mapping file CSV file contains the ICE Chat imid and corresponding Microsoft 365 mailbox address for each user. If you enable automatic user mapping and provide a custom mapping, for every chat item, the connector will first look at custom mapping file. If it doesn't find a valid Microsoft 365 user that corresponds to a user's ICE Chat imid, the connector will import the item to the mailboxes for the users specified in the *SenderEmail* and *RecipientEmail* properties of the chat item. If the connector doesn't find a valid Microsoft 365 user by either automatic or custom user mapping, the item won't be imported.

8. Click **Next**, review your settings, and then click **Finish** to create the connector.
9. Go to the **Data connectors** page to see the progress of the import process for the new connector.

Set up a connector to archive Instant Bloomberg data

2/18/2021 • 8 minutes to read • Edit Online

Use a native connector in the Microsoft 365 compliance center to import and archive financial services chat data from the Instant Bloomberg collaboration tool. After you set up and configure a connector, it connects to your organization's Bloomberg secure FTP site (SFTP) once every day, converts the content of chat messages to an email message format, and then imports those items to mailboxes in Microsoft 365.

After Instant Bloomberg data is stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, Content Search, In-Place Archiving, Auditing, Communication compliance, and Microsoft 365 retention policies to Instant Bloomberg data. For example, you can search Instant Bloomberg chat messages using Content Search or associate the mailbox that contains the Instant Bloomberg data with a custodian in an Advanced eDiscovery case. Using an Instant Bloomberg connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving Instant Bloomberg data

The following overview explains the process of using a connector to archive Instant Bloomberg chat data in Microsoft 365.

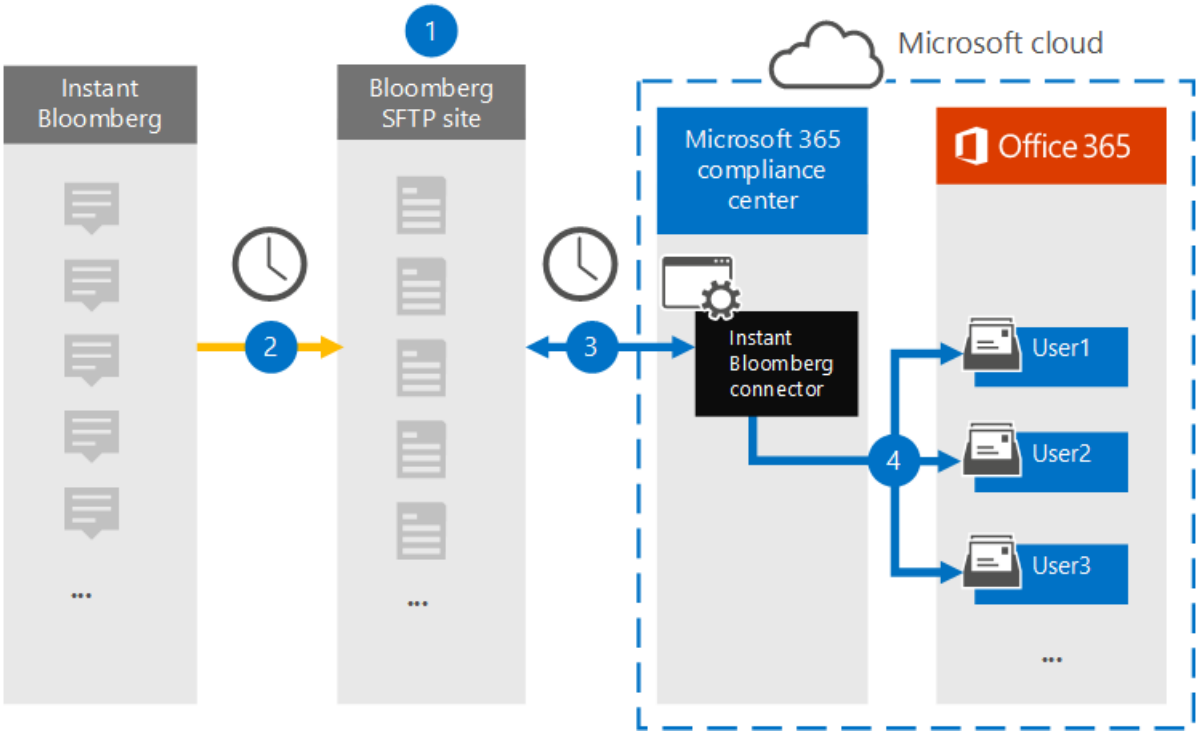

1. Your organization works with Bloomberg to set up a Bloomberg SFTP site. You'll also work with Bloomberg to configure Instant Bloomberg to copy chat messages to your Bloomberg SFTP site.
2. Once every 24 hours, chat messages from Instant Bloomberg are copied to the Bloomberg SFTP site.
3. The Instant Bloomberg connector that you create in the Microsoft 365 compliance center connects to the Bloomberg SFTP site every day and transfers the chat messages from the previous 24 hours to a secure Azure Storage area in the Microsoft Cloud. The connector also converts the content of a chat massage to an email message format.

4. The connector imports the chat message items to the mailbox of a specific user. A new folder named InstantBloomberg is created in the specific user's mailbox and the items will be imported to it. The connector does this by using the value of the *CorporateEmailAddress* property. Every chat message contains this property, which is populated with the email address of every participant of the chat message. In addition to automatic user mapping using the value of the *CorporateEmailAddress* property, you can also define a custom mapping by uploading a CSV mapping file. This mapping file should contain a Bloomberg UUID and the corresponding Microsoft 365 mailbox address for each user. If you enable automatic user mapping and provide a custom mapping, for every chat item the connector will first look at custom-mapping file. If it doesn't find a valid Microsoft 365 user that corresponds to a user's Bloomberg UUID, the connector will use the *CorporateEmailAddress* property of the chat item. If the connector doesn't find a valid Microsoft 365 user in either the custom-mapping file or the *CorporateEmailAddress* property of the chat item, the item won't be imported.

Before you begin

Some of the implementation steps required to archive Instant Bloomberg data are external to Microsoft 365 and must be completed before you can create the connector in the compliance center.

- Subscribe to Bloomberg Anywhere. This is required so that you can log in to Bloomberg Anywhere to access the Bloomberg SFTP site that you have to set up and configure.
- Set up a Bloomberg SFTP (Secure file transfer protocol) site. After working with Bloomberg to set up the SFTP site, data from Instant Bloomberg is uploaded to the SFTP site every day. The connector you create in Step 2 connects to this SFTP site and transfers the chat data to Microsoft 365 mailboxes. SFTP also encrypts the Instant Bloomberg chat data that is sent to mailboxes during the transfer process.

  For information about Bloomberg SFTP (also called *BB-SFTP*):

  - See the "SFTP Connectivity Standards" document at Bloomberg Support.
  - Contact Bloomberg customer support.

  After you work with Bloomberg to set up an SFTP site, Bloomberg will provide some information to you after you respond to the Bloomberg implementation email message. Save a copy of the following information. You use it to set up a connector in Step 3.

  - Firm code, which is an ID for your organization and is used to log in to the Bloomberg SFTP site.
  - Password for your Bloomberg SFTP site
  - URL for Bloomberg SFTP site (for example, sftp.bloomberg.com)
  - Port number for Bloomberg SFTP site
- The Instant Bloomberg connector can import a total of 200,000 items in a single day. If there are more than 200,000 items on the SFTP site, none of those items will be imported to Microsoft 365.
- The user who creates an Instant Bloomberg connector in Step 3 (and who downloads the public keys and IP address in Step 1) must be assigned the Mailbox Import Export role in Exchange Online. This is required to add connectors in the **Data connectors** page in the Microsoft 365 compliance center. By default, this role isn't assigned to any role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Step 1: Obtain SSH and PGP public keys

The first step is to obtain a copy of the public keys for Secure Shell (SSH) and Pretty Good Privacy (PGP). You use these keys in Step 2 to configure the Bloomberg SFTP site to allow the connector (that you create in Step 3) to connect to the SFTP site and transfer the Instant Bloomberg chat data to Microsoft 365 mailboxes. You also obtain an IP address in this step, which you use when configuring the Bloomberg SFTP site.

1. Go to https://compliance.microsoft.com and then click **Data connectors** > **Instant Bloomberg**.
2. On the **Instant Bloomberg** product description page, click **Add connector**
3. On the **Terms of service** page, click **Accept**.
4. On the **Add credentials for Bloomberg SFTP site** under step 1, click the **Download SSH key**, **Download PGP key**, and **Download IP address** links to save a copy of each file to your local computer. These files contain the following items that are used to configure the Bloomberg SFTP site in Step 2:
   - SSH public key: This key is used to configure Secure Shell (SSH) to enable a secure remote login when the connector connects to the Bloomberg SFTP site.
   - PGP public key: This key is used to configure the encryption of data that's transferred from the Bloomberg SFTP site to Microsoft 365.
   - IP address: The Bloomberg SFTP site is configured to accept a connection request only from this IP address, which is used by the Instant Bloomberg connector that you create in Step 3.
5. Click **Cancel** to close the wizard. You come back to this wizard in Step 3 to create the connector.

Step 2: Configure the Bloomberg SFTP site

The next step is to use the SSH and PGP public keys and the IP address that you obtained in Step 1 to configure SSH authentication and PGP encryption for the Bloomberg SFTP site. This lets the Instant Bloomberg connector that you create in Step 3 connect to the Bloomberg SFTP site and transfer Instant Bloomberg data to Microsoft 365. You need to work with Bloomberg customer support to set up your Bloomberg SFTP site. Contact Bloomberg customer support for assistance.

> **IMPORTANT**
> Bloomberg recommends that you attach the three files that you downloaded in Step 1 to an email message and send it to their customer support team when working with them to set up your Bloomberg SFTP site.

Step 3: Create an Instant Bloomberg connector

The last step is to create an Instant Bloomberg connector in the Microsoft 365 compliance center. The connector uses the information you provide to connect to the Bloomberg SFTP site and transfer chat messages to the corresponding user mailbox boxes in Microsoft 365.

1. Go to https://compliance.microsoft.com and then click **Data connectors** > **Instant Bloomberg**.
2. On the **Instant Bloomberg** product description page, click **Add connector**
3. On the **Terms of service** page, click **Accept**.
4. On the **Add credentials for Bloomberg SFTP site** page, under Step 3, enter the required information in the following boxes and then click **Next**.
   - **Firm code:** The ID for your organization that is used as the username for the Bloomberg SFTP site.

- **Password:** Password for Bloomberg SFTP site.
- **SFTP URL:** The URL for Bloomberg SFTP site (for example, sftp.bloomberg.com).
- **SFTP port:** The port number for Bloomberg SFTP site. The connector uses this port to connect to the SFTP site.

5. On the **Select data types to import** page, select the required data types to be imported apart from **Messages**
6. On the **User-mapping** page, enable automatic user mapping and provide custom user mapping as required

> **NOTE**
>
> The connector imports the chat message items to the mailbox of a specific user. A new folder named **InstantBloomberg** is created in the specific user's mailbox and the items will be imported to it. The connector does by using the value of the *CorporateEmailAddress* property. Every chat message contains this property, and the property is populated with the email address of every participant of the chat message. In addition to automatic user mapping using the value of the *CorporateEmailAddress* property, you can also define custom mapping by uploading a CSV mapping file. The mapping file should contain the Bloomberg UUID and corresponding Microsoft 365 mailbox address for each user. If you enable automatic user mapping and provide a custom mapping, for every chat item the connector will first look at custom mapping file. If it doesn't find a valid Microsoft 365 user that corresponds to a user's Bloomberg UUID, the connector will use the *CorporateEmailAddress* property of the chat item. If the connector doesn't find a valid Microsoft 365 user in either the custom mapping file or the *CorporateEmailAddress* property of the chat item, the item won't be imported.

7. Click **Next**, review your settings, and then click **prepare** to create the connector.
8. Go to the **Data connectors** page to see the progress of the import process for the new connector.

Set up a connector to archive LinkedIn data

2/18/2021 • 3 minutes to read • Edit Online

Use a connector in the Microsoft 365 compliance center to import and archive data from LinkedIn Company pages. After you set up and configure a connector, it connects to the account for the specific LinkedIn Company page once every 24 hours. The connector converts the messages posted to the Company page to an email message, and then imports those items to a mailbox in Microsoft 365.

After the LinkedIn Company page data is stored in a mailbox, you can apply Microsoft 365 compliance features such as Litigation Hold, Content Search, In-Place Archiving, Auditing, and Microsoft 365 retention policies to LinkedIn data. For example, you can search for these items using Content Search or associate the storage mailbox with a custodian in an Advanced eDiscovery case. Creating a connector to import and archive LinkedIn data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Before you set up a connector

- The user who creates a LinkedIn Company Page connector must be assigned the Mailbox Import Export role in Exchange Online. This is required to add connectors in the **Data connectors** page in the Microsoft 365 compliance center. By default, this role isn't assigned to any role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".
- You must have the sign-in credentials (email address or phone number and password) of a LinkedIn user account that is an admin for the LinkedIn Company Page that you want to archive. You use these credentials to sign into LinkedIn when setting up the connector.
- The LinkedIn connector can import a total of 200,000 items in a single day. If there are more than 200,000 LinkedIn items in a day, none of those items will be imported to Microsoft 365.

Create a LinkedIn connector

1. Go to https://compliance.microsoft.com and then click **Data connectors** > **LinkedIn Company pages**.
2. On the **LinkedIn company pages** product page, click **Add connector**.
3. On the **Terms of service** page, select **Accept**.
4. On the **Sign in with LinkedIn** page, click **Sign in with LinkedIn**.

   The LinkedIn sign-in page is displayed.

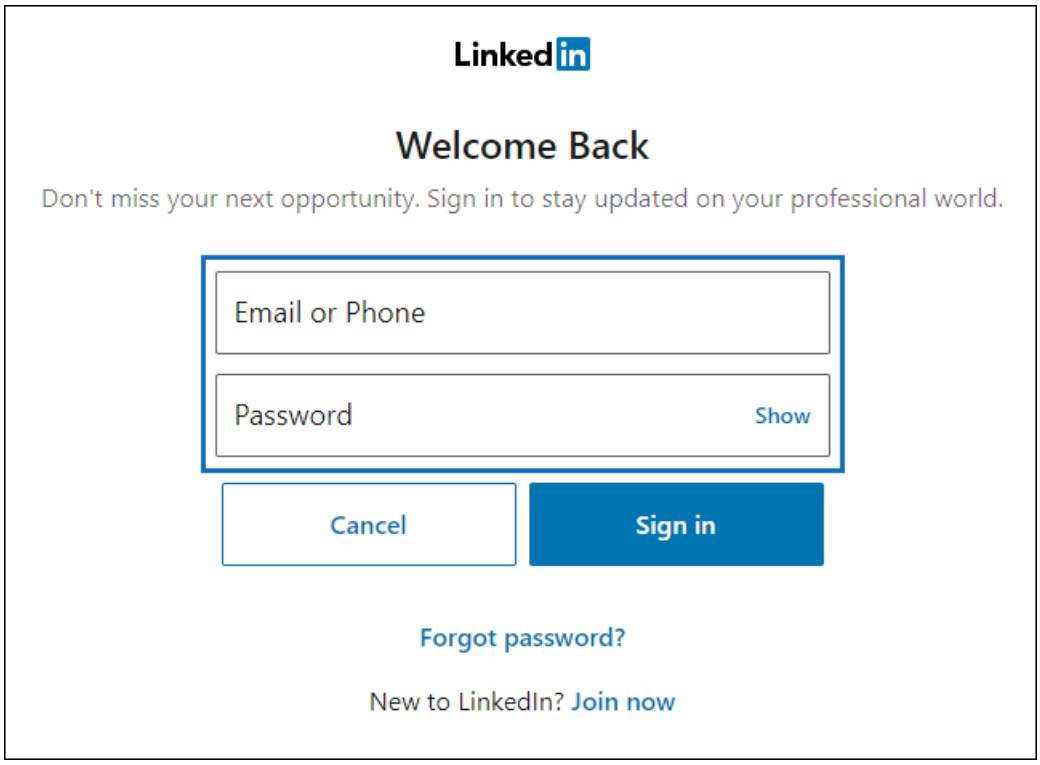

5. On the LinkedIn sign in page, enter the email address (or phone number) and password for the LinkedIn account associated with the company page that you want to archive, and then click **Sign in**.

   A wizard page is displayed with a list of all LinkedIn Company Pages associated with the account that you signed in to. A connector can only be configured for one company page. If your organization has multiple LinkedIn Company Pages, you have to create a connector for each one.

6. Select the company page that you want to archive items from, and then click **Next**.

7. On the **Choose storage location** page, click in the box, select the email address of a Microsoft 365 mailbox that the LinkedIn items will be imported to, and then click **Next**. Items are imported to the inbox folder in this mailbox.

8. Click **Next** to review the connector settings and then click **Finish** to complete the connector setup.

After you create the connector, you can go back to the **Data connectors** page to see the progress of the import process for the new connector (select **Refresh** if necessary to update the list of connectors). The value in the **Status** column is **Waiting to start**. It takes up to 24 hours for the initial import process to be started. After the first time the connector runs and imports the LinkedIn items, the connector will run once every 24 hours and import any new items that are created on the LinkedIn Company Page in the previous 24 hours.

To view more details, select the connector in the list on the **Data connectors** page to display the flyout page. Under **Status**, the date range that's displayed indicates the age filter that was selected when the connector was created.

More information

LinkedIn items are imported to the LinkedIn subfolder in the inbox of the storage mailbox in Microsoft 365. They appear as email messages.

Set up a connector to import physical badging data (preview)

2/18/2021 • 12 minutes to read • Edit Online

You can set up a data connector in the Microsoft 365 compliance center to import physical badging data, such as employee's raw physical access events or any physical access alarms generated by your organization's badging system. Examples of physical access points are an entry to a building or an entry to server room or data center. Physical badging data can be used by the Microsoft 365 insider risk management solution to help protect your organization from malicious activity or data theft inside your organization.

Setting up a physical badging connector consists of the following tasks:

- Creating an app in Azure Active Directory (Azure AD) to access an API endpoint that accepts a JSON payload that contains physical badging data.
- Creating the JSON payload with a schema defined by physical badging data connector.
- Creating a physical badging data connector in the Microsoft 365 compliance center.
- Running a script to push the physical badging data to the API endpoint.
- Optionally, scheduling the script to run automatically to import currently physical badging data.

Before you set up the connector

- The user who creates the physical badging connector in Step 3 must be assigned the Mailbox Import Export role in Exchange Online. By default, this role isn't assigned to any role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a new role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".
- You need to determine how to retrieve or export the data from your organization's physical badging system (on a daily basis) and create a JSON file that's described in Step 2. The script that you run in Step 4 will push the data in the JSON file to the API endpoint.
- The sample script that you run in Step 4 pushes the physical badging data from JSON file to the connector API so that it can be used by the insider risk management solution. This sample script isn't supported under any Microsoft standard support program or service. The sample script is provided AS IS without warranty of any kind. Microsoft further disclaims all implied warranties including, without limitation, any implied warranties of merchantability or of fitness for a particular purpose. The entire risk arising out of the use or performance of the sample script and documentation remains with you. In no event shall Microsoft, its authors, or anyone else involved in the creation, production, or delivery of the scripts be liable for any damages whatsoever (including, without limitation, damages for loss of business profits, business interruption, loss of business information, or other pecuniary loss) arising out of the use of or inability to use the sample scripts or documentation, even if Microsoft has been advised of the possibility of such damages.

Step 1: Create an app in Azure Active Directory

The first step is to create and register a new app in Azure Active Directory (Azure AD). The app will correspond to the physical badging connector that you create in Step 3. Creating this app will allow Azure AD to authenticate

the push request for JSON payload containing physical badging data. During the creation of this Azure AD app, be sure to save the following information. These values will be used in later steps.

- Azure AD application ID (also called the *app Id* or *client Id*)
- Azure AD application secret (also called the *client secret*)
- Tenant Id (also called the *directory Id*)

For step-by-step instructions for creating an app in Azure AD, see Register an application with the Microsoft identity platform.

Step 2: Prepare a JSON file with physical badging data

The next step is to create a JSON file that contains information about employees' physical access data. As explained in the before you begin section, you'll need to determine how to generate this JSON file from your organization's physical badging system.

The JSON file must conform to the schema definition required by the connector. Here are descriptions of the required schema properties for the JSON file:

PROPERTY	DESCRIPTION	DATA TYPE
UserId	An employee can have multiple digital identities across the systems. The input needs to have the Azure AD ID already resolved by the source system.	UPN or email address
AssetId	The reference ID of the physical asset or physical access point.	Alphanumeric string
AssetName	The friendly name of the physical asset or physical access point.	Alphanumeric string
EventTime	The time stamp of access.	Date and time, in UTC format
AccessStatus	Value of `Success` or `Failed`	String

Here's an example of a JSON file that conforms to the required schema:

```
[
    {
        "UserId":"sarad@contoso.com"
        "AssetId":"Mid-Sec-7",
        "AssetName":"Main Building 1st Floor Mid Section",
        "EventTime":"2019-07-04T01:57:49",
        "AccessStatus":"Failed",
    },
    {
        "UserId":"pilarp@contoso.com",
        "AssetId":"Mid-Sec-7",
        "AssetName":"Main Building 1st Floor Mid Section",
        "EventTime":"2019-07-04T02:57:49",
        "AccessStatus":"Success",
    }
]
```

You can also download the following schema definition for the JSON file from the wizard when you create the

physical badging connector in Step 3.

```
{
 "title" : "Physical Badging Signals",
 "description" : "Access signals from physical badging systems",
 "DataType" : {
  "description" : "Identify what is the data type for input signal",
  "type" : "string",
 },
 "type" : "object",
 "properties": {
  "UserId" : {
   "description" : "Unique identifier AAD Id resolved by the source system",
   "type" : "string",
  },
  "AssetId": {
   "description" : "Unique ID of the physical asset/access point",
   "type" : "string",
  },
  "AssetName": {
   "description" : "friendly name of the physical asset/access point",
   "type" : "string",
  },
  "EventTime" : {
   "description" : "timestamp of access",
   "type" : "string",
  },
  "AccessStatus" : {
   "description" : "what was the status of access attempt - Success/Failed",
   "type" : "string",
  },
 }
 "required" : ["UserId", "AssetId", "EventTime" "AccessStatus"]
}
```

Step 3: Create the physical badging connector

The next step is to create a physical badging connector in the Microsoft 365 compliance center. After you run the script in Step 4, the JSON file that you created in Step 3 will be processed and pushed to the API endpoint you configured in Step 1. In this step, be sure to copy the JobId that's generated when you create the connector. You'll use the JobId when you run the script.

1. Go to https://compliance.microsoft.com and then click **Data connectors** in the left nav.
2. On the **Data connectors** page under **Physical badging**, click **View**.
3. On the **Physical badging** page, click **Add connector**.
4. On the **Authentication credentials** page, do the following and then click **Next**:
   a. Type or paste the Azure AD application ID for the Azure app that you created in Step 1.
   b. Download the sample schema for your reference to create the JSON file.
   c. Type a unique name for the physical badging connector.
5. On the **Review** page, review your settings and then click **Finish** to create the connector.
6. A status page is displayed that confirms the connector was created. This page also contains the job ID. You can copy job ID from this page or from the flyout page for the connector. You need this job ID when running the script.

   The status page also contains a link to the script. Refer to this script to understand how to post the JSON

file to the API endpoint.

7. Click **Done**.

   The new connector is displayed in the list on the **Connectors** tab.

8. Click the physical badging connector that you just created to display the flyout page, which contains properties and other information about the connector.

Step 4: Run the script to POST your JSON file containing physical badging data

The next step in setting up a physical badging connector is to run a script that will push the physical badging data in the JSON file (that you created in Step 2) to the API endpoint you created in Step 1. We provide a sample script for your reference and you can choose to use it or create your own script to post the JSON file to the API endpoint.

After you run the script, the JSON file containing the physical badging data is pushed to your Microsoft 365 organization where it can accessed by the insider risk management solution. We recommend you post physical badging data daily. You can do this by automating the process to generate the JSON file every day from your physical badging system and then scheduling the script to push the data.

> **NOTE**
> The maximum number of records in the JSON file that can be processed by the API is 50,000 records.

1. Go to this GitHub site to access the sample script.
2. Click the **Raw** button to display the script in text view
3. Copy all the lines in the sample script and then save them to a text file.
4. Modify the sample script for your organization, if necessary.
5. Save the text file as a Windows PowerShell script file by using a filename suffix of .ps1; for example, PhysicalBadging.ps1.
6. Open a Command Prompt on your local computer, and go to the directory where you saved the script.
7. Run the following command to push the physical badging data in the JSON file to the Microsoft cloud; for example:

```
.\PhysicalBadging.ps1 -tenantId "<Tenant Id>" -appId "<Azure AD App Id>" -appSecret "<Azure AD App
Secret>" -jobId "Job Id" -jsonFilePath "<records file path>"
```

   The following table describes the parameters to use with this script and their required values. Information you obtained in the previous steps is used in the values for these parameters.

PARAMETER	DESCRIPTION
tenantId	This is the Id for your Microsoft 365 organization that you obtained in Step 1. You can also obtain the tenantId for your organization on the **Overview** blade in the Azure AD admin center. This is used to identify your organization.

PARAMETER	DESCRIPTION
appId	This is the Azure AD application Id for the app that you created in Azure AD in Step 1. This is used by Azure AD for authentication when the script attempts to accesses your Microsoft 365 organization.
appSecret	This is the Azure AD application secret for the app that you created in Azure AD in Step 1. This is also used for authentication.
jobId	This is the Job Id for the physical badging connector that you created in Step 3. This is used to associate the physical badging data that is pushed to the Microsoft cloud with the physical badging connector.
JsonFilePath	This is the file path on the local computer (the one you're using to run the script) for the JSON file that you created in Step 2. This file must follow the sample schema described in Step 3.

Here's an example of the syntax for the physical badging connector script using actual values for each parameter:

```
.\PhysicalBadging.ps1 -tenantId d5723623-11cf-4e2e-b5a5-01d1506273g9 -appId 29ee526e-f9a7-4e98-a682-67f41bfd643e -appSecret MNubVGbcQDkGCnn -jobId b8be4a7d-e338-43eb-a69e-c513cd458eba -csvFilePath 'C:\Users\contosoadmin\Desktop\Data\physical_badging_data.json'
```

If the upload is successful, the script displays the **Upload Successful** message.

If you have multiple JSON files, you have to run the script for each file.

> **NOTE**
> You can also choose to push the physical badging data to the API endpoint by methods other than running the previous script. For example, here's a sample for using Postman to push your data to the API endpoint.

Step 5: Monitor the physical badging connector

After you create the physical badging connector and push your physical badging data, you can view the connector and upload status in the Microsoft 365 compliance center. If you schedule the script to run automatically on a regular basis, you can also view the current status after the last time the script ran.

1. Go to https://compliance.microsoft.com and click **Data connectors** in the left nav.
2. Click the **Connectors** tab and then select the physical badging connector to display the flyout page. This page contains the properties and information about the connector.

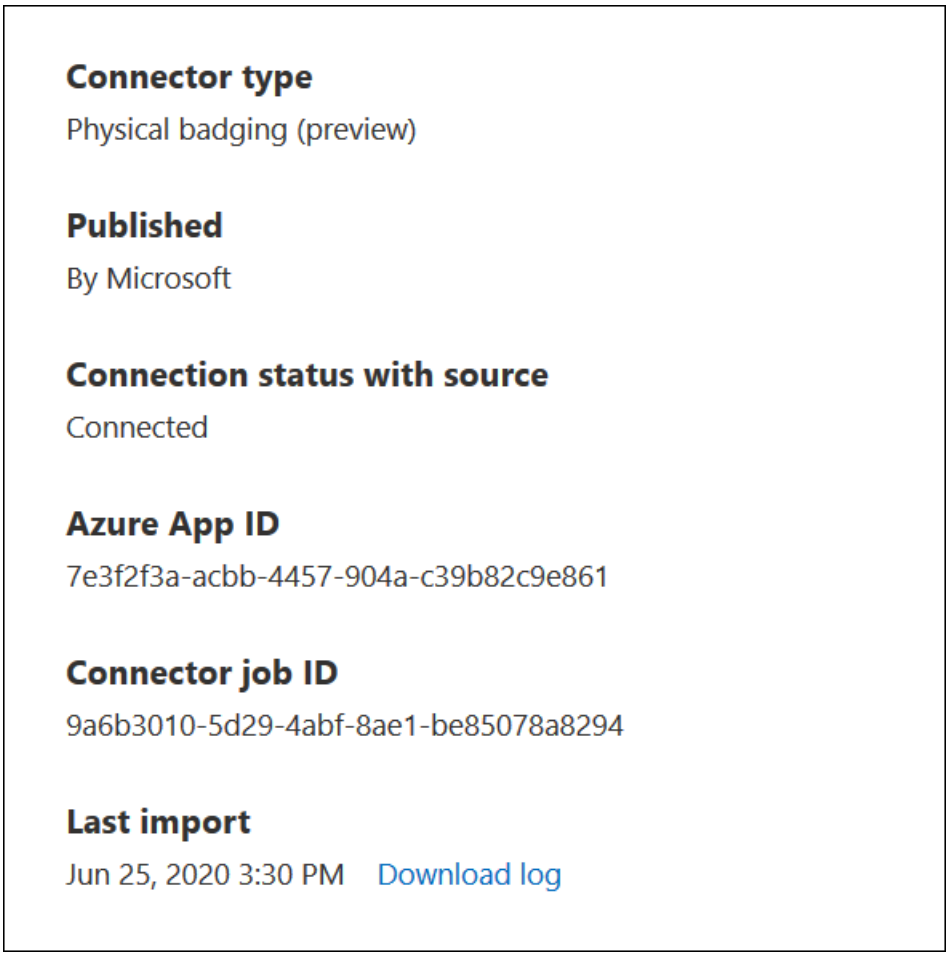

3. Under **Last import**, click the **Download log** link to open (or save) the status log for the connector. This log contains information about each time the script runs and uploads the data from the CSV file to the Microsoft cloud.

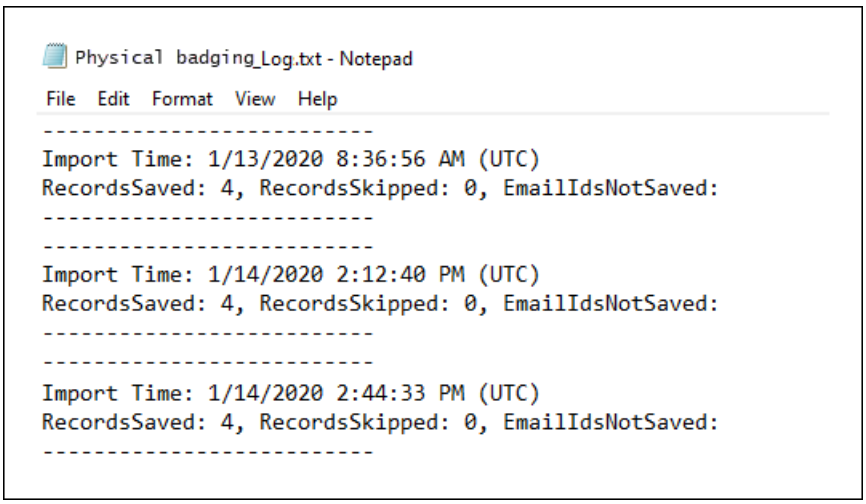

   The **RecordsSaved** field indicates the number of rows in the CSV file that uploaded. For example, if the CSV file contains four rows, then the value of the **RecordsSaved** fields is 4, if the script successfully uploaded all the rows in the CSV file.

If you've haven't run the script in Step 4, a link to download the script is displayed under **Last import**. You can download the script and then follow the steps in Step 4 to run it.

(Optional) Step 6: Schedule the script to run automatically

To make sure the latest physical badging data from your organization is available to tools like the insider risk management solution, we recommend that you schedule the script to run automatically on a recurring basis,

such as once a day. This also requires that you update the physical badging data to JSON file on a similar (if not the same) schedule so that it contains the latest information about employees who leave your organization. The goal is to upload the most current physical badging data so that the physical badging connector can make it available to the insider risk management solution.

You can user the Task Scheduler app in Windows to automatically run the script every day.

1. On your local computer, click the Windows **Start** button and then type **Task Scheduler**.
2. Click the **Task Scheduler** app to open it.
3. In the **Actions** section, click **Create Task**.
4. On the **General** tab, type a descriptive name for the scheduled task; for example, **physical badging connector Script**. You can also add an optional description.
5. Under **Security options**, do the following things:
   a. Determine whether to run the script only when you're logged on to the computer or run it when you're logged on or not.
   b. Make sure that the **Run with the highest privileges** checkbox is selected.
6. Select the **Triggers** tab, click **New**, and then do the following things:
   a. Under **Settings**, select the **Daily** option, and then choose a date and time to run the script for the first time. The script will every day at the same specified time.
   b. Under **Advanced settings**, make sure the **Enabled** checkbox is selected.
   c. Click **Ok**.
7. Select the **Actions** tab, click **New**, and then do the following things:

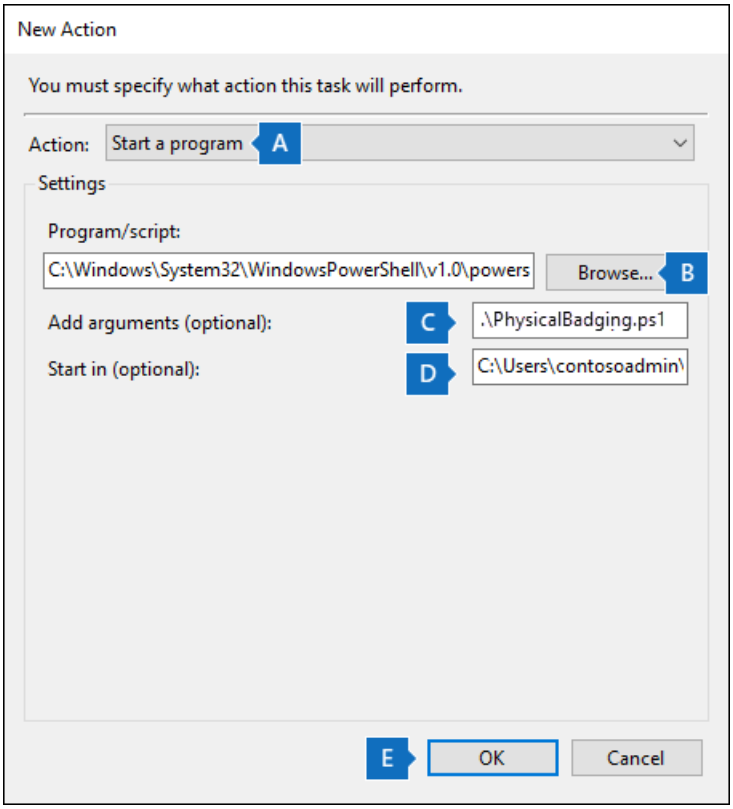

   a. In the **Action** dropdown list, make sure that **Start a program** is selected.
   b. In the **Program/script** box, click **Browse**, and go to the following location and select it so the path is displayed in the box: C:\Windows\System32\WindowsPowerShell\v1.0\powershell.exe.

c. In the **Add arguments (optional)** box, paste the same script command that you ran in Step 4. For example, .\PhysicalBadging.ps1-tenantId "d5723623-11cf-4e2e-b5a5-01d1506273g9" -appId "c12823b7-b55a-4989-faba-02de41bb97c3" -appSecret "MNubVGbcQDkGCnn" -jobId "e081f4f4-3831-48d6-7bb3-fcfab1581458" -jsonFilePath "C:\Users\contosoadmin\Desktop\Data\physical_badging_data.csv"

d. In the **Start in (optional)** box, paste the folder location of the script that you ran in Step 4. For example, C:\Users\contosoadmin\Desktop\Scripts.

e. Click **Ok** to save the settings for the new action.

8. In the **Create Task** window, click **Ok** to save the scheduled task. You might be prompted to enter your user account credentials.

   The new task is displayed in the Task Scheduler Library.

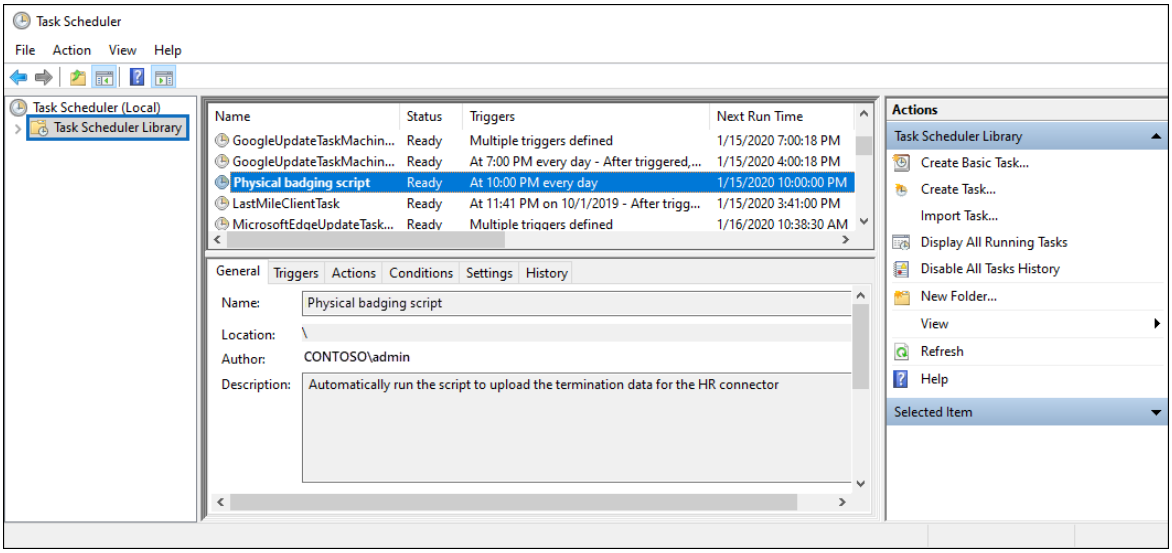

The last time the script ran and the next time it's scheduled to run is displayed. You can double-click the task to edit it.

You can also verify the last time the script ran on the flyout page of the corresponding physical badging connector in the compliance center.

Set up a connector to archive Twitter data (preview)

2/18/2021 • 6 minutes to read • Edit Online

Use a connector in the Microsoft 365 compliance center to import and archive data from Twitter to Microsoft 365. After you set up and configure the connector, it connects to your organization's Twitter account (on a scheduled basis), converts the content of an item to an email message format, and then imports those items to a mailbox in Microsoft 365.

After the Twitter data is imported, you can apply Microsoft 365 compliance features such as Litigation Hold, Content Search, In-Place Archiving, Auditing, and Microsoft 365 retention policies to the Twitter data. For example, when a mailbox is placed on Litigation Hold or assigned to a retention policy, the Twitter data is preserved. You can search third-party data using Content Search or associate the mailbox where the Twitter data is stored with a custodian in an Advanced eDiscovery case. Using a connector to import and archive Twitter data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

After Twitter data is imported, you can apply Microsoft 365 compliance features such as Litigation Hold, Content Search, In-Place Archiving, Auditing, Communication compliance, and Microsoft 365 retention policies to the data stored in the mailbox. For example, you can search Twitter data using Content Search or associate the mailbox where the data is stored with a custodian in an Advanced eDiscovery case. Using a connector to import and archive Twitter data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Before you set up a connector

Complete the following prerequisites before you can set up and configure a connector in the Microsoft 365 compliance center to import and archive data from your organization's Twitter account.

- You need a Twitter account for your organization; you need to sign in to this account when setting up the connector.
- Your organization must have a valid Azure subscription. If you don't have an existing Azure subscription, you can sign up for one of these options:
  - Sign up for a free one year Azure subscription
  - Sign up for a Pay-As-You-Go Azure subscription

  > **NOTE**
  > The free Azure Active Directory subscription that's included with your Microsoft 365 subscription doesn't support the connectors in the Security & Compliance Center.

- The Twitter connector can import a total of 200,000 items in a single day. If there are more than 200,000 Twitter items in a day, none of those items will be imported to Microsoft 365.
- The user who sets up the Twitter connector in the Microsoft 365 compliance center (in Step 5) must be assigned the Mailbox Import Export role in Exchange Online. By default, this role isn't assigned to any role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Step 1: Create an app in Azure Active Directory

The first step is to register a new app in Azure Active Directory (AAD). This app corresponds to the web app resource that you implement in Step 2 for the Twitter connector.

For step-by-step instructions, see Create an app in Azure Active Directory.

During the completion of this step (by following the step-by-step instructions), you'll save the following information to a text file. These values will be used in later steps in the deployment process.

- AAD application ID
- AAD application secret
- Tenant Id

Step 2: Deploy connector web service from GitHub repository to your Azure account

The next step is to deploy the source code for the Twitter connector app that will use Twitter API to connect to your Twitter account and extract data so you can import it to Microsoft 365. The Twitter connector that you deploy for your organization will upload the items from your organization's Twitter account to the Azure Storage location that is created in this step. After you create a Twitter connector in the Microsoft 365 compliance center (in Step 5), the Microsoft 365 Import service will copy the Twitter data from the Azure Storage location to a mailbox in Microsoft 365. As previous explained in the Before you set up a connector section, you must have a valid Azure subscription to create an Azure Storage account.

To deploy the source code for the Twitter connector app:

1. Go to this GitHub site.
2. Click **Deploy to Azure**.

For step-by-step instructions, see Deploy the connector web service from GitHub to your Azure account.

While you follow the step-by-step instructions to complete this step, you provide the following information

- APISecretKey: You create this secret during the completion of this step. It's used in Step 5.
- tenantId: The tenant ID of your Microsoft 365 organization that you copied after creating the Twitter app in Azure Active Directory in Step 1.

After completing this step, be sure to copy the app Service URL (for example, `https://twitterconnector.azurewebsites.net`). You need to use this URL to complete Step 3, Step 4, and Step 5).

Step 3: Create developer app on Twitter

The next step is to create and configure a developer app on Twitter. The custom connector that you create in Step 7 uses the Twitter app to interact with the Twitter API to obtain data from your organization's Twitter account.

For step-by-step instructions, see Create the Twitter app.

During the completion of this step (by following the step-by-step instructions), you save the following information to a text file. These values will be used to configure the Twitter connector app in Step 4.

- Twitter API Key
- Twitter API Secret Key
- Twitter Access Token

- Twitter Access Token Secret

Step 4: Configure the Twitter connector app

The next step is to add configurations settings to the Twitter connector app that you deployed in Step 2. You do this by going to the home page of your connector app and configuring it.

For step-by-step instructions, see Configure the connector web app.

During the completion of this step (by following the step-by-step instructions), you'll provide the following information (that you've copied to a text file after completing the previous steps):

- Twitter API Key (obtained in Step 3)
- Twitter API Secret Key (obtained in Step 3)
- Twitter Access Token (obtained in Step 3)
- Twitter Access Token Secret (obtained in Step 3)
- Azure Active Directory application ID (the AAD application ID obtained in Step 1)
- Azure Active Directory application secret (the AAD application secret obtained in Step 1)

Step 5: Set up a Twitter connector in the Microsoft 365 compliance center

The final step is to set up the Twitter connector in the Microsoft 365 compliance center that will import data from your organization's Twitter account to a specified mailbox in Microsoft 365. After you complete this step, the Microsoft 365 Import service will start importing data from your organization's Twitter account to Microsoft 365.

For step-by-step instructions, see Set up a Twitter connector in the Microsoft 365 compliance center.

During the completion of this step (by following the step-by-step instructions), you'll provide the following information (that you've copied to a text file after completing the steps).

- Azure app service URL (obtained in Step 2; for example, `https://twitterconnector.azurewebsites.net`)
- APISecretKey (that you created in Step 2)

Set up a connector to archive CellTrust data

2/18/2021 • 4 minutes to read • Edit Online

Use a Globanet connector in the Microsoft 365 compliance center to import and archive data from the CellTrust platform to user mailboxes in your Microsoft 365 organization. Globanet provides a CellTrust connector that captures items from the third-party data source and imports those items to Microsoft 365. The connector converts the content of SMS messages from CellTrust accounts to an email message format and then imports those items to the user's mailbox in Microsoft 365.

After CellTrust data is stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, eDiscovery, retention policies and retention labels, and communication compliance. Using a CellTrust connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving CellTrust data

The following overview explains the process of using a connector to archive CellTrust data in Microsoft 365.

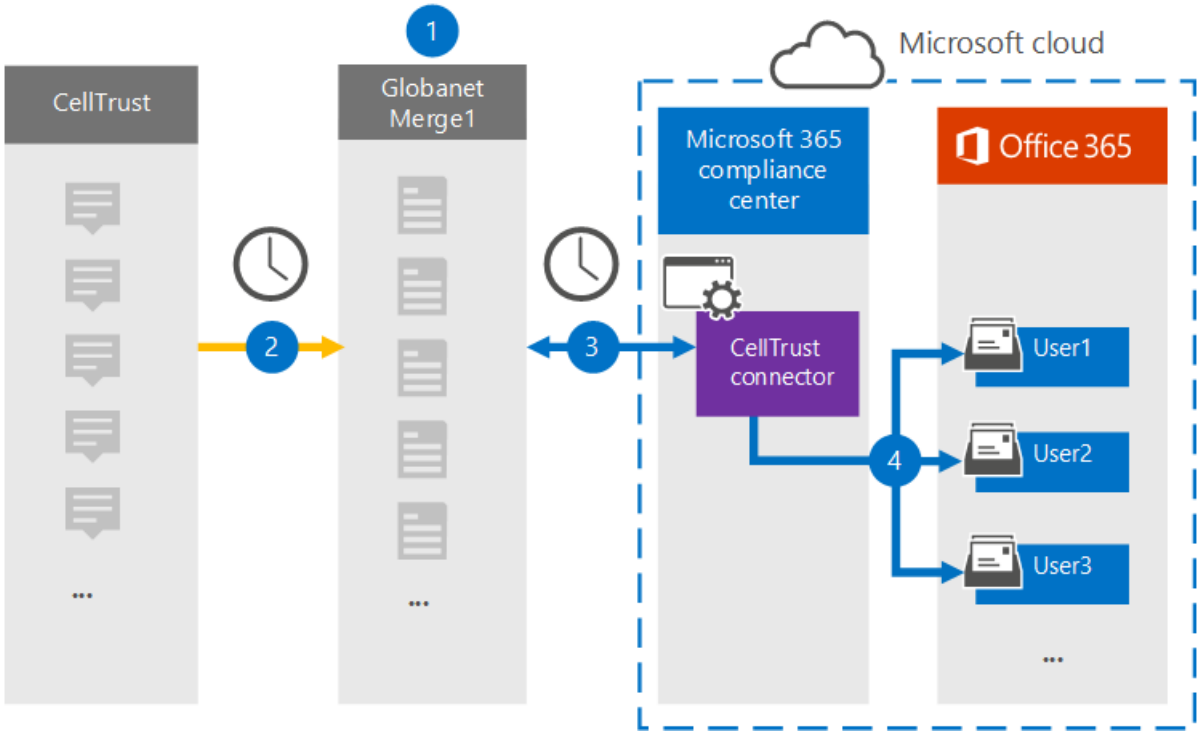

1. Your organization works with CellTrust to set up and configure a CellTrust site.
2. Once every 24 hours, CellTrust items are copied to the Globanet Merge1 site. The connector also converts the content of a message to an email message format.
3. The CellTrust connector that you create in the Microsoft 365 compliance center connects to the Globanet Merge1 site every day and transfers the messages to a secure Azure Storage location in the Microsoft cloud.
4. The automatic user mapping as connector imports items to the mailboxes of specific users by using the value of the *Email* property of the described in Step 3. A subfolder in the Inbox folder named **CellTrust** is created in the user mailboxes, and the message items are imported to that folder. The connector determines which mailbox to import items to by using the value of the *Email* property. Every CellTrust item contains this property, which is populated with the email address of every participant.

Before you begin

- Create a Merge1 account for Microsoft connectors. To create an account, contact Globanet Customer Support. You need to sign into this account when you create the connector in Step 1.
- The user who creates the CellTrust connector in Step 1 (and completes it in Step 3) must be assigned to the Mailbox Import Export role in Exchange Online. This role is required to add connectors on the **Data connectors** page in the Microsoft 365 compliance center. By default, this role isn't assigned to any role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Step 1: Set up the CellTrust connector

The first step is to access to the **Data Connectors** in the Microsoft 365 compliance center and create a connector for CellTrust data.

1. Go to https://compliance.microsoft.com and then click **Data connectors** > **CellTrust**.
2. On the **CellTrust** product description page, click **Add connector**.
3. On the **Terms of service** page, click **Accept**.
4. Enter a unique name that identifies the connector and then click **Next**.
5. Sign in to your Merge1 account to configure the connector.

Step 2: Configure the CellTrust connector on the Globanet Merge1 site

The second step is to configure the CellTrust connector on the Globanet Merge1 site. For information about how to configure the CellTrust connector, see Merge1 Third-Party Connectors User Guide.

After you click **Save & Finish**, the **User mapping** page in the connector wizard in the Microsoft 365 compliance center is displayed.

Step 3: Map users and complete the connector setup

To map users and complete the connector set up in the Microsoft 365 compliance center, follow these steps:

1. On the **Map CellTrust users to Microsoft 365 users** page, enable automatic user mapping. The CellTrust items include a property called *Email*, which contains email addresses for users in your organization. If the connector can associate this address with a Microsoft 365 user, the items are imported to that user's mailbox.
2. Click **Next**, review your settings, and go to the **Data connectors** page to see the progress of the import process for the new connector.

Step 4: Monitor the CellTrust connector

After you create the CellTrust connector, you can view the connector status in the Microsoft 365 compliance center.

1. Go to https://compliance.microsoft.com and click **Data connectors** in the left nav.
2. Click the **Connectors** tab and then select the **CellTrust** connector to display the flyout page, which

contains the properties and information about the connector.

3. Under **Connector status with source**, click the **Download log** link to open (or save) the status log for the connector. This log contains data that has been imported to the Microsoft cloud.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.

Set up a connector to archive Cisco Jabber data

2/18/2021 • 4 minutes to read • Edit Online

Use a Globanet connector in the Microsoft 365 compliance center to import and archive data from the Cisco Jabber platform to user mailboxes in your Microsoft 365 organization. Globanet provides you with a Cisco Jabber connector that is configured to capture items from the Jabber's MS SQL Database, such as 1:1 chat messages and group chats and then import those items to Microsoft 365. The connector retrieves data from the Cisco Jabber's MS SQL Database, processes it, and the converts the content from a user's Cisco Jabber account to an email message format and then imports those items to the user's mailbox in Microsoft 365.

After Cisco Jabber data is stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, eDiscovery, retention policies and retention labels, and communication compliance. Using a Cisco Jabber connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving Cisco Jabber data

The following overview explains the process of using a connector to archive Cisco Jabber data in Microsoft 365.

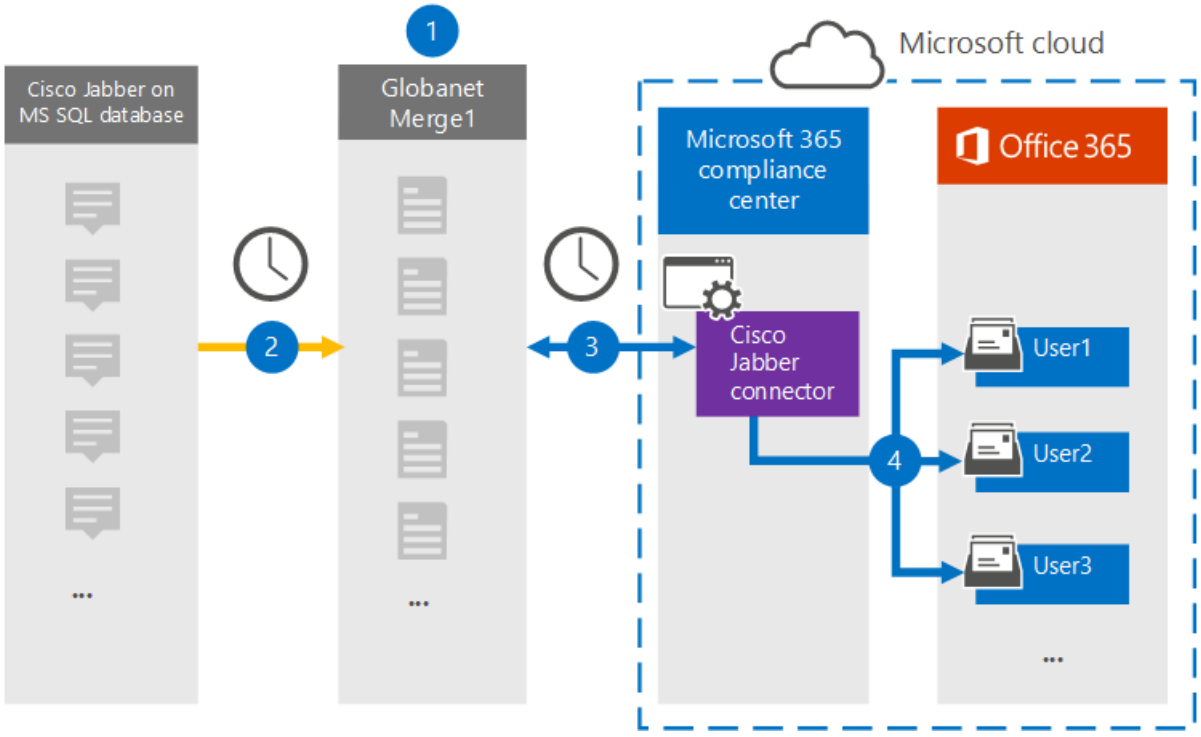

1. Your organization works with Cisco to set up and configure a Cisco Jabber on MS SQL Database.
2. Once every 24 hours, Cisco Jabber items are copied from the MS SQL Database to the Globanet Merge1 site. The connector also converts the content of chat messages to an email message format.
3. The Cisco Jabber connector that you create in the Microsoft 365 compliance center connects to the Globanet Merge1 site every day and transfers the items to a secure Azure Storage location in the Microsoft cloud.
4. The automatic user mapping as connector imports items to the mailboxes of specific users by using the value of the *Email* property of the described in Step 3. A subfolder in the Inbox folder named **Cisco Jabber on MS SQL** is created in the user mailboxes, and the message items are imported to that folder. The connector determines which mailbox to import items to by using the value of the *Email* property.

Every Cisco Jabber item contains this property, which is populated with the email address of every participant.

Before you begin

- Create a Globanet Merge1 account for Microsoft connectors. To create this account, contact Globanet Customer Support. You will sign into this account when you create the connector in Step 1.
- Set up an MS SQL Database to retrieve Jabber items from before creating the connector in Step 1. You will specify the connection settings for the MS SQL Database when configuring the Cisco Jabber connector in Step 2. For more information, see the Merge1 Third-Party Connectors User Guide.
- The user who creates the Cisco Jabber connector in Step 1 (and completes it in Step 3) must be assigned to the Mailbox Import Export role in Exchange Online. This role is required to add connectors on the **Data connectors** page in the Microsoft 365 compliance center. By default, this role is not assigned to a role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Step 1: Set up the Cisco Jabber connector

The first step is to access to the **Data Connectors** in the Microsoft 365 compliance center and create a connector for Cisco Jabber on MS SQL data.

1. Go to https://compliance.microsoft.comand then click **Data connectors** > **Cisco Jabber on MS SQL**.
2. On the **Cisco Jabber on MS SQL** product description page, click **Add connector**.
3. On the **Terms of service** page, click **Accept**.
4. Enter a unique name that identifies the connector and then click **Next**.
5. Sign in to your Merge1 account to configure the connector.

Step 2: Configure the Cisco Jabber connector on the Globanet Merge1 site

The second step is to configure the Cisco Jabber on MS SQL connector on the Globanet Merge1 site. For information about how to configure the Cisco Jabber on MS SQL connector, see Merge1 Third-Party Connectors User Guide.

After you click **Save & Finish**, the **User mapping** page in the connector wizard in the Microsoft 365 compliance center is displayed.

Step 3: Map users and complete the connector setup

To map users and complete the connector set up in the Microsoft 365 compliance center, follow these steps:

1. On the **Map Cisco Jabber on MS SQL users to Microsoft 365 users** page, enable automatic user mapping. The Cisco Jabber on MS SQL items include a property called *Email*, which contains email addresses for users in your organization. If the connector can associate this address with a Microsoft 365 user, the items are imported to that user's mailbox.
2. Click **Next**, review your settings, and go to the **Data connectors** page to see the progress of the import process for the new connector.

Step 4: Monitor the Cisco Jabber connector

After you create the Cisco Jabber on MS SQL connector, you can view the connector status in the Microsoft 365 compliance center.

1. Go to https://compliance.microsoft.com and click **Data connectors** in the left nav.
2. Click the **Connectors** tab and then select the **Cisco Jabber on MS SQL** connector to display the flyout page. This page contains the properties and information about the connector.
3. Under **Connector status with source**, click the **Download log** link to open (or save) the status log for the connector. This log contains data that has been imported to the Microsoft cloud.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.

Set up a connector to archive EML data

2/18/2021 • 4 minutes to read • Edit Online

Use a Globanet connector in the Microsoft 365 compliance center to import and archive EML data to user mailboxes in your Microsoft 365 organization. EML is the file extension for an email message saved to a file. The connector converts the content of an item from the source format to an email message format and then imports the item to a user mailbox.

After EML messages are stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, eDiscovery, and retention policies and retention labels. Using an EML connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving EML data

The following overview explains the process of using a connector to archive EML data in Microsoft 365.

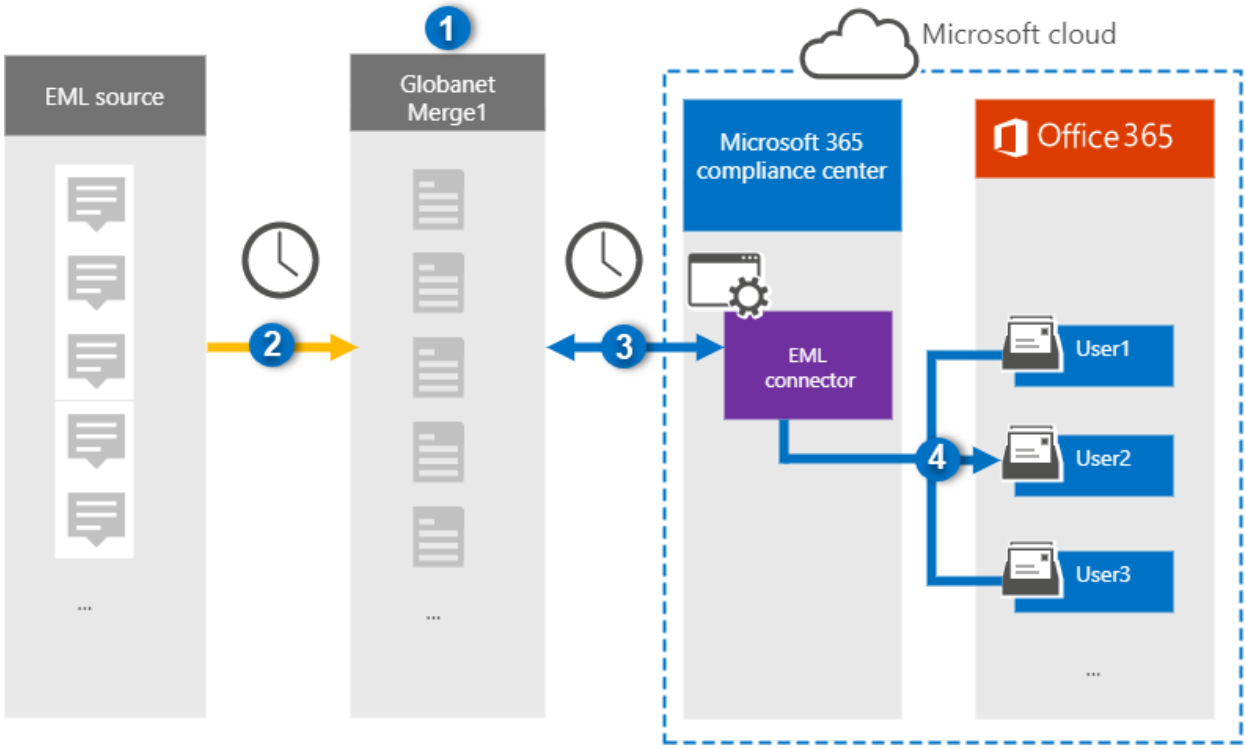

1. Your organization works with the EML source to set up and configure an EML site.
2. Once every 24 hours, content items from the EML source are copied to the Globanet Merge1 site. During this process, the content of an EML file is converted to an email message format.
3. The EML connector that you create in the Microsoft 365 compliance center, connects to the Globanet Merge1 site every day and transfers the messages to a secure Azure Storage location in the Microsoft cloud.
4. The connector imports the converted message items to the mailboxes of specific users using the value of the *Email* property of the automatic user mapping process that's described in Step 3. During this process, a subfolder in the Inbox folder named **EML** is created in the user mailboxes, and the EML items are imported to that folder. The connector determines which mailbox to import items to by using the value of the *Email* property. Every message contains this property, which is populated with the email address of

every participant of the content item.

Before you begin

- Create a Globanet Merge1 account for Microsoft connectors. To create an account, contact Globanet Customer Support. You will sign into this account when you create the connector in Step 1.
- The user who creates the EML connector in Step 1 (and completes it in Step 3) must be assigned to the Mailbox Import Export role in Exchange Online. This role is required to add connectors on the **Data connectors** page in the Microsoft 365 compliance center. By default, this role is not assigned to a role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Step 1: Set up an EML Connector

The first step is to access to the **Data Connectors** page in the Microsoft 365 compliance center and create a connector for EML data.

1. Go to https://compliance.microsoft.com and then click **Data connectors** > **EML**.
2. On the **EML** product description page, click **Add connector**.
3. On the **Terms of service** page, click **Accept**.
4. Enter a unique name that identifies the connector, and then click **Next**.
5. Sign in to your Merge1 account to configure the connector.

Step 2: Configure the EML connector on the Globanet Merge1 site

The second step is to configure the EML connector on the Globanet Merge1 site. For information about configuring the EML connector, see Merge1 Third-Party Connectors User Guide.

After you click **Save & Finish**, the **User mapping** page in the connector wizard in the Microsoft 365 compliance center is displayed.

Step 3: Map users and complete the connector setup

To map users and complete the connector setup in the Microsoft 365 compliance center, follow these steps:

1. On the **Map external users to Microsoft 365 users** page, enable automatic user mapping. The EML source items include a property called *Email*, which contains email addresses for users in your organization. If the connector can associate this address with a Microsoft 365 user, the EML items are imported to that user's mailbox.
2. Click **Next**, review your settings, and then go to the **Data connectors** page to see the progress of the import process for the new connector.

Step 4: Monitor the EML connector

After you create the EML connector, you can view the connector status in the Microsoft 365 compliance center.

1. Go to https://compliance.microsoft.com and click **Data connectors** in the left nav.
2. Click the **Connectors** tab and then select the **EML** connector to display the flyout page. This page contains the properties and information about the connector.

3. Under **Connector status with source**, click the **Download log** link to open (or save) the status log for the connector. This log contains information about the data that has been imported to the Microsoft cloud.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.

Set up a connector to archive FX Connect data

2/18/2021 • 4 minutes to read • Edit Online

Use a Globanet connector in the Microsoft 365 compliance center to import and archive data from the FX Connect collaboration platform to user mailboxes in your Microsoft 365 organization. Globanet provides an FX Connect connector that is configured to capture FX Connect items and import those items to Microsoft 365. The connector converts the content from FX Connect, such as trades, messages, and other details from your organization's FX Connect account, to an email message format and then imports those items to the user's mailbox in Microsoft 365.

After FX Connect data is stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, eDiscovery, retention policies and retention labels, and communication compliance. Using an FX Connect connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving FX Connect data

The following overview explains the process of using a connector to archive the FX Connect information in Microsoft 365.

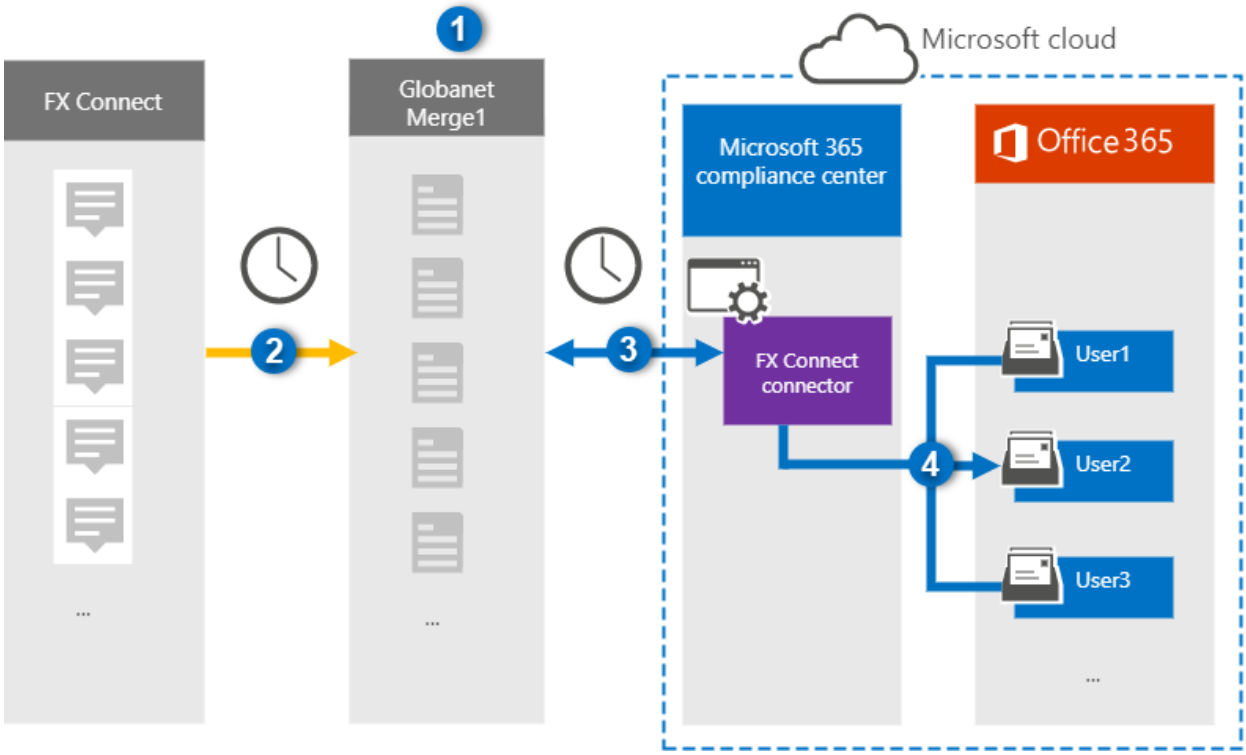

1. Your organization works with FX Connect to set up and configure an FX Connect site.
2. Once every 24 hours, items from FX Connect accounts are copied to the Globanet Merge1 site. The connector also converts the FX Connect items to an email message format.
3. The FX Connect connector that you create in the Microsoft 365 compliance center, connects to the Globanet Merge1 site every day and transfers the FX Connect items to a secure Azure Storage location in the Microsoft cloud.
4. The connector imports items to the mailboxes of specific users by using the value of the *Email* property of the automatic user mapping as described in Step 3. A subfolder in the Inbox folder named **FX Connect**

is created in the user mailboxes, and the items are imported to that folder. The connector does this by using the value of the *Email* property. Every FX Connect item contains this property, which is populated with the email address of every participant of the item.

Before you begin

- Create a Globanet Merge1 account for Microsoft connectors. To create an account, contact Globanet Customer Support. You will sign into this account when you create the connector in Step 1.
- The user who creates the FX Connect connector in Step 1 (and completes it in Step 3) must be assigned to the Mailbox Import Export role in Exchange Online. This role is required to add connectors on the **Data connectors** page in the Microsoft 365 compliance center. By default, this role is not assigned to a role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Step 1: Set up the FX Connect connector

The first step is to access to the **Data Connectors** page in the Microsoft 365 compliance center and create a connector for FX Connect data.

1. Go to https://compliance.microsoft.com and then click **Data connectors** > **FX Connect**.
2. On the **FX Connect** product description page, click **Add connector**.
3. On the **Terms of service** page, click **Accept**.
4. Enter a unique name that identifies the connector, and then click **Next**.
5. Sign in to your Merge1 account to configure the connector.

Step 2: Configure the FX Connect connector on the Globanet Merge1 site

The second step is to configure the FX Connect connector on the Merge1 site. For information about how to configure the FX Connect connector, see Merge1 Third-Party Connectors User Guide.

After you click **Save & Finish**, the **User mapping** page in the connector wizard in the Microsoft 365 compliance center is displayed.

Step 3: Map users and complete the connector setup

To map users and complete the connector setup in the Microsoft 365 compliance center, follow these steps:

1. On the **Map FX Connect users to Microsoft 365 users** page, enable automatic user mapping. The FX Connect items include a property called *Email*, which contains email addresses for users in your organization. If the connector can associate this address with a Microsoft 365 user, the items are imported to that user's mailbox.
2. Click **Next**, review your settings, and then go to the **Data connectors** page to see the progress of the import process for the new connector.

Step 4: Monitor the FX Connect connector

After you create the FX Connect connector, you can view the connector status in the Microsoft 365 compliance center.

1. Go to https://compliance.microsoft.com/ and click **Data connectors** in the left nav.
2. Click the **Connectors** tab and then select the **FX Connect** connector to display the flyout page. This page contains the properties and information about the connector.
3. Under **Connector status with source**, click the **Download log** link to open (or save) the status log for the connector. This log contains data that has been imported to the Microsoft cloud.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.

Set up a connector to archive Jive data

2/18/2021 • 4 minutes to read • Edit Online

Use a Globanet connector in the Microsoft 365 compliance center to import and archive data from the collaboration platform to user mailboxes in your Microsoft 365 organization. Globanet provides a Jive connector that is configured to capture items from the third-party data source (on a regular basis) and then import those items to Microsoft 365. The connector converts content such as email messages, chats, and attachments from a user's Jive account to an email message format and then imports those items to the user's mailbox in Microsoft 365.

After Jive data is stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, eDiscovery, retention policies and retention labels, and communication compliance. Using a Jive connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving Jive data

The following overview explains the process of using a connector to archive the Jive data in Microsoft 365.

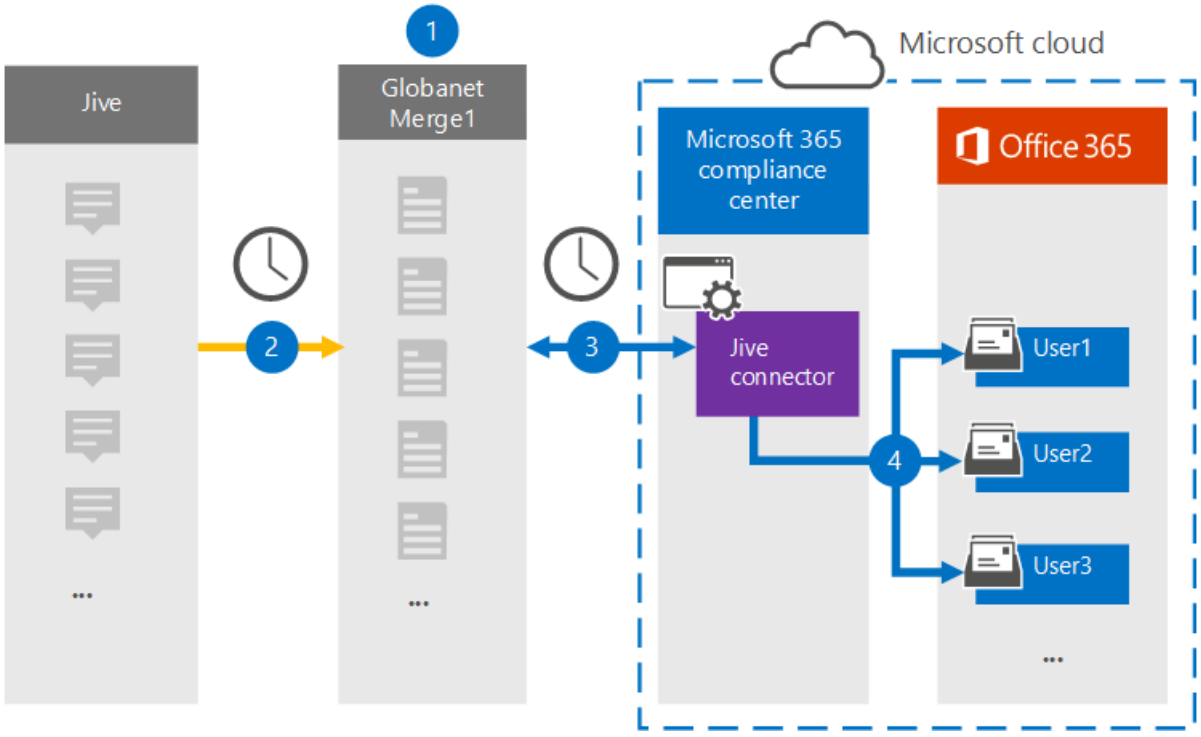

1. Your organization works with Jive to set up and configure a Jive site.
2. Once every 24 hours, items from Jive are copied to the Globanet Merge1 site. The connector also converts the content of Jive items to an email message format.
3. The Jive connector that you create in the Microsoft 365 compliance center connects to the Globanet Merge1 site every day and transfers the content to a secure Azure Storage location in the Microsoft cloud.
4. The connector imports the converted items to the mailboxes of specific users by using the value of the *Email* property of the automatic user mapping as described in Step 3. A new subfolder in the Inbox folder named **Jive** is created in the user mailboxes, and the items are imported to that folder. The connector does this by using the value of the *Email* property. Every Jive item contains this property, which is populated with the email address of every participant of the item.

Before you begin

- Create a Globanet Merge1 account for Microsoft connectors. To create this account, contact globanet customer support. You will sign into this account when you create the connector in Step 1.
- The user who creates the Jive connector in Step 1 (and completes it in Step 3) must be assigned to the Mailbox Import Export role in Exchange Online. This role is required to add connectors on the **Data connectors** page in the Microsoft 365 compliance center. By default, this role is not assigned to a role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Step 1: Set up the Jive connector

The first step is to access to the **Data Connectors** page in the Microsoft 365 compliance center and create a connector for Jive data.

1. Go to https://compliance.microsoft.com and then click **Data connectors** > **Jive**.
2. On the **Jive** product description page, click **Add connector**.
3. On the **Terms of service** page, click **Accept**.
4. Enter a unique name that identifies the connector, and then click **Next**.
5. Sign in to your Merge1 account to configure the connector.

Step 2: Configure the Jive connector

The second step is to configure the Jive connector on the Merge1 site. For information about how to configure the Jive connector, see Merge1 Third-Party Connectors User Guide.

After you click **Save & Finish**, the **User mapping** page in the connector wizard in the Microsoft 365 compliance center is displayed.

Step 3: Map users and complete the connector setup

To map users and complete the connector setup in the Microsoft 365 compliance center, follow the steps below:

1. On the **Map Jive users to Microsoft 365 users** page, enable automatic user mapping. The Jive items include a property called *Email*, which contains email addresses for users in your organization. If the connector can associate this address with a Microsoft 365 user, the items are imported to that user's mailbox.
2. Click **Next**, review your settings, and go to the **Data connectors** page to see the progress of the import process for the new connector.

Step 4: Monitor the Jive connector

After you create the Jive connector, you can view the connector status in the Microsoft 365 compliance center.

1. Go to https://compliance.microsoft.com and click **Data connectors** in the left nav.
2. Click the **Connectors** tab and then select the **Jive** connector to display the flyout page. This page contains the properties and information about the connector.
3. Under **Connector status with source**, click the **Download log** link to open (or save) the status log for

the connector. This log contains information about the data that has been imported to the Microsoft cloud.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.

Set up a connector to archive data from MS SQL Database

2/18/2021 • 4 minutes to read • Edit Online

Use a Globanet connector in the Microsoft 365 compliance center to import and archive data from MS SQL Database to user mailboxes in your Microsoft 365 organization. Globanet provides you with an MS SQL Database Importer connector that's configured to capture items from a database using an XML configuration file and import those items to Microsoft 365. The connector converts content from MS SQL Database to an email message format and then imports those items to user mailboxes in Microsoft 365.

After content from MS SQL Database stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, eDiscovery, retention policies and retention labels. Using an MS SQL Database connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving the MS SQL data

The following overview explains the process of using a connector to archive MS SQL data in Microsoft 365.

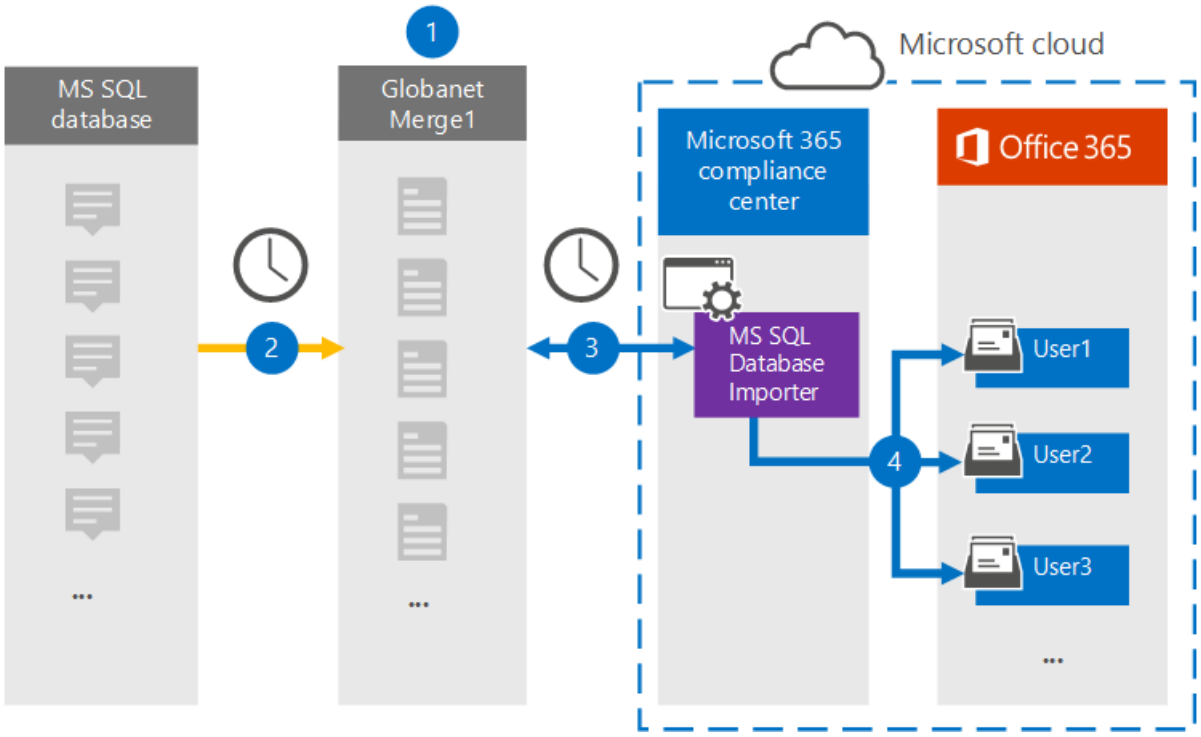

1. Your organization works with an MS SQL Database provider to set up and configure an MS SQL Database site.
2. Once every 24 hours, MS SQL Database items are copied to the Globanet Merge1 site. The connector also converts this content to an email message format.
3. The MS SQL Database Importer connector that you create in the Microsoft 365 compliance center, connects to the Globanet Merge1 site every day and transfers the messages to a secure Azure Storage location in the Microsoft cloud.
4. The connector imports the converted MS SQL Database items to the mailboxes of specific users using the value of the *Email* property of the automatic user mapping as described in Step 3. A subfolder in the

Inbox folder named **MS SQL Database Importer** is created in the user mailboxes, and the items are imported to that folder. The connector determines which mailbox to import items to by using the value of the *Email* property. Every item from the MS SQL Database contains this property, which is populated with the email address of every participant of the item.

Before you begin

- Create a Globanet Merge1 account for Microsoft connectors. To create an account, contact Globanet Customer Support. You need to sign into this account when you create the connector in Step 1.
- The user who creates the MS SQL Database Importer connector in Step 1 (and completes it in Step 3) must be assigned to the Mailbox Import Export role in Exchange Online. This role is required to add connectors on the Data connectors page in the Microsoft 365 compliance center. By default, this role is not assigned to any role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Step 1: Set up the MS SQL Database Importer connector

The first step is to access to the **Data Connectors** page in the Microsoft365 compliance center and create a connector for the MS SQL Database.

1. Go to https://compliance.microsoft.com and then click **Data connectors** > **MS SQL Database Importer**.
2. On the **MS SQL Database Importer** product description page, click **Add new connector**.
3. On the **Terms of service** page, click **Accept**.
4. Enter a unique name that identifies the connector, and then click **Next**.
5. Sign in to your Merge1 account to configure the connector.

Step 2: Configure the MS SQL Database Importer connector on the Globanet Merge1 site

The second step is to configure the MS SQL Database Importer connector on the Merge1 site. For information about how to configure the MS SQL Database Importer, see Merge1 Third-Party Connectors User Guide.

After you click **Save & Finish**, the **User mapping** page in the connector wizard in the Microsoft 365 compliance center is displayed.

Step 3: Map users and complete the connector setup

To map users and complete the connector setup, follow these steps:

1. On the **Map MS SQL Database Importer users to Microsoft 365 users** page, enable automatic user mapping. The MS SQL Database items include a property called *Email*, which contains email addresses for users in your organization. If the connector can associate this address with a Microsoft 365 user, the items are imported to that user's mailbox.
2. Click **Next**, review your settings, and go to the **Data connectors** page to see the progress of the import process for the new connector.

Step 4: Monitor the MS SQL Database Importer connector

After you create the MS SQL Database Importer connector, you can view the connector status in the Microsoft 365 compliance center.

1. Go to https://compliance.microsoft.com/ and click **Data connectors** in the left nav.
2. Click the **Connectors** tab and then select the **MS SQL Database Importer** connector to display the flyout page, which contains the properties and information about the connector.
3. Under **Connector status with source**, click the **Download log** link to open (or save) the status log for the connector. This log contains data that has been imported to the Microsoft cloud.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.

Set up a connector to archive Pivot data

2/18/2021 • 4 minutes to read • Edit Online

Use a Globanet connector in the Microsoft 365 compliance center to import and archive data from the Pivot platform to user mailboxes in your Microsoft 365 organization. Globanet provides you with a Pivot connector that is configured to capture items from the third-party data source (on a regular basis) and then import those items to Microsoft 365. Pivot is an instant messaging platform that allows collaboration with financial market participants. The connector converts items such as chat messages, from a users' Pivot accounts to an email message format and then imports those items to the user mailboxes in Microsoft 365.

After Pivot data is stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, eDiscovery, retention policies and retention labels, and communication compliance. Using a Pivot connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving Pivot data

The following overview explains the process of using a connector to archive the Pivot data in Microsoft 365.

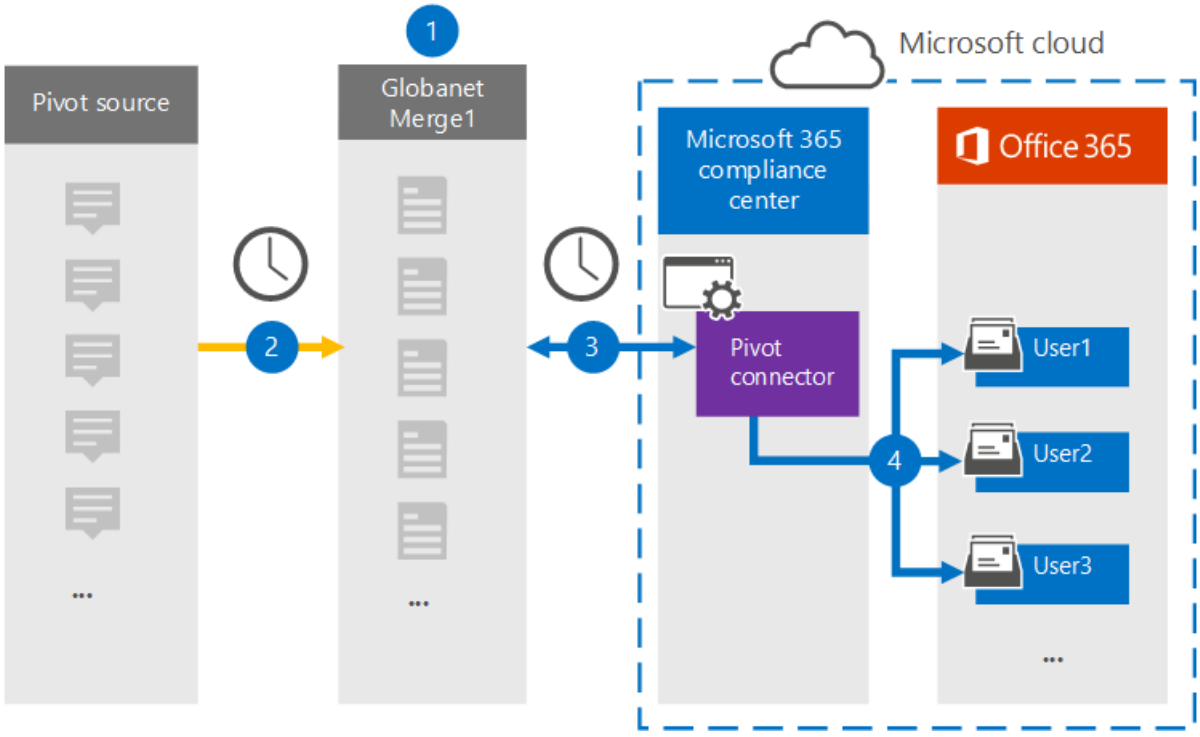

1. Your organization works with Pivot to set up and configure a Pivot source site.
2. Once every 24 hours, Pivot items are copied to the Globanet Merge1 site. The connector also converts the Pivot items to an email message format.
3. The Pivot connector that you create in the Microsoft 365 compliance center, connects to the Globanet Merge1 site every day and transfers the Pivot items to a secure Azure Storage location in the Microsoft cloud.
4. The connector imports the Pivot items to the mailboxes of specific users by using the value of the *Email* property of the automatic user mapping as described in Step 3. A subfolder in the Inbox folder named **Pivot** is created in the user mailboxes, and the items are imported to that folder. The connector does this by using the value of the *Email* property. Every Pivot item contains this property, which is populated with

the email address of every participant of the item.

Before you begin

- Create a Globanet Merge1 account for Microsoft connectors. To create this account, contact Globanet Customer Support. You will sign into this account when you create the connector in Step 1.
- The user who creates the Pivot connector in Step 1 (and completes it in Step 3) must be assigned to the Mailbox Import Export role in Exchange Online. This role is required to add connectors on the Data connectors page in the Microsoft 365 compliance center. By default, this role is not assigned to a role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Step 1: Set up the Pivot connector

The first step is to access to the **Data Connectors** page in the Microsoft compliance center and create a connector for Pivot data.

1. Go to https://compliance.microsoft.com and then click **Data connectors** > **Pivot**.
2. On the **Pivot** product description page, click **Add connector**.
3. On the **Terms of service** page, click **Accept**.
4. Enter a unique name that identifies the connector and then click **Next**.
5. Sign in to your Merge1 account to configure the connector.

Step 2: Configure the Pivot connector on the Globanet Merge1 site

The second step is to configure the Pivot connector on the Merge1 site. For information about how to configure the Pivot connector on the Globanet Merge1 site, see Merge1 Third-Party Connectors User Guide.

After you click **Save & Finish**, the **User mapping** page in the connector wizard in the Microsoft 365 compliance center is displayed.

Step 3: Map users and complete the connector setup

To map users and complete the connector setup in the Microsoft 356 compliance center, follow these steps:

1. On the **Map Pivot users to Microsoft 365 users** page, enable automatic user mapping. The Pivot items include a property called *Email*, which contains email addresses for users in your organization. If the connector can associate this address with a Microsoft 365 user, the items are imported to that user's mailbox.
2. Click **Next**, review your settings, and go to the **Data connectors** page to see the progress of the import process for the new connector.

Step 4: Monitor the Pivot connector

After you create the Pivot connector, you can view the connector status in the Microsoft 365 compliance center.

1. Go to https://compliance.microsoft.com and click **Data connectors** in the left nav.
2. Click the **Connectors** tab and then select the **Pivot** connector to display the flyout page. This page contains the properties and information about the connector.

3. Under **Connector status with source**, click the **Download log** link to open (or save) the status log for the connector. This log contains data that has been imported to the Microsoft cloud.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.

Set up a connector to archive Redtail Speak data

2/18/2021 • 4 minutes to read • Edit Online

Use a Globanet connector in the Microsoft 365 compliance center to import and archive data from the Redtail Speak to user mailboxes in your Microsoft 365 organization. Globanet provides you with a Redtail Speak connector that's configured to capture items from your organization's SFTP server where the items are received from Redtail. The connector converts the content from Redtail Speak to an email message format and then imports those items to the user's mailbox in Microsoft 365.

After Redtail Speak data is stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, eDiscovery, retention policies, and retention labels. Using a Redtail Speak connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving the Redtail Speak data

The following overview explains the process of using a connector to archive the Redtail Speak data in Microsoft 365.

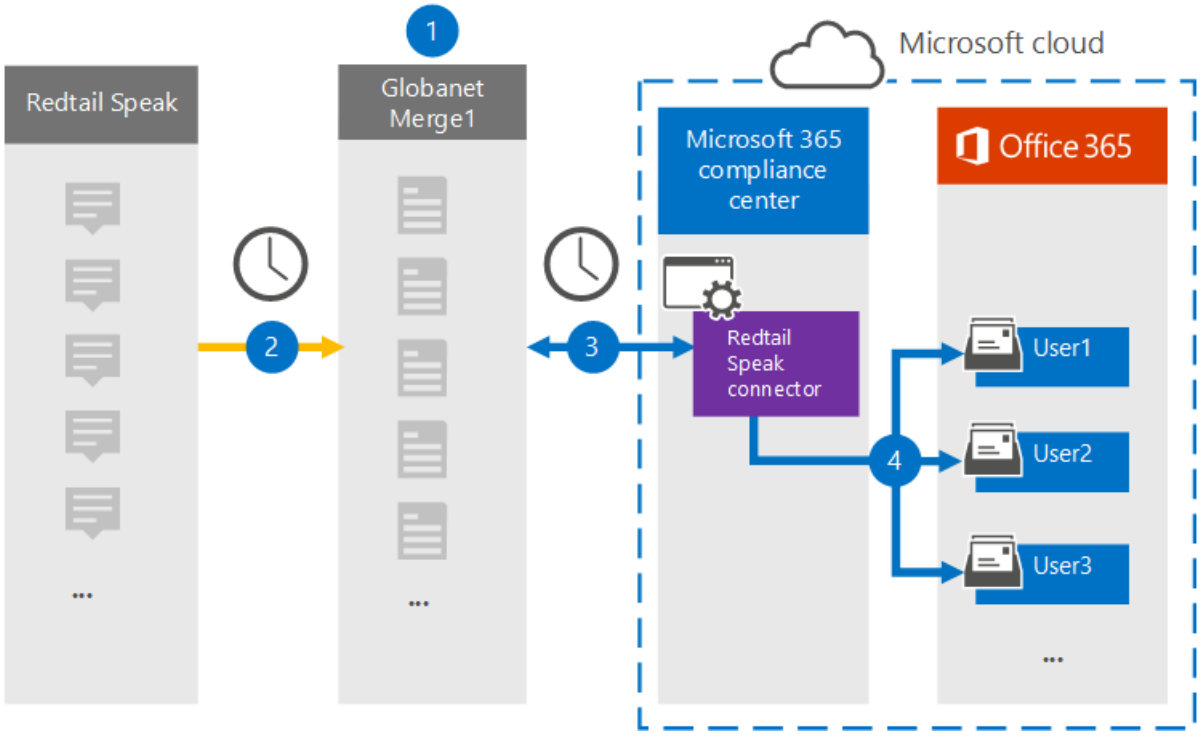

1. Your organization works with Redtail Speak to set up and configure an SMTP gateway where messages are forwarded from Redtail Speak to your organization's SFTP server on a daily basis.

2. Once every 24 hours, the Redtail Speak items are copied to the Globanet Merge1 site. The connector also converts the Redtail Speak items to an email message format.

3. The Redtail Speak connector that you create in the Microsoft 365 compliance center connects to the Globanet Merge1 site every day and transfers the messages to a secure Azure Storage location in the Microsoft cloud.

4. The connector imports the converted Redtail Speak items to the mailboxes of specific users using the value of the *Email* property of the automatic user mapping as described in Step 3. A subfolder in the Inbox folder named **Redtail Speak** is created in the user mailboxes, and the items are imported to that

folder. The connector determines which mailbox to import items to by using the value of the *Email* property. Every Redtail Speak item contains this property, which is populated with the email address of every participant of the item.

Before you begin

- Create a Globanet Merge1 account for Microsoft connectors. To create an account, contact Globanet Customer Support. You need to sign into this account when you create the connector in Step 1.
- In Step 2, you need to specify your organization's SFTP server. This step is necessary so that Globanet Merge1 can contact it to collect Redtail Speak data via SFTP.
- The user who creates the Redtail Speak Importer connector in Step 1 (and completes it in Step 3) must be assigned to the Mailbox Import Export role in Exchange Online. This role is required to add connectors on the Data connectors page in the Microsoft 365 compliance center. This role is not assigned to any role group in Exchange Online by default. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Step 1: Set up the Redtail Speak connector

The first step is to access to the **Data Connectors** page in the Microsoft 365 compliance center and create a connector for the Redtail Speak data.

1. Go to https://compliance.microsoft.com and select **Data connectors** > **Redtail Speak**.
2. On the **Redtail Speak** product description page, select **Add new connector**.
3. On the **Terms of service** page, select **Accept**.
4. Enter a unique name that identifies the connector, and then select **Next**.
5. Sign in to your Merge1 account to configure the connector.

Step 2: Configure the Redtail Speak connector on the Globanet Merge1 site

The second step is to configure the Redtail Speak connector on the Merge1 site. For information about how to configure the Redtail Speak connector, see Merge1 Third-Party Connectors User Guide.

After you select **Save & Finish**, the **User mapping** page in the connector wizard in the Microsoft 365 compliance center is displayed.

Step 3: Map users and complete the connector setup

To map users and complete the connector setup, follow these steps:

1. On the **Map Redtail Speak users to Microsoft 365 users** page, enable automatic user mapping. The Redtail Speak items include a property called *Email*, which contains email addresses for users in your organization. If the connector can associate this address with a Microsoft 365 user, the items are imported to that user's mailbox.
2. Select **Next**, review your settings, and go to the **Data connectors** page to see the progress of the import process for the new connector.

Step 4: Monitor the Redtail Speak connector

After you create the Redtail Speak connector, you can view the connector status in the Microsoft 365 compliance center.

1. Go to https://compliance.microsoft.com and select **Data connectors** in the left nav.
2. Select the **Connectors** tab and then select the **Redtail Speak** connector to display the flyout page. This page displays properties and information about the connector.
3. Under **Connector status with source**, select the **Download log** link to open (or save) the status log for the connector. This log contains data that has been imported to the Microsoft cloud.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.

Set up a connector to archive Reuters Dealing data

2/18/2021 • 4 minutes to read • Edit Online

Use a Globanet connector in the Microsoft 365 compliance center to import and archive data from the Reuters Dealing platform to user mailboxes in your Microsoft 365 organization. Globanet provides you with a Reuters Dealing connector that's configured to capture items from the third-party data source (on a regular basis) and then import those items to Microsoft 365. The connector converts Dealing communications from the Reuters Dealing account to an email message format and then imports those items to the user's mailbox in Microsoft 365.

After Reuters Dealing data is stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, eDiscovery, retention policies and retention labels, and communication compliance. Using a Reuters Dealing connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving Reuters Dealing data

The following overview explains the process of using a connector to archive the Reuters Dealing data in Microsoft 365.

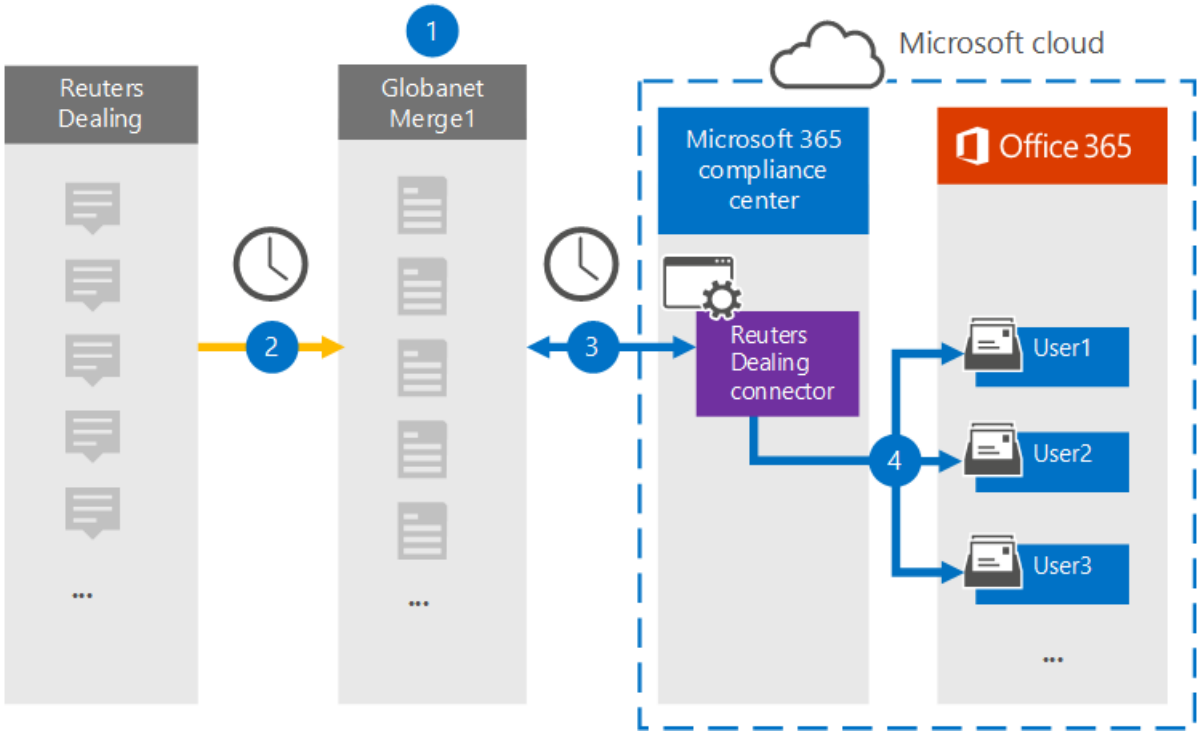

1. Your organization works with Reuters Dealing to set up and configure a Reuters Dealing site.
2. Once every 24 hours, Reuters Dealing items are copied to the Globanet Merge1 site. The connector also converts the items to an email message format.
3. The Reuters Dealing connector that you create in the Microsoft 365 compliance center connects to the Globanet Merge1 site every day and transfers the content to a secure Azure Storage location in the Microsoft cloud.
4. The connector imports items to the mailboxes of specific users by using the value of the *Email* property of the automatic user mapping as described in Step 3. A subfolder in the Inbox folder named **Reuters Dealing** is created in the user mailboxes, and the items are imported to that folder. The connector

determines which mailbox to import items to by using the value of the *Email* property. Every Reuters Dealing item contains this property, which is populated with the email address of every participant of the item.

Before you begin

- Create a Globanet Merge1 account for Microsoft connectors. To create an account, contact Globanet Customer Support. You need to sign into this account when you create the connector in Step 1.
- The user who creates the Reuters Dealing connector in Step 1 (and completes it in Step 3) must be assigned to the Mailbox Import Export role in Exchange Online. This role is required to add connectors on the **Data connectors** page in the Microsoft 365 compliance center. By default, this role is not assigned to any role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Step 1: Set up the Reuters Dealing connector

The first step is to access to the **Data Connectors** page in the Microsoft 365 and create a connector for Reuters Dealing data.

1. Go to https://compliance.microsoft.com and then click **Data connectors** > **Reuters Dealing**.
2. On the **Reuters Dealing** product description page, click **Add connector**.
3. On the **Terms of service** page, click **Accept**.
4. Enter a unique name that identifies the connector, and then click **Next**.
5. Sign to your Merge1 account to configure the connector.

Step 2: Configure the Reuters Dealing connector on the Globanet Merge1 site

The second step is to configure the Reuters Dealing connector on Globanet the Merge1 site. For information about configuring the Reuters Dealing connector, see Merge1 Third-Party Connectors User Guide.

After you click **Save & Finish**, the **User mapping** page in the connector wizard in the Microsoft 365 compliance center is displayed.

Step 3: Map users and complete the connector setup

To map users and complete the connector setup in the Microsoft 365 compliance center, follow these steps:

1. On the **Map Reuters Dealing users to Microsoft 365 users** page, enable automatic user mapping.

   Reuters Dealing items include a property called *Email*, which contains email addresses for users in your organization. If the connector can associate this address with a Microsoft 365 user, the items are imported to that user's mailbox.

2. Click **Next**, review your settings, and go to the **Data connectors** page to see the progress of the import process for the new connector.

Step 4: Monitor the Reuters Dealing connector

After you create the Reuters Dealing connector, you can view the connector status in the Microsoft 365

compliance center.

1. Go to https://compliance.microsoft.com and click **Data connectors** in the left nav.
2. Click the **Connectors** tab and then select the **Reuters Dealing** connector to display the flyout page, which contains the properties and information about the connector.
3. Under **Connector status with source**, click the **Download log** link to open (or save) the status log for the connector. This log contains data that has been imported to the Microsoft cloud.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.

Set up a connector to archive Reuters Eikon data

2/18/2021 • 4 minutes to read • Edit Online

Use a Globanet connector in the Microsoft 365 compliance center to import and archive data from the Reuters Eikon platform to user mailboxes in your Microsoft 365 organization. Globanet provides a Reuters Eikon connector that is configured to capture items from the third-party data source (on a regular basis) and import those items to Microsoft 365. The connector converts the content such as person-to-person messages, group chats, attachments, and disclaimers from a user's Reuters Eikon account to an email message format and then imports those items to the user's mailbox in Microsoft 365.

After Reuters Eikon data is stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, eDiscovery, retention policies and retention labels, and communication compliance. Using a Reuters Eikon connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving Reuters Eikon data

The following overview explains the process of using a connector to archive Reuters Eikon data in Microsoft 365.

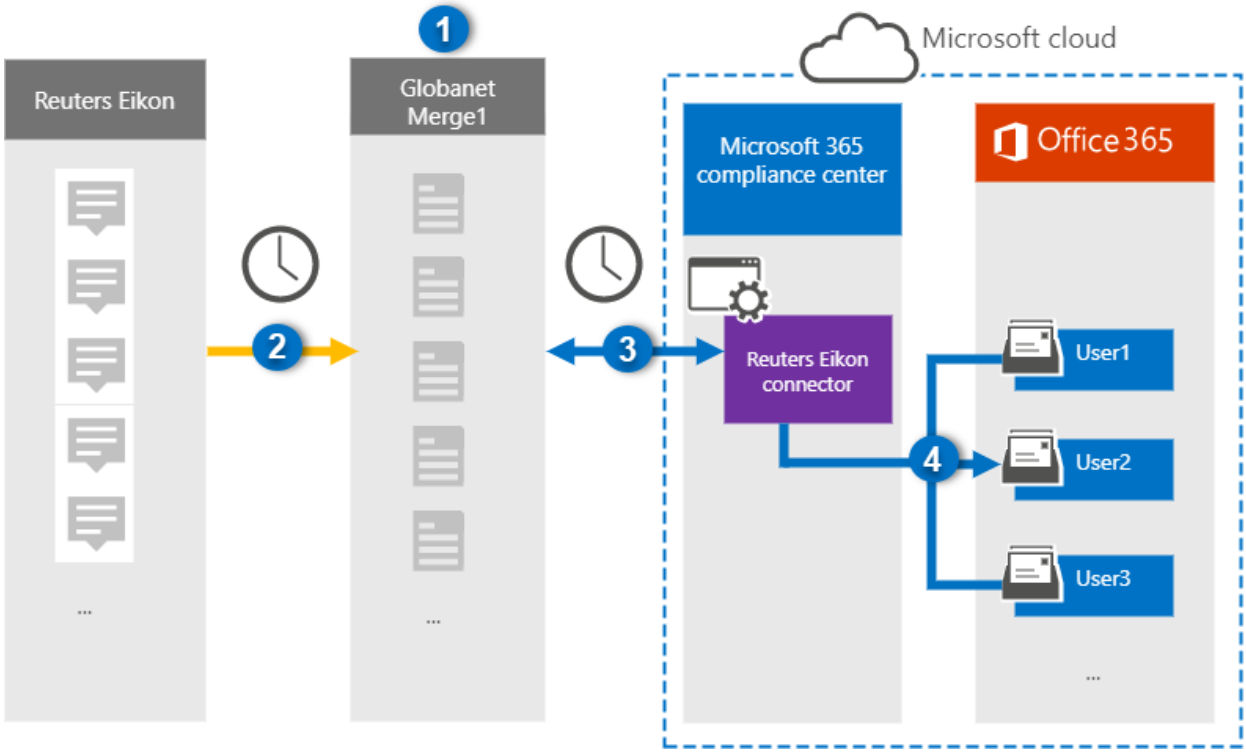

1. Your organization works with Reuters Eikon to set up and configure a Reuters Eikon site.
2. Once every 24 hours, Reuters Eikon items are copied to the Globanet Merge1 site. The connector also converts Reuters Eikon items to an email message format.
3. The Reuters Eikon connector that you create in the Microsoft 365 compliance center connects to the Globanet Merge1 site every day and transfers the content to a secure Azure Storage location in the Microsoft cloud.
4. The connector imports items to the mailboxes of specific users by using the value of the *Email* property of the automatic user mapping as described in Step 3. A subfolder in the Inbox folder named **Reuters Eikon** is created in the user mailboxes, and the items are imported to that folder. The connector

determines which mailbox to import items to by using the value of the *Email* property. Every Reuters Eikon item contains this property, which is populated with the email address of every participant of the item.

Before you begin

- Create a Globanet Merge1 account for Microsoft connectors. To create an account, contact Globanet Customer Support. You will sign into this account when you create the connector in Step 1.
- The user who creates the Reuters Eikon connector in Step 1 (and completes it in Step 3) must be assigned to the Mailbox Import Export role in Exchange Online. This role is required to add connectors on the **Data connectors** page in the Microsoft 365 compliance center. By default, this role is not assigned to a role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Step 1: Set up the Reuters Eikon connector

The first step is to access to the **Data Connectors** page in the Microsoft 365 compliance center and create a connector for Reuters Eikon data.

1. Go to https://compliance.microsoft.com and then click **Data connectors** > **Reuters Eikon**.
2. On the **Reuters Eikon** product description page, click **Add connector**.
3. On the **Terms of service** page, click **Accept**.
4. Enter a unique name that identifies the connector, and then click **Next**.
5. Sign in to your Merge1 account to configure the connector.

Step 2: Configure the Reuters Eikon connector on the Globanet Merge1 site

The second step is to configure the Reuters Eikon connector on the Merge1 site. For information about how to configure the Reuters Eikon connector on the Globanet Merge1 site, see Merge1 Third-Party Connectors User Guide.

After you click **Save & Finish**, the **User mapping** page in the connector wizard in the Microsoft 365 compliance center is displayed.

Step 3: Map users and complete the connector setup

To map users and complete the connector setup in the Microsoft 365 compliance center, follow these steps:

1. On the **Map external users to Microsoft 365 users** page, enable automatic user mapping. The Reuters Eikon items include a property called *Email*, which contains email addresses for users in your organization. If the connector can associate this address with a Microsoft 365 user, the items are imported to that user's mailbox.
2. Click **Next**, review your settings, and then go to the **Data connectors** page to see the progress of the import process for the new connector.

Step 4: Monitor the Reuters Eikon connector

After you create the Reuters Eikon connector, you can view the connector status in the Microsoft 365 compliance

center.

1. Go to https://compliance.microsoft.com and click **Data connectors** in the left nav.
2. Click the **Connectors** tab and then select the **Reuters Eikon** connector to display the flyout page. This page contains the properties and information about the connector.
3. Under **Connector status with source**, click the **Download log** link to open (or save) the status log for the connector. This log contains information about the data that has been imported to the Microsoft cloud.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.

Set up a connector to archive Reuters FX data

2/18/2021 • 4 minutes to read • Edit Online

Use a Globanet connector in the Microsoft 365 compliance center to import and archive data from the Reuters FX platform to user mailboxes in your Microsoft 365 organization. Globanet provides you with a Reuters FX connector that is configured to capture items from the third-party data source (on a regular basis) and then import those items to Microsoft 365. The connector converts the currencies and FX rates from the Reuters FX account to an email message format and then imports those items to the user's mailbox in Microsoft 365.

After Reuters FX data is stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, eDiscovery, retention policies and retention labels, and communication compliance. Using a Reuters FX connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving Reuters FX data

The following overview explains the process of using a connector to archive Reuters FX data in Microsoft 365.

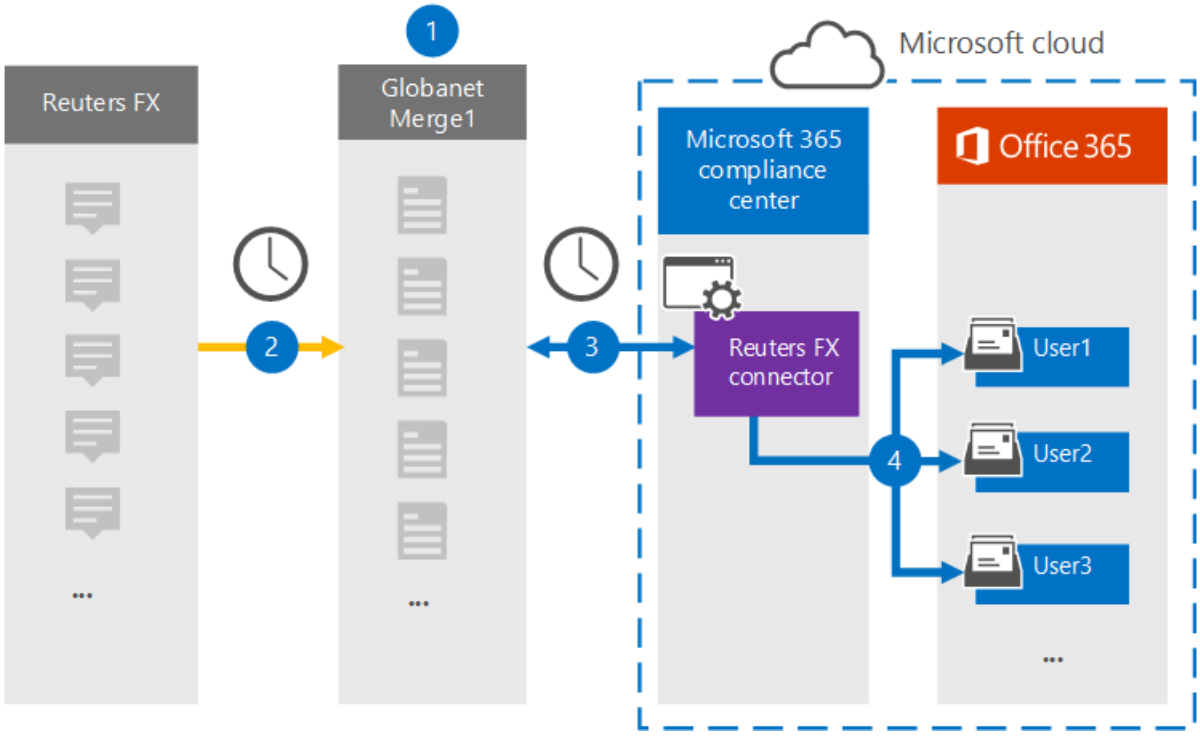

1. Your organization works with Reuters FX to set up and configure a Reuters FX site.
2. Once every 24 hours, Reuters FX items are copied to the Globanet Merge1 site. The connector also converts the items to an email message format.
3. The Reuters FX connector that you create in the Microsoft 365 compliance center, connects to the Globanet Merge1 site every day and transfers the content to a secure Azure Storage location in the Microsoft cloud.
4. The connector imports the items to the mailboxes of specific users using the value of the *Email* property of the automatic user mapping as described in Step 3. A subfolder in the Inbox folder named **Reuters FX** is created in the user mailboxes, and the items are imported to that folder. The connector determines which mailbox to import items to by using the value of the *Email* property. Every Reuters FX item contains this property, which is populated with the email address of every participant of the item.

Before you begin

- Create a Globanet Merge1 account for Microsoft connectors. To create an account, contact Globanet Customer Support. You need to sign into this account when you create the connector in Step 1.
- The user who creates the Reuters FX connector in Step 1 (and completes it in Step 3) must be assigned to the Mailbox Import Export role in Exchange Online. This role is required to add connectors on the **Data connectors** page in the Microsoft 365 compliance center. By default, this role is not assigned to any role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Step 1: Set up the Reuters FX connector

The first step is to access to the **Data Connectors** page in the Microsoft 365 and create a connector for Reuters FX data.

1. Go to https://compliance.microsoft.com and then click **Data connectors** > **Reuters FX**.
2. On the **Reuters FX** product description page, click **Add connector**.
3. On the **Terms of service** page, click **Accept**.
4. Enter a unique name that identifies the connector, and then click **Next**.
5. Sign to your Merge1 account to configure the connector.

Step 2: Configure the Reuters FX connector on the Globanet Merge1 site

The second step is to configure the Reuters FX connector on the Globanet Merge1 site. For information about configuring the Reuters FX connector, see Merge1 Third-Party Connectors User Guide.

After you click **Save & Finish**, the **User mapping** page in the connector wizard in the Microsoft 365 compliance center is displayed.

Step 3: Map users and complete the connector setup

To map users and complete the connector setup in the Microsoft 365 compliance center, follow the steps below:

1. On the **Map Reuters FX users to Microsoft 365 users** page, enable automatic user mapping.

   Reuters FX items include a property called *Email*, which contains email addresses for users in your organization. If the connector can associate this address with a Microsoft 365 user, the items are imported to that user's mailbox.

2. Click **Next**, review your settings, and go to the **Data connectors** page to see the progress of the import process for the new connector.

Step 4: Monitor the Reuters FX connector

After you create the Reuters FX connector, you can view the connector status in the Microsoft 365 compliance center.

1. Go to https://compliance.microsoft.com/ and click **Data connectors** in the left nav.
2. Click the **Connectors** tab and then select the **Reuters FX** connector to display the flyout page, which

contains the properties and information about the connector.

3. Under **Connector status with source**, click the **Download log** link to open (or save) the status log for the connector. This log contains data that has been imported to the Microsoft cloud.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.

Set up a connector to archive Salesforce Chatter data

2/18/2021 • 4 minutes to read • Edit Online

Use a Globanet connector in the Microsoft 365 compliance center to import and archive data from the Salesforce Chatter platform to user mailboxes in your Microsoft 365 organization. Globanet provides a Salesforce Chatter connector that captures items from the third-party data source and imports those items to Microsoft 365. The connector converts the content such as chats, attachments, and posts from Salesforce Chatter to an email message format and then imports those items to the user's mailbox in Microsoft 365.

After Salesforce Chatter data is stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, eDiscovery, retention policies and retention labels. Using a Salesforce Chatter connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving Salesforce Chatter data

The following overview explains the process of using a connector to archive the Salesforce Chatter data in Microsoft 365.

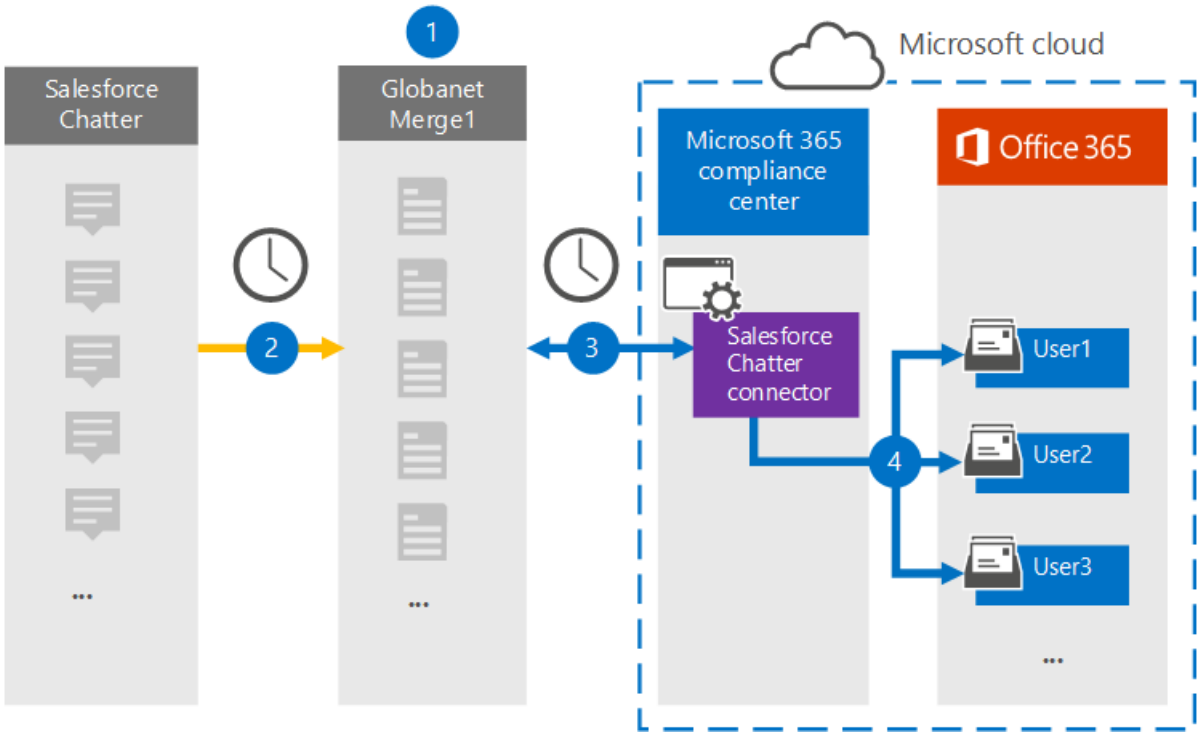

1. Your organization works with Salesforce Chatter to set up and configure a Salesforce Chatter site.
2. Once every 24 hours, Salesforce Chatter items are copied to the Globanet Merge1 site. The connector also Salesforce Chatter items to an email message format.
3. The Salesforce Chatter connector that you create in the Microsoft 365 compliance center, connects to the Globanet Merge1 site every day and transfers the Chatter content to a secure Azure Storage location in the Microsoft cloud.
4. The connector imports the converted items to the mailboxes of specific users using the value of the *Email* property of the automatic user mapping as described in Step 3. A subfolder in the Inbox folder named

Salesforce Chatter is created in the user mailboxes, and items are imported to that folder. The connector determines which mailbox to import items to by using the value of the *Email* property. Every Chatter item contains this property, which is populated with the email address of every participant of the item.

Before you begin

- Create a Merge1 account for Microsoft connectors. To create an account, contact Globanet Customer Support. You need to sign into this account when you create the connector in Step 1.
- Create a Salesforce application and acquire a token at https://salesforce.com. You'll need to log into the Salesforce account as an admin and get a user personal token to import data. Also, triggers need to be published on the Chatter site to capture updates, deletes, and edits. These triggers will create a post on a channel, and Merge1 will capture the information from the channel. For step-by-step instructions about how to create the application and acquire the token, see Merge1 Third-Party Connectors User Guide.
- The user who creates the Salesforce Chatter connector in Step 1 (and completes it in Step 3) must be assigned to the Mailbox Import Export role in Exchange Online. This role is required to add connectors on the **Data connectors** page in the Microsoft 365 compliance center. By default, this role isn't assigned to any role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Step 1: Set up the Salesforce Chatter connector

The first step is to access to the **Data Connectors** page in the Microsoft 365 compliance center and create a connector for Chatter data.

1. Go to https://compliance.microsoft.com and then click **Data connectors** > **Salesforce Chatter**.
2. On the **Salesforce Chatter** product description page, click **Add connector**.
3. On the **Terms of service** page, click **Accept**.
4. Enter a unique name that identifies the connector, and then click **Next**.
5. Sign in to your Merge1 account to configure the connector.

Step 2: Configure the Salesforce Chatter on the Globanet Merge1 site

The second step is to configure the Salesforce Chatter connector on the Globanet Merge1 site. For information about how to configure the Salesforce Chatter connector, see Merge1 Third-Party Connectors User Guide.

After you click **Save & Finish,** the **User mapping** page in the connector wizard in the Microsoft 365 compliance center is displayed.

Step 3: Map users and complete the connector setup

To map users and complete the connector setup in the Microsoft 365 compliance center, follow these steps:

1. On the **Map Salesforce Chatter users to Microsoft 365 users** page, enable automatic user mapping. The Salesforce Chatter items include a property called *Email*, which contains email addresses for users in your organization. If the connector can associate this address with a Microsoft 365 user, the items are imported to that user's mailbox.
2. click **Next**, review your settings, and then go to the **Data connectors** page to see the progress of the

import process for the new connector.

Step 4: Monitor the Salesforce Chatter connector

After you create the Salesforce Chatter connector, you can view the connector status in the Microsoft 365 compliance center.

1. Go to https://compliance.microsoft.com and click **Data connectors** in the left nav.
2. click the **Connectors** tab and then click the **Salesforce Chatter** connector to display the flyout page, which contains the properties and information about the connector.
3. Under **Connector status with source**, click the **Download log** link to open (or save) the status log for the connector. This log contains data that's been imported to the Microsoft cloud.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.

Set up a connector to archive ServiceNow data

2/18/2021 • 4 minutes to read • Edit Online

Use a Globanet connector in the Microsoft 365 compliance center to import and archive data from the ServiceNow platform to user mailboxes in your Microsoft 365 organization. Globanet provides a ServiceNow connector that captures items from the third-party data source and import those items to Microsoft 365. The connector converts the content such as live messages, attachments, and posts from ServiceNow to an email message format and then imports those items to user mailboxes in Microsoft 365.

After ServiceNow data is stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, eDiscovery, retention policies, and retention labels. Using a ServiceNow connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving ServiceNow data

The following overview explains the process of using a connector to archive the ServiceNow data in Microsoft 365.

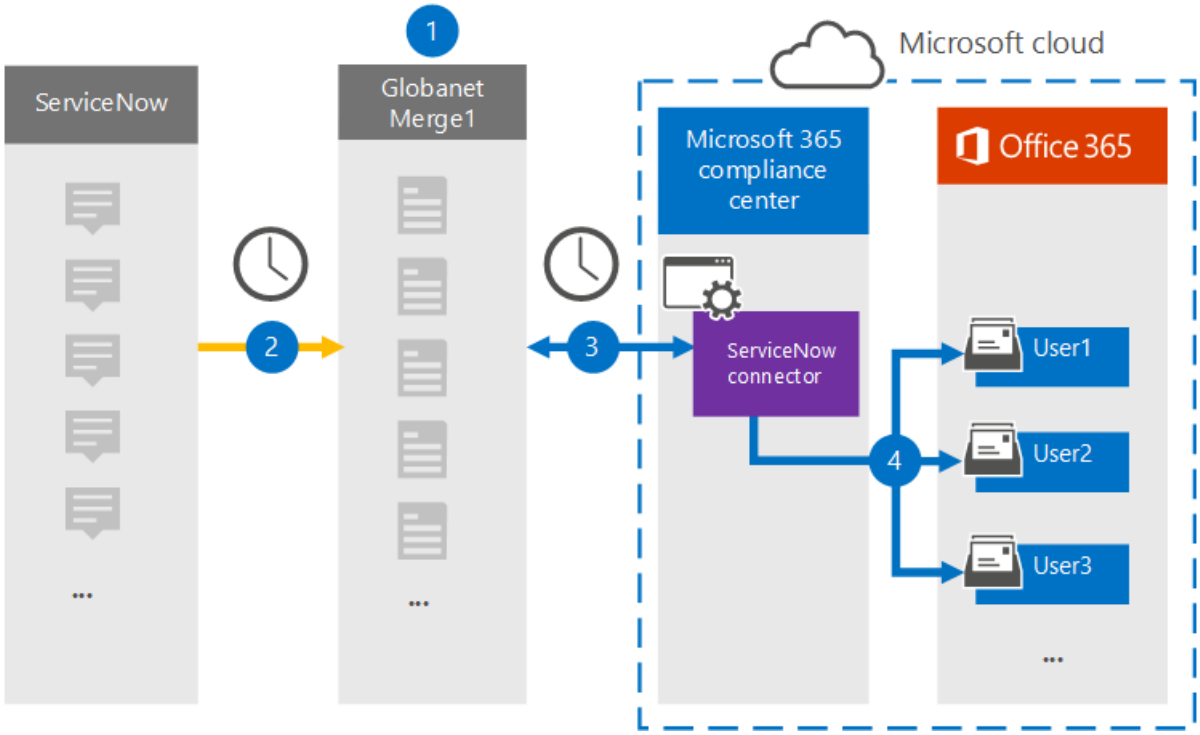

1. Your organization works with ServiceNow to set up and configure a ServiceNow site.
2. Once every 24 hours, ServiceNow items are copied to the Globanet Merge1 site. The connector also converts ServiceNow items to an email message format.
3. The ServiceNow connector that you create in the Microsoft 365 compliance center connects to the Globanet Merge1 site every day and transfers the ServiceNow content to a secure Azure Storage location in the Microsoft cloud.
4. The connector imports the converted items to the mailboxes of specific users using the value of the *Email* property of the automatic user mapping as described in Step 3. A subfolder in the Inbox folder named **ServiceNow** is created in the user mailboxes, and items are imported to that folder. The connector determines which mailbox to import items to by using the value of the *Email* property. Every ServiceNow

item contains this property, which is populated with the email address of every participant of the item.

Before you begin

- Create a Merge1 account for Microsoft connectors. To create an account, contact Globanet Customer Support. You need to sign into this account when you create the connector in Step 1.
- Create a ServiceNow application to fetch data from your ServiceNow account. For step-by step instructions about creating the application, see Merge1 Third-Party Connectors User Guide.
- The user who creates the ServiceNow connector in Step 1 (and completes it in Step 3) must be assigned to the Mailbox Import Export role in Exchange Online. This role is required to add connectors on the **Data connectors** page in the Microsoft 365 compliance center. By default, this role isn't assigned to any role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Step 1: Set up the ServiceNow connector

The first step is to access to the **Data Connectors** page in the Microsoft 365 compliance center and create a connector for ServiceNow data.

1. Go to https://compliance.microsoft.com and then click **Data connectors** > **ServiceNow**.
2. On the **ServiceNow** product description page, click **Add connector**.
3. On the **Terms of service** page, click **Accept**.
4. Enter a unique name that identifies the connector, and then click **Next**.
5. Sign in to your Merge1 account to configure the connector.

Step 2: Configure the ServiceNow on the Globanet Merge1 site

The second step is to configure the ServiceNow connector on the Globanet Merge1 site. For information about how to configure the ServiceNow connector, see Merge1 Third-Party Connectors User Guide.

After you click **Save & Finish,** the **User mapping** page in the connector wizard in the Microsoft 365 compliance center is displayed.

Step 3: Map users and complete the connector setup

To map users and complete the connector setup in the Microsoft 365 compliance center, follow these steps:

1. On the **Map ServiceNow users to Microsoft 365 users** page, enable automatic user mapping. The ServiceNow items include a property called *Email*, which contains email addresses for users in your organization. If the connector can associate this address with a Microsoft 365 user, the items are imported to that user's mailbox.
2. Click **Next**, review your settings, and then go to the **Data connectors** page to see the progress of the import process for the new connector.

Step 4: Monitor the ServiceNow connector

After you create the ServiceNow connector, you can view the connector status in the Microsoft 365 compliance center.

1. Go to https://compliance.microsoft.com and click **Data connectors** in the left nav.
2. Click the **Connectors** tab and then select the **ServiceNow** connector to display the flyout page, which contains the properties and information about the connector.
3. Under **Connector status with source**, click the **Download log** link to open (or save) the status log for the connector. This log contains data that has been imported to the Microsoft cloud.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.

Set up a connector to archive Slack eDiscovery data

2/18/2021 • 4 minutes to read • Edit Online

Use a Globanet connector in the Microsoft 365 compliance center to import and archive third-party data from social media, instant messaging, and document collaboration platforms to mailboxes in your Microsoft 365 organization. Globanet provides a Slack connector that's configured to capture items from the third-party data source (on a regular basis) and then import those items to Microsoft 365. Slack pulls messages and files from the Slack API and converts them to an email message format and then imports the item to user mailboxes.

After Slack eDiscovery data is stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, eDiscovery, retention policies and retention labels, and communication compliance. Using a Slack connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving Slack eDiscovery data

The following overview explains the process of using a connector to archive the Slack information in Microsoft 365.

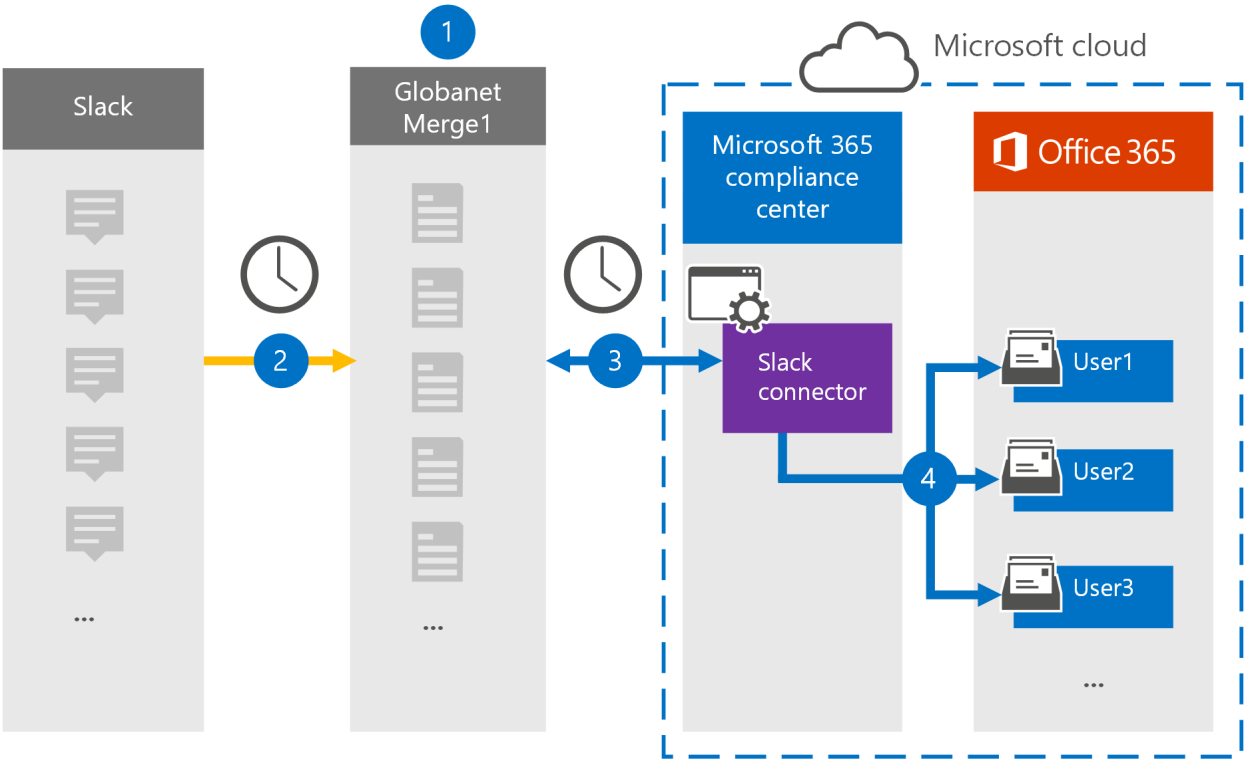

1. Your organization works with Slack to set up and configure a Slack site.
2. Once every 24 hours, chat messages from Slack eDiscovery are copied to the Globanet Merge1 site. The connector also converts the content of a chat message to an email message format.
3. The Slack eDiscovery connector that you create in the Microsoft 365 compliance center, connects to the Globanet Merge1 site every day and transfers the chat messages to a secure Azure Storage location in the Microsoft cloud.
4. The connector imports the converted chat message items to the mailboxes of specific users using the value of the *Email* property and automatic user mapping, as described in Step 3. A new subfolder in the Inbox folder named **Slack eDiscovery** is created in the user mailboxes, and the chat message items are

imported to that folder. The connector determines which mailbox to import items to by using the value of the *Email* property. Every chat message contains this property, which is populated with the email address of every participant of the chat message.

Before you begin

- Create a Globanet Merge1 account for Microsoft connectors. To create an account, contact Globanet Customer Support. You will sign into this account when you create the connector in Step 1.
- Obtain the username and password for your organization's Slack enterprise account. You'll need to sign into this account in Step 2 when you configure Slack.
- The user who creates the Slack eDiscovery connector in Step 1 (and completes it in Step 3) must be assigned to the Mailbox Import Export role in Exchange Online. This role is required to add connectors on the **Data connectors** page in the Microsoft 365 compliance center. By default, this role is not assigned to a role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Step 1: Set up the Slack eDiscovery connector

The first step is to access to the **Data Connectors** page in the Microsoft 365 compliance center and create a connector for Slack data.

1. Go to https://compliance.microsoft.com and then click **Data connectors** > **Slack eDiscovery**.
2. On the **Slack eDiscovery** product description page, click **Add connector**.
3. On the **Terms of service** page, click **Accept**.
4. Enter a unique name that identifies the connector, and then click **Next**.
5. Sign in to your Merge1 account to configure the connector.

Step 2: Configure Slack eDiscovery

The second step is to configure the Slack eDiscovery connector on the Merge1 site. For more information about how to configure the Slack eDiscovery connector on the Globanet Merge1 site, see Merge1 Third-Party Connectors User Guide.

After you click **Save & Finish**, the **User mapping** page in the connector wizard in the Microsoft 365 compliance center is displayed.

Step 3: Map users and complete the connector setup

1. On the **Map external users to Microsoft 365 users** page, enable automatic user mapping.

   Slack eDiscovery items include a property called *Email*, which contains email addresses for users in your organization. If the connector can associate this address with a Microsoft 365 user, the items are imported to that user's mailbox.

2. Click **Next**, review your settings, and go to the **Data connectors** page to see the progress of the import process for the new connector.

Step 4: Monitor the Slack eDiscovery connector

After you create the Slack eDiscovery connector, you can view the connector status in the Microsoft 365 compliance center.

1. Go to https://compliance.microsoft.com and click **Data connectors** in the left nav.
2. Click the **Connectors** tab and then select the **Slack eDiscovery** connector to display the flyout page. This page contains the properties and information about the connector.
3. Under **Connector status with source**, click the **Download log** link to open (or save) the status log for the connector. This log contains information about the data that has been imported to the Microsoft cloud.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.

Set up a connector to archive Symphony data

2/18/2021 • 4 minutes to read • Edit Online

Use a Globanet connector in the Microsoft 365 compliance center to import and archive Symphony data to user mailboxes in your Microsoft 365 organization. Symphony is a messaging and collaboration platform used in the financial services industry. Globanet provides a Symphony data connector in the Microsoft 365 compliance center that you can configure to capture items from the third-party data source (on a regular basis) and then import those items to user mailboxes. The connector converts the content of an item from the Symphony account to an email message format and then imports the item to a mailbox in Microsoft 365.

After Symphony communications are stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, eDiscovery, retention policies and retention labels, and communication compliance. Using a Symphony connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving Symphony data

The following overview explains the process of using a data connector to archive Symphony communications in Microsoft 365.

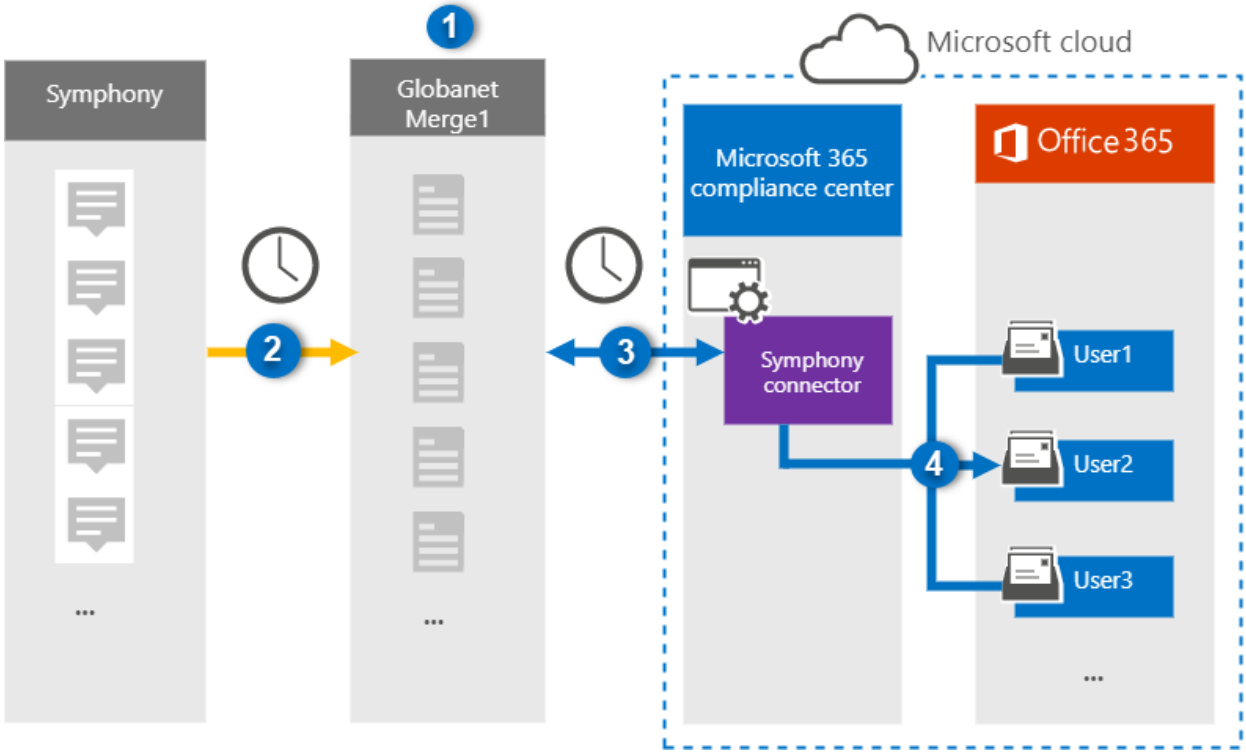

1. Your organization works with Symphony to set up and configure a Symphony site.
2. Once every 24 hours, chat messages from Symphony are copied to the Globanet Merge1 site. The connector also converts the content of a chat message to an email message format.
3. The Symphony connector that you create in the Microsoft 365 compliance center, connects to the Globanet Merge1 site every day and transfers the messages to a secure Azure Storage location in the Microsoft cloud.
4. The connector imports the converted message items to the mailboxes of specific users using the value of the *Email* property of the automatic user mapping as described in Step 3. A new subfolder in the Inbox

folder named **Symphony** is created in the user mailboxes, and the message items are imported to that folder. The connector determines which mailbox to import items to by using the value of the *Email* property. Every chat message contains this property, which is populated with the email address for every participant.

Before you begin

- Create a Globanet Merge1 account for Microsoft connectors. To create an account, contact Globanet Customer Support. You will sign into this account when you create the connector in Step 1.
- The user who creates the Symphony connector in Step 1 (and completes it in Step 3) must be assigned to the Mailbox Import Export role in Exchange Online. This role is required to add connectors on the **Data connectors** page in the Microsoft 365 compliance center. By default, this role is not assigned to a role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Step 1: Set up the Symphony connector

The first step is to access to the **Data Connectors** page in the Microsoft 365 compliance center and create a connector for Symphony data.

1. Go to https://compliance.microsoft.com and then click **Data connectors** > **Symphony**.
2. On the **Symphony** product description page, click **Add connector**.
3. On the **Terms of service** page, click **Accept**.
4. Enter a unique name that identifies the connector, and then click **Next**.
5. Sign in to your Merge1 account to configure the connector.

Configure the Symphony connector on the Globanet Merge1 site

The second step is to configure the Symphony connector on the Merge1 site. For information about configuring the Symphony connector on the Globanet Merge1 site, see Merge1 Third-Party Connectors User Guide.

After you click **Save & Finish**, the **User mapping** page in the connector wizard in the Microsoft 365 compliance center is displayed.

Step 3: Map users and complete the connector setup

To map users and complete the connector setup in the Microsoft 365 compliance center, follow these steps:

1. On the **Map external users to Microsoft 365 users** page, enable automatic user mapping. The Symphony items include a property called *Email*, which contains email addresses for users in your organization. If the connector can associate this address with a Microsoft 365 user, the items are imported to that user's mailbox.
2. Click **Next**, review your settings, and then go to the **Data connectors** page to see the progress of the import process for the new connector.

Step 4: Monitor the Symphony connector

After you create the Symphony connector, you can view the connector status in the Microsoft 365 compliance center.

1. Go to https://compliance.microsoft.com and click **Data connectors** in the left nav.
2. Click the **Connectors** tab and then select the **Symphony** connector to display the flyout page. This page contains the properties and information about the connector.
3. Under **Connector status with source**, click the **Download log** link to open (or save) the status log for the connector. This log contains information about the data that has been imported to the Microsoft cloud.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.

Set up a connector to archive text-delimited data

2/18/2021 • 4 minutes to read • Edit Online

Use a Globanet connector in the Microsoft 365 compliance center to import and archive text-delimited data to user mailboxes in your Microsoft 365 organization. Globanet provides a text-delimited connector that's configured to capture items from a third-party data source (on a regular basis) and import those items to Microsoft 365. The connector converts content from the text-delimited data source to an email message format and then imports those items to the user's mailbox in Microsoft 365.

After text-delimited data is stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, eDiscovery, and retention policies and retention labels. Using a text-delimited data connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving the text-delimited data

The following overview explains the process of using a connector to archive text-delimited source information in Microsoft 365.

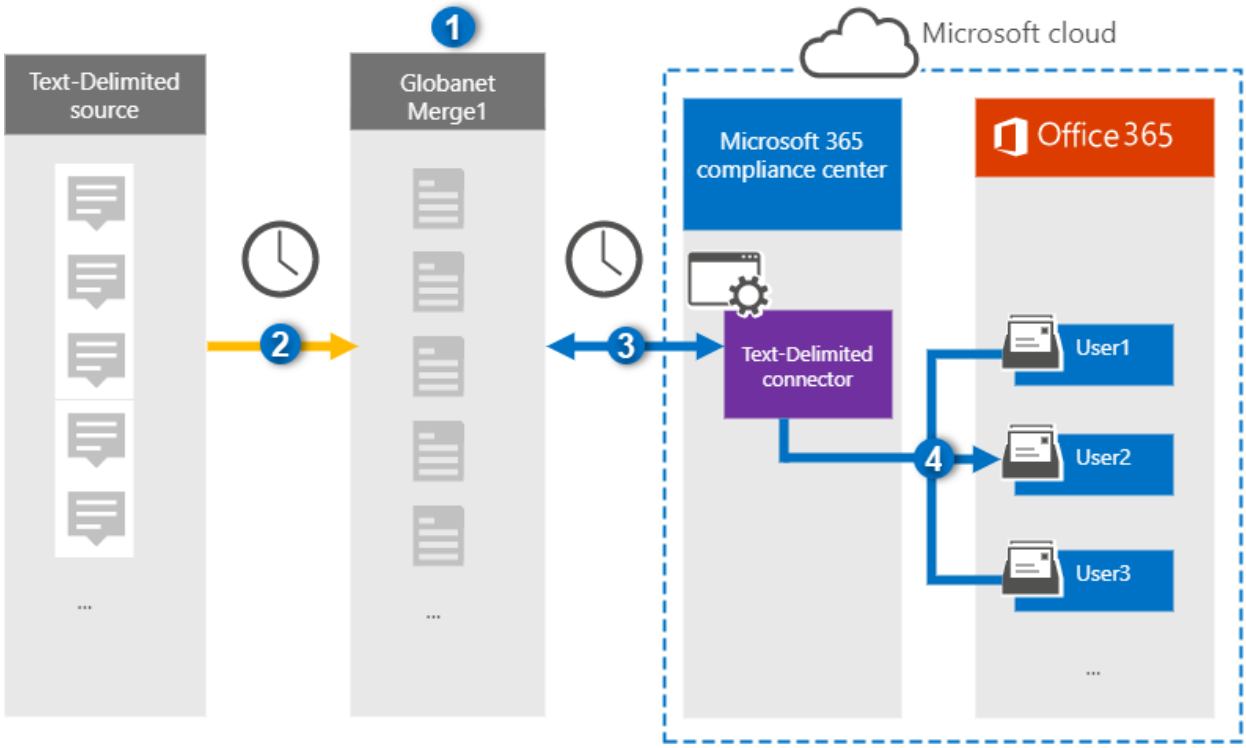

1. Your organization works with the text-delimited source to set up and configure a text-delimited site.
2. Once every 24 hours, chat messages from the text-delimited source are copied to the Globanet Merge1 site. The connector also converts the content to an email message format.
3. The text-delimited connector that you create in the Microsoft 365 compliance center connects to the Globanet Merge1 site every day and transfers the messages to a secure Azure Storage location in the Microsoft cloud.
4. The connector imports the converted message items to the mailboxes of specific users using the value of the *Email* property of the automatic user mapping as described in Step 3. A new subfolder in the Inbox folder named **Text- Delimited** is created in the user mailboxes, and the message items are imported to

that folder. The connector determines which mailbox to import items to by using the value of the *Email* property. Every message contains this property, which is populated with the email address of every participant.

Before you begin

- Create a Globanet Merge1 account for Microsoft connectors. To create this account, contact Globanet Customer Support. You will sign into this account when you create the connector in Step 1.
- The user who creates the text-delimited connector in Step 1 (and completes it in Step 3) must be assigned to the Mailbox Import Export role in Exchange Online. This role is required to add connectors on the **Data connectors** page in the Microsoft 365 compliance center. By default, this role is not assigned to a role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Step 1: Set up the text-delimited connector

The first step is to access to the **Data Connectors** page in the Microsoft 365 compliance center and create a connector for text-delimited data.

1. Go to https://compliance.microsoft.com and then click **Data connectors** > **Text-Delimited**.
2. On the **text-delimited** product description page, click **Add connector**.
3. On the **Terms of service** page, click **Accept**.
4. Enter a unique name that identifies the connector, and then click **Next**.
5. Sign in to your Merge1 account to configure the connector.

Step 2: Configure the Text-delimited connector on the Globanet Merge1 site

The second step is to configure the text-delimited connector on the Merge1 site. For information about configuring the text-delimited connector on the Globanet Merge1 site, see Merge1 Third-Party Connectors User Guide.

After you click **Save & Finish**, the **User mapping** page in the connector wizard in the Microsoft 365 compliance center is displayed.

Step 3: Map users and complete the connector setup

To map users and complete the connector setup in the Microsoft 365 compliance center, follow these steps:

1. On the **Map external users to Microsoft 365 users** page, enable automatic user mapping. The Text-Delimited source items include a property called *Email*, which contains email addresses for users in your organization. If the connector can associate this address with a Microsoft 365 user, the items are imported to that user's mailbox.
2. Click **Next**, review your settings, and then go to the **Data connectors** page to see the progress of the import process for the new connector.

Step 4: Monitor the text-delimited connector

After you create the Text- Delimited connector, you can view the connector status in the Microsoft 365

compliance center.

1. Go to https://compliance.microsoft.com and click **Data connectors** in the left nav.
2. Click the **Connectors** tab and then select the **Text- Delimited** connector to display the flyout page. This page contains the properties and information about the connector.
3. Under **Connector status with source**, click the **Download log** link to open (or save) the status log for the connector. This log contains information about the data that has been imported to the Microsoft cloud.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.

Set up a connector to archive webpage data

2/18/2021 • 4 minutes to read • Edit Online

Use a Globanet connector in the Microsoft 365 compliance center to import and archive data from webpages to user mailboxes in your Microsoft 365 organization. Globanet provides a Webpage Capture connector that captures specific webpages (and any links on those pages) in a specific website or an entire domain. The connector converts the webpage content to a PDF, PNG, or custom file format and then attaches the converted files to an email message and then imports those email items to user mailboxes in Microsoft 365.

After webpage content is stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, eDiscovery, and retention policies and retention labels. Using a Webpage Capture connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving webpage data

The following overview explains the process of using a connector to archive webpage content in Microsoft 365.

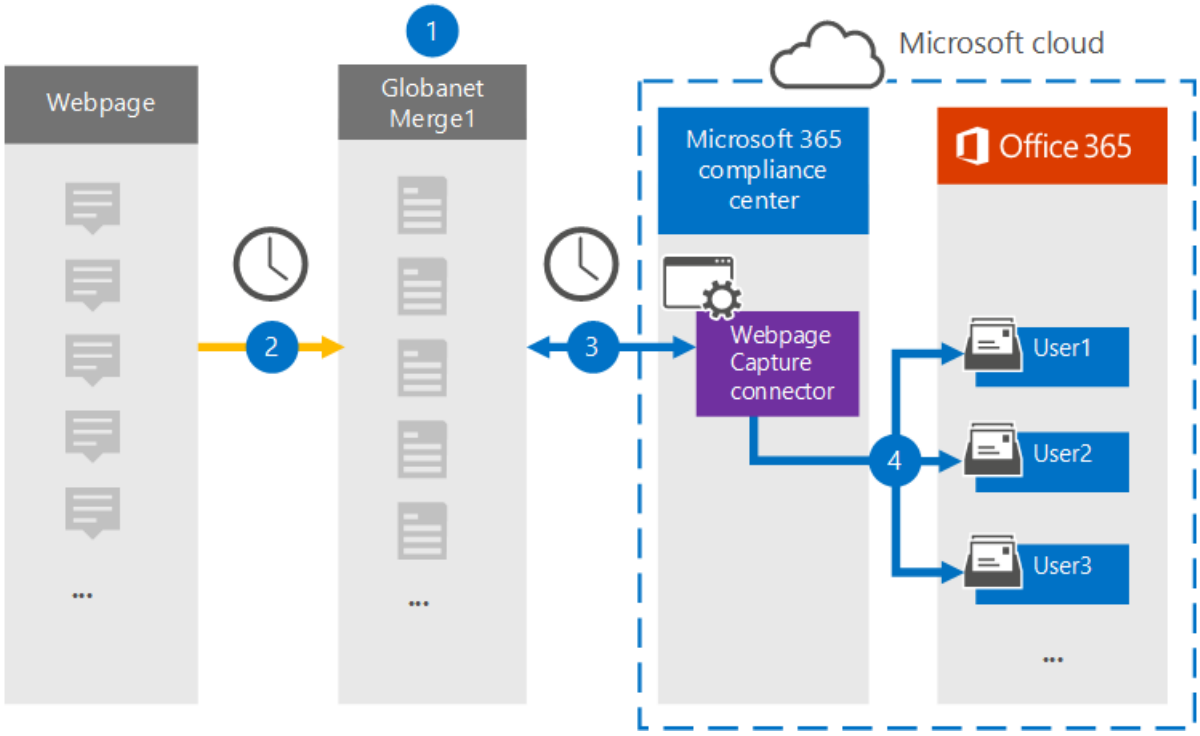

1. Your organization works with the webpage source to set up and configure a Webpage Capture site.
2. Once every 24 hours, the webpage sources items are copied to the Globanet Merge1 site. The connector also converts and attaches the content of a webpage to an email message.
3. The Webpage Capture connector that you create in the Microsoft 365 compliance center, connects to the Globanet Merge1 site every day and transfers the webpage items to a secure Azure Storage location in the Microsoft cloud.
4. The connector imports the converted webpage items to the mailboxes of specific users by using the value of the *Email* property of the automatic user mapping as described in Step 3. A subfolder in the Inbox folder named **Webpage Capture** is created in the user mailboxes, and the webpage items are imported to that folder. The connector does this by using the value of the *Email* property. Every webpage item contains this property, which is populated with the email addresses provided when you configure the

Webpage Capture connector in Step 2.

Before you begin

- Create a Globanet Merge1 account for Microsoft connectors. To create this account, contact Globanet Customer Support. You will sign into this account when you create the connector in Step 1.
- You need to work with Globanet support to set up a custom file format to convert the webpage items to. For more information, see the Merge1 Third-Party Connectors User Guide in
- The user who creates the Webpage Capture connector in Step 1 (and completes it in Step 3) must be assigned to the Mailbox Import Export role in Exchange Online. This role is required to add connectors on the **Data connectors** page in the Microsoft 365 compliance center. By default, this role is not assigned to a role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Step 1: Set up the Webpage Capture connector

The first step is to access to the **Data Connectors** and create a connector for Web Page source data.

1. Go to https://compliance.microsoft.com and then click **Data connectors** > **Webpage Capture**.
2. On the **Webpage Capture** product description page, click **Add connector**.
3. On the **Terms of service** page, click **Accept**.
4. Enter a unique name that identifies the connector, and then click **Next**.
5. Sign in to your Merge1 account to configure the connector.

Step 2: Configure the Webpage Capture connector on the Globanet Merge1 site

The second step is to configure the Webpage Capture connector on the Globanet Merge1 site. For information about how to configure the Webpage Capture connector, see Merge1 Third-Party Connectors User Guide.

After you click **Save & Finish**, the **User mapping** page in the connector wizard in the Microsoft 365 compliance center is displayed.

Step 3: Map users and complete the connector setup

To map users and complete the connector setup in the Microsoft 365 compliance center, follow the steps below:

1. On the **Map Webpage Capture users to Microsoft 365 users** page, enable automatic user mapping. The Webpage Capture items include a property called *Email*, which contains email addresses for users in your organization. If the connector can associate this address with a Microsoft 365 user, the items are imported to that user's mailbox.
2. Click **Next**, review your settings, and go to the **Data connectors** page to see the progress of the import process for the new connector.

Step 4: Monitor the Webpage Capture connector

After you create the Webpage Capture connector, you can view the connector status in the Microsoft 365 compliance center.

1. Go to https://compliance.microsoft.com and click **Data connectors** in the left nav.
2. Click the **Connectors** tab and then select the **Webpage Capture** connector to display the flyout page. This page contains the properties and information about the connector.
3. Under **Connector status with source**, click the **Download log** link to open (or save) the status log for the connector. This log contains data that has been imported to the Microsoft cloud.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.

Set up a connector to archive Webex Teams data

2/18/2021 • 4 minutes to read • Edit Online

Use a Globanet connector in the Microsoft 365 compliance center to import and archive data from Webex Teams to user mailboxes in your Microsoft 365 organization. Globanet provides a Webex Teams connector that is configured to capture Webex Teams communication items and import them to Microsoft 365. The connector converts content from Webex Teams, such as 1:1 chats, group conversations, channel conversations, and attachments from your organization's Webex Teams account, to an email message format and then imports those items to the user's mailbox in Microsoft 365.

After Webex Teams data is stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, eDiscovery, retention policies and retention labels, and communication compliance. Using a Webex Teams connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving Webex Teams data

The following overview explains the process of using a connector to archive Webex Teams data in Microsoft 365.

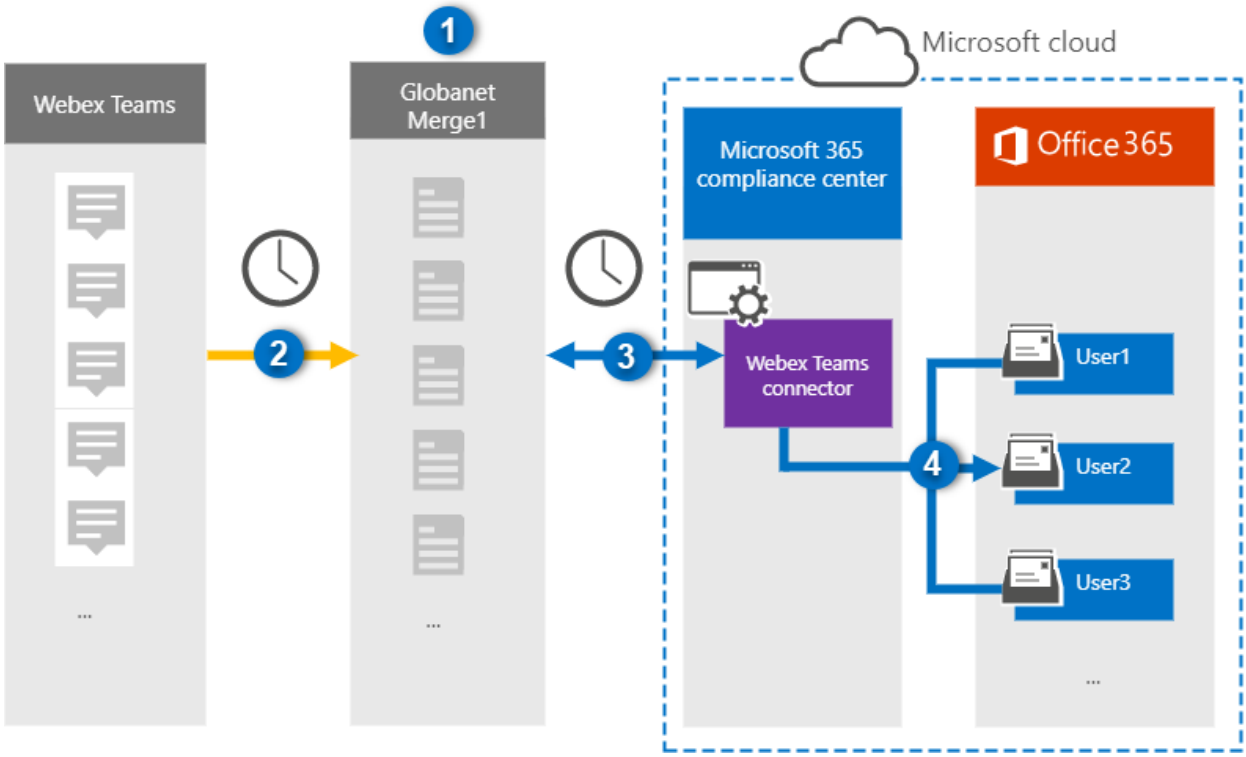

1. Your organization works with Webex Teams to set up and configure a Webex Teams site.
2. Once every 24 hours, Webex Teams items are copied to the Globanet Merge1 site. The connector also converts the Webex Teams items to an email message format.
3. The Webex Teams connector that you create in the Microsoft 365 compliance center, connects to the Globanet Merge1 every day, and transfers the Webex Teams items to a secure Azure Storage location in the Microsoft cloud.
4. The connector imports items to the mailboxes of specific users by using the value of the *Email* property of the automatic user mapping as described in Step 3. A subfolder in the Inbox folder named **Webex Teams** is created in the user mailboxes, and the items are imported to that folder. The connector does this

by using the value of the *Email* property. Every Webex Teams item contains this property, which is populated with the email address of every participant of the item.

Before you begin

- Create a Globanet Merge1 account for Microsoft connectors. To create this account, contact Globanet Customer Support. You will sign into this account when you create the connector in Step 1.
- Create an application at https://developer.webex.com/ to fetch data from your Webex Teams account. For step-by step instructions about creating the application, see Merge1 Third-Party Connectors User Guide

  When you create this application, the Webex platform generates a set of unique credentials. These credentials are used in Step 2 when you configure the Webex Teams connector on the Global Merge1 site.
- The user who creates the Webex Teams connector in Step 1 (and completes it in Step 3) must be assigned to the Mailbox Import Export role in Exchange Online. This role is required to add connectors on the **Data connectors** page in the Microsoft 365 compliance center. By default, this role is not assigned to a role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Step 1: Set up the Webex Teams connector

The first step is to gain access to the **Data Connectors** and set up the Webex Teams connector.

1. Go to https://compliance.microsoft.com and then click **Data connectors** > **Webex Teams**.
2. On the **Webex Teams** product description page, click **Add connector**.
3. On the **Terms of service** page, click **Accept**.
4. Enter a unique name that identifies the connector, and then click **Next**.
5. Sign in to your Merge1 account to configure the connector.

Step 2: Configure the Webex Teams connector on the Globanet Merge1 site

The second step is to configure the Webex Teams connector on the Merge1 site. For information about how to configure the Webex Teams connector, see Merge1 Third-Party Connectors User Guide.

After you click **Save & Finish**, the **User mapping** page in the connector wizard in the Microsoft 365 compliance center is displayed.

Step 3: Map users and complete the connector setup

To map users and complete the connector setup in the Microsoft 365 compliance center, follow these steps:

1. On the **Map Webex Teams users to Microsoft 365 users** page, enable automatic user mapping. The Webex Teams items include a property called *Email*, which contains email addresses for users in your organization. If the connector can associate this address with a Microsoft 365 user, the items are imported to that user's mailbox.
2. Click **Next**, review your settings, and then go to the **Data connectors** page to see the progress of the import process for the new connector.

Step 4: Monitor the Webex Teams connector

After you create the Webex Teams connector, you can view the connector status in the Microsoft 365 compliance center.

1. Go to https://compliance.microsoft.com and click **Data connectors** in the left nav.
2. Click the **Connectors** tab and then select the **Webex Teams** connector to display the flyout page. This page contains the properties and information about the connector.
3. Under **Connector status with source**, click the **Download log** link to open (or save) the status log for the connector. This log contains information about the data that has been imported to the Microsoft cloud.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.

Set up a connector to archive Workplace from Facebook data

2/18/2021 • 4 minutes to read • Edit Online

Use a Globanet connector in the Microsoft 365 compliance center to import and archive data from Workplace from Facebook to user mailboxes in your Microsoft 365 organization. Globanet provides a Workplace from Facebook connector that is configured to capture items from the third-party data source (on a regular basis) and import those items to Microsoft 365. The connector converts the content such as chats, attachments, posts, and videos from Workplace to an email message format and then imports those items to user mailboxes in Microsoft 365.

After Workplace data is stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, eDiscovery, retention policies and retention labels, and communication compliance. Using Workplace from Facebook connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving Workplace from Facebook data

The following overview explains the process of using a connector to archive Workplace data in Microsoft 365.

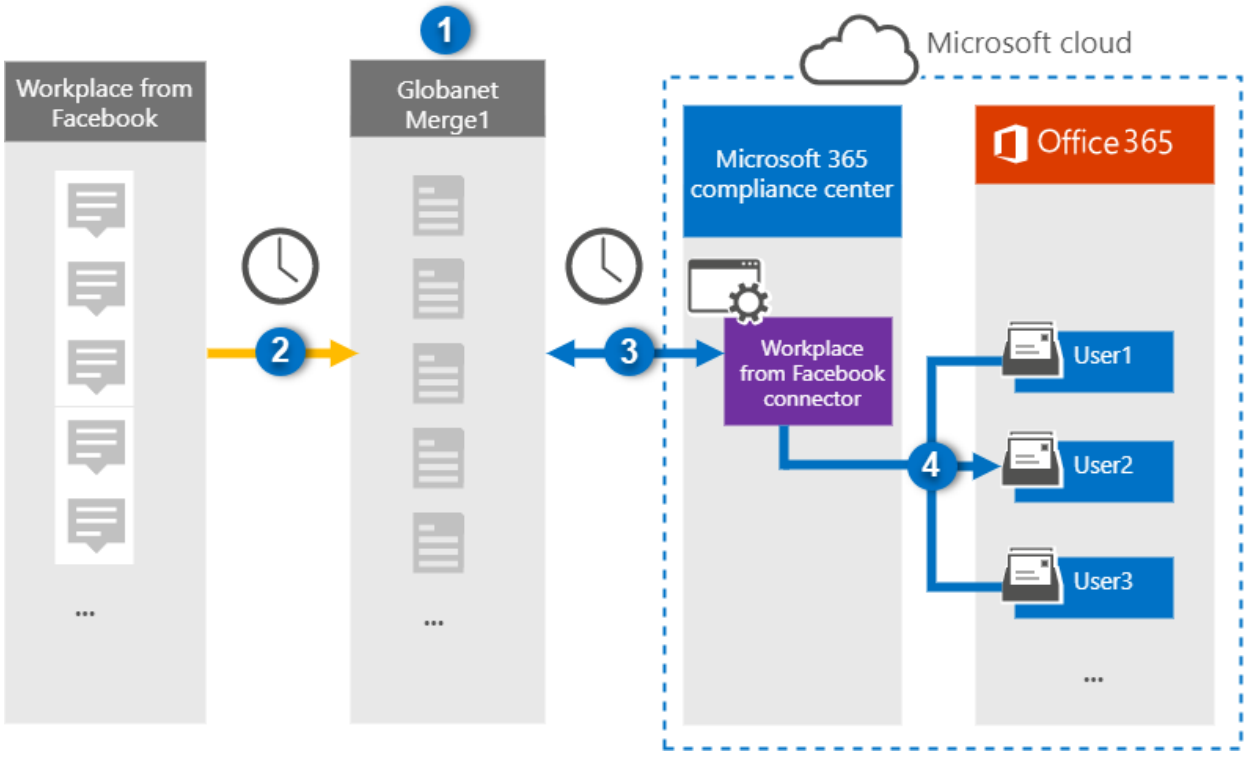

1. Your organization works with Workplace from Facebook to set up and configure a Workplace site.
2. Once every 24 hours, items from Workplace are copied to the Globanet Merge1 site. The connector also converts the content of these items to an email message format.
3. The Workplace from Facebook connector that you create in the Microsoft 365 compliance center, connects to the Globanet Merge1 every day, and transfers the Workplace items to a secure Azure Storage location in the Microsoft cloud.
4. The connector imports the converted items to the mailboxes of specific users using the value of the *Email*

property of the automatic user mapping as described in Step 3. A subfolder in the Inbox folder named **Workplace from Facebook** is created, and the Workplace items are imported to that folder. The connector does this by using the value of the *Email* property. Every Workplace item contains this property, which is populated with the email address of every chat or post participant.

Before you begin

- Create a Globanet Merge1 account for Microsoft connectors. To create this account, contact Globanet Customer Support. You will sign into this account when you create the connector in Step 1.
- Create a custom integration at https://my.workplace.com/work/admin/apps/ to retrieve data from Workplace via APIs for compliance and eDiscovery purposes.

  When creating the integration, the Workplace platform generates a set of unique credentials used to generate tokens that are used for authentication. These tokens are used in the Workplace from Facebook connector configuration wizard in Step 2. For step-by step instructions about how to create the applications, see Merge1 Third-Party Connectors User Guide.
- The user who creates the Workplace from Facebook connector in Step 1 (and completes it in Step 3) must be assigned to the Mailbox Import Export role in Exchange Online. This role is required to add connectors on the **Data connectors** page in the Microsoft 365 compliance center. By default, this role is not assigned to a role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Step 1: Set up the Workplace from Facebook connector

The first step is to access to the **Data Connectors** page in the Microsoft 365 compliance center and create a connector for Workplace data.

1. Go to https://compliance.microsoft.com and then click **Data connectors** > **Workplace from Facebook**.
2. On the **Workplace from Facebook** product description page, click **Add connector**.
3. On the **Terms of service** page, click **Accept**.
4. Enter a unique name that identifies the connector, and then click **Next**.
5. Sign in to your Merge1 account to configure the connector.

Step 2: Configure the Workplace from Facebook connector on the Globanet Merge1 site

The second step is to configure the Workplace from Facebook connector on the Merge1 site. For information about how to configure the Workplace from Facebook connector, see Merge1 Third-Party Connectors User Guide.

After you click **Save & Finish**, the **User mapping** page in the connector wizard in the Microsoft 365 compliance center is displayed.

Step 3: Map users and complete the connector setup

To map users and complete the connector setup in the Microsoft 365 compliance center, follow these steps:

1. On the **Map external users to Microsoft 365 users** page, enable automatic user mapping. The Workplace items include a property called *Email* that contains email addresses for users in your organization. If the connector can associate this address with a Microsoft 365 user, the items are imported to that user's mailbox.
2. Click **Next**, review your settings, and then go to the **Data connectors** page to see the progress of the import process for the new connector.

Step 4: Monitor the Workplace from Facebook connector

After you create the Workplace from Facebook connector, you can view the connector status in the Microsoft 365 compliance center.

1. Go to https://compliance.microsoft.com and click **Data connectors** in the left nav.
2. Click the **Connectors** tab and then select the **Workplace from Facebook** connector to display the flyout page. This page contains the properties and information about the connector.
3. Under **Connector status with source**, click the **Download log** link to open (or save) the status log for the connector. This log contains information about the data that has been imported to the Microsoft cloud.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.

Set up a connector to archive XIP source data

2/18/2021 • 4 minutes to read • Edit Online

Use a Globanet connector in the Microsoft 365 compliance center to import and archive data from the XIP source platform to user mailboxes in your Microsoft 365 organization. Globanet provides a XIP connector that allows using an XIP file to import items to Microsoft 365. An XIP file is similar to a ZIP file, but allows for a digital signature to be used. The digital signature is verified by the Globanet Merge 1 before the XIP source file is extracted. The connector converts the content from the XIP source file to an email message format and then imports those items to the user's mailbox in Microsoft 365.

After XIP source data is stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, eDiscovery, retention policies and retention labels, and communication compliance. Using an XIP connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving the XIP source data

The following overview explains the process of using a connector to archive the XIP source data in Microsoft 365.

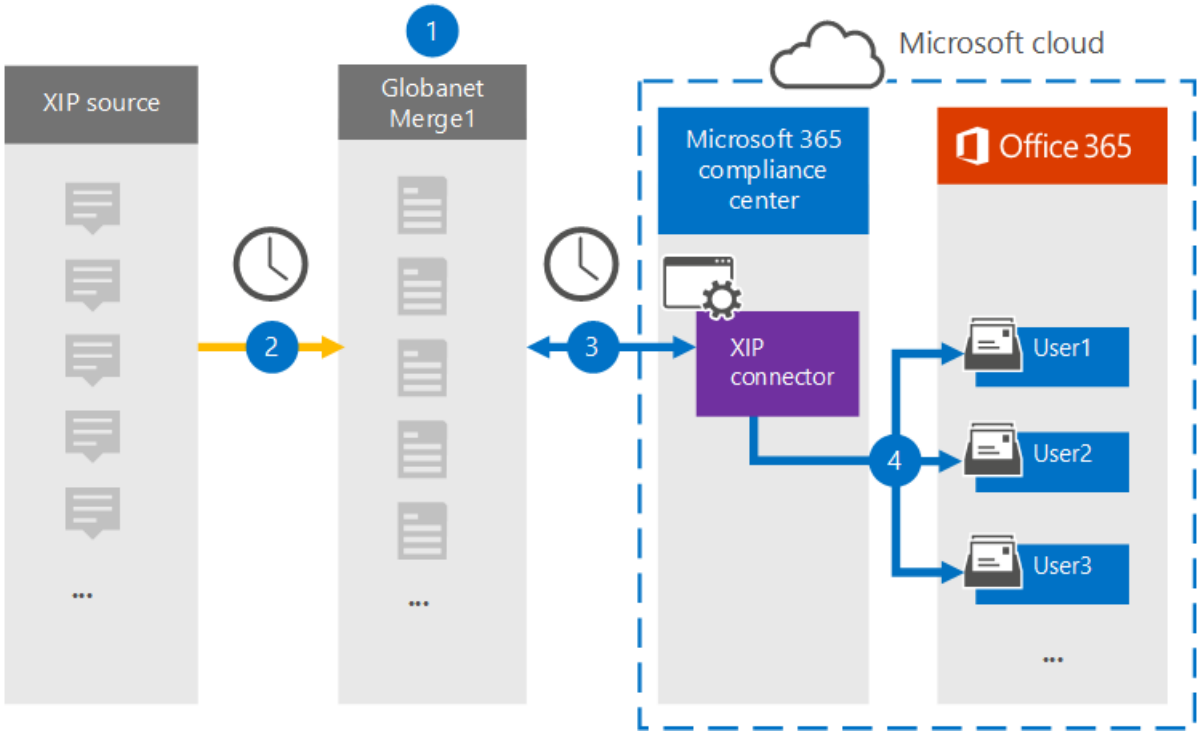

1. Your organization works with the XIP source to set up and configure an XIP site.
2. Once every 24 hours, XIP source items are copied to the Globanet Merge1 site. The connector also converts the content to an email message format.
3. The XIP connector that you create in the Microsoft 365 compliance center, connects to the Globanet Merge1 site every day and transfers the messages to a secure Azure Storage location in the Microsoft cloud.
4. The connector imports the converted message items to the mailboxes of specific users using the value of the *Email* property of the automatic user mapping as described in Step 3. A subfolder in the Inbox folder named **XIP** is created in the user mailboxes, and the items are imported to that folder. The connector

determines which mailbox to import items to by using the value of the *Email* property. Every source item contains this property, which is populated with the email address of every participant.

Before you begin

- Create a Globanet Merge1 account for Microsoft connectors. To create an account, contact Globanet Customer Support. You need to sign into this account when you create the connector in Step 1.
- The user who creates the XIP connector in Step 1 (and completes it in Step 3) must be assigned to the Mailbox Import Export role in Exchange Online. This role is required to add connectors on the Data connectors page in the Microsoft 365 compliance center. By default, this role is not assigned to any role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Step 1: Set up the XIP connector

The first step is to access to the **Data Connectors** page in the Microsoft365 compliance center and create a connector for the XIP source data.

1. Go to https://compliance.microsoft.com and then click **Data connectors** > **XIP**.
2. On the **XIP** product description page, click **Add new connector**.
3. On the **Terms of service** page, click **Accept**.
4. Enter a unique name that identifies the connector, and then click **Next**.
5. Sign in to your Merge1 account to configure the connector.

Step 2: Configure the XIP connector on the Globanet Merge1 site

The second step is to configure the XIP connector on the Merge1 site. For information about how to configure the XIP connector, see Merge1 Third-Party Connectors User Guide.

After you click **Save & Finish**, the **User mapping** page in the connector wizard in the Microsoft 365 compliance center is displayed.

Step 3: Map users and complete the connector setup

To map users and complete the connector setup, follow these steps:

1. On the **Map XIP users to Microsoft 365 users** page, enable automatic user mapping. The XIP source items include a property called *Email*, which contains email addresses for users in your organization. If the connector can associate this address with a Microsoft 365 user, the items are imported to that user's mailbox.
2. Click **Next**, review your settings, and go to the **Data connectors** page to see the progress of the import process for the new connector.

Step 4: Monitor the XIP connector

After you create the XIP connector, you can view the connector status in the Microsoft 365 compliance center.

1. Go to https://compliance.microsoft.com and click **Data connectors** in the left nav.
2. Click the **Connectors** tab and then select the **XIP** connector to display the flyout page, which contains

the properties and information about the connector.

3. Under **Connector status with source**, click the **Download log** link to open (or save) the status log for the connector. This log contains data that has been imported to the Microsoft cloud.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.

Set up a connector to archive XSLT/XML data

2/18/2021 • 4 minutes to read • Edit Online

Use a Globanet connector in the Microsoft 365 compliance center to import and archive data from the Web page source to user mailboxes in your Microsoft 365 organization. Globanet provides you with an XSLT/XML connector that allows the rapid development of files created by using XSLT (Extensible Style sheet Language Transformations) to transform XML files into other file formats (such as HTML or text) that can be imported to Microsoft 365. The connector converts the content of an item from the XSLT/XML source to an email message format and then imports the converted item to Microsoft 365 mailboxes.

After XSLT/XML data is stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, eDiscovery, and retention policies and retention labels. Using an XSLT/XML connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving XSLT/XML data

The following overview explains the process of using a connector to archive XSLT/XML source data in Microsoft 365.

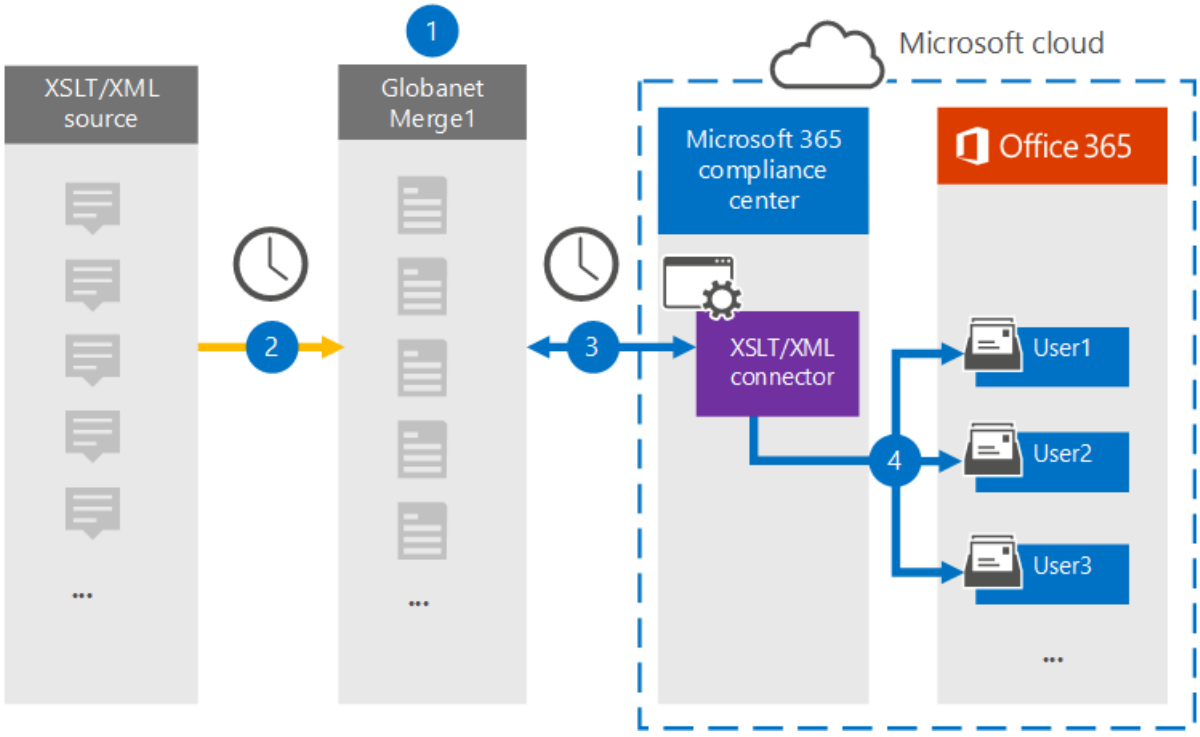

1. Your organization works with the XSLT/XML source to set up and configure an XSLT/XML site.
2. Once every 24 hours, chat messages from the XSLT/XML source are copied to the Globanet Merge1 site. The connector also converts the content to an email message format.
3. The XSLT/XML connector that you create in the Microsoft 365 compliance center, connects to the Globanet Merge1 site every day and transfers the messages to a secure Azure Storage location in the Microsoft cloud.
4. The connector imports the converted message items to the mailboxes of specific users using the value of the *Email* property of the automatic user mapping as described in Step 3. A new subfolder in the Inbox folder named **XSLT/XML** is created in the user mailboxes, and the message items are imported to that

folder. The connector does this by using the value of the *Email* property. Every message contains this property, which is populated with the email address of every participant of the message.

Before you begin

- Create a Globanet Merge1 account for Microsoft connectors. To create this account, contact Globanet Customer Support. You will sign into this account when you create the connector in Step 1.
- The user who creates the XSLT/XML connector in Step 1 (and completes it in Step 3) must be assigned to the Mailbox Import Export role in Exchange Online. This role is required to add connectors on the **Data connectors** page in the Microsoft 365 compliance center. By default, this role is not assigned to a role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Step 1: Set up an XSLT/XML connector

The first step is to access to the **Data Connectors** in the Microsoft 365 compliance center and create a connector for XSLT/XML data.

1. Go to https://compliance.microsoft.com and then click **Data connectors** > **XSLT/XML**.
2. On the **XSLT/XML** product description page, click **Add new connector**.
3. On the **Terms of service** page, click **Accept**.
4. Enter a unique name that identifies the connector, and then click **Next**.
5. Sign in to your Merge1 account to configure the connector.

Step 2: Configure an XSLT/XML connector

The second step is to configure the XSLT/XML connector on the Merge1 site. For information about how to configure the XSLT/XML connector on the Globanet Merge1 site, see Merge1 Third-Party Connectors User Guide.

After you click **Save & Finish**, the **User mapping** page in the connector wizard in the Microsoft 365 compliance center is displayed.

Step 3: Map users and complete the connector setup

1. To map users and complete the connector setup in the Microsoft 365 compliance center, follow the steps below:
2. On the **Map XSLT/XML users to Microsoft 365 users** page, enable automatic user mapping. The XSLT/XML items include a property called *Email*, which contains email addresses for users in your organization. If the connector can associate this address with a Microsoft 365 user, the items are imported to that user's mailbox.
3. Click **Next**, review your settings, and go to the **Data connectors** page to see the progress of the import process for the new connector.

Step 4: Monitor the XSLT/XML connector

After you create the XSLT/XML connector, you can view the connector status in the Microsoft 365 compliance center.

1. Go to https://compliance.microsoft.com and click **Data connectors** in the left nav.
2. Click the **Connectors** tab and then select the **XSLT/XML** connector to display the flyout page. This page contains the properties and information about the connector.
3. Under **Connector status with source**, click the **Download log** link to open (or save) the status log for the connector. This log contains data that has been imported to the Microsoft cloud.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.

Set up a connector to archive Yieldbroker data

2/18/2021 • 4 minutes to read • Edit Online

Use a Globanet connector in the Microsoft 365 compliance center to import and archive data from the Yieldbroker to user mailboxes in your Microsoft 365 organization. Globanet provides you with a Yieldbroker connector that's configured to capture items from the third-party data source and import those items to Microsoft 365. The connector converts the content from Yieldbroker to an email message format and then imports those items to the user's mailbox in Microsoft 365.

After Yieldbroker is stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, eDiscovery, retention policies, and retention labels. Using a Yieldbroker connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving Yieldbroker data

The following overview explains the process of using a connector to archive the Yieldbroker data in Microsoft 365.

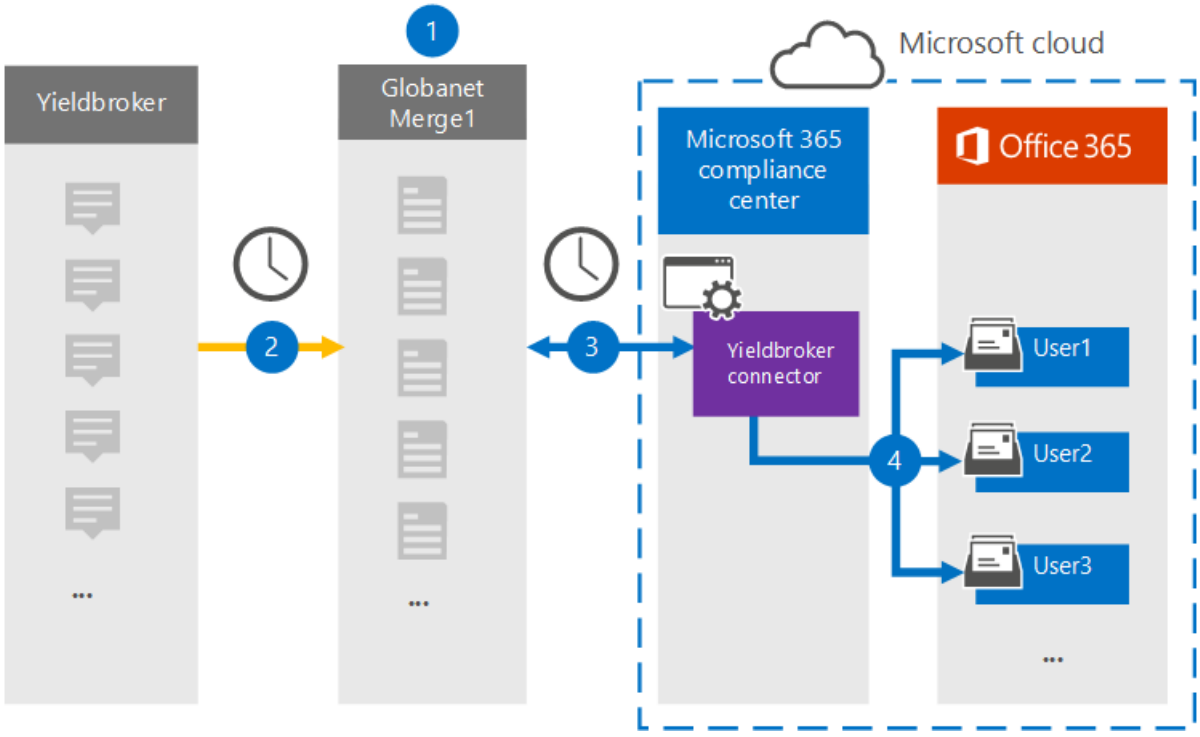

1. Your organization works with the Yieldbroker to set up and configure a Yieldbroker site.
2. Once every 24 hours, Yieldbroker items are copied to the Globanet Merge1 site. The connector also converts the content to an email message format.
3. The Yieldbroker connector that you create in the Microsoft 365 compliance center, connects to the Globanet Merge1 site every day and transfers the messages to a secure Azure Storage location in the Microsoft cloud.
4. The connector imports the converted Yieldbroker items to the mailboxes of specific users using the value of the *Email* property of the automatic user mapping as described in Step 3. A subfolder in the Inbox folder named **Yieldbroker** is created in the user mailboxes, and the items are imported to that folder. The connector determines which mailbox to import items to by using the value of the *Email* property.

Every Yieldbroker contains this property, which is populated with the email address of every participant of the item.

Before you begin

- Create a Globanet Merge1 account for Microsoft connectors. To create an account, contact Globanet Customer Support. You need to sign into this account when you create the connector in Step 1.
- The user who creates the Yieldbroker connector in Step 1 (and completes it in Step 3) must be assigned to the Mailbox Import Export role in Exchange Online. This role is required to add connectors on the Data connectors page in the Microsoft 365 compliance center. By default, this role is not assigned to any role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Step 1: Set up the Yieldbroker connector

The first step is to access to the **Data Connectors** page in the Microsoft 365 compliance center and create a connector for the Yieldbroker.

1. Go to https://compliance.microsoft.com and then click **Data connectors** > **Yieldbroker**.
2. On the **Yieldbroker** product description page, click **Add new connector**.
3. On the **Terms of service** page, click **Accept**.
4. Enter a unique name that identifies the connector, and then click **Next**.
5. Sign in to your Merge1 account to configure the connector.

Step 2: Configure the Yieldbroker connector on the Globanet Merge1 site

The second step is to configure the Yieldbroker connector on the Merge1 site. For information about how to configure the Yieldbroker, see Merge1 Third-Party Connectors User Guide.

After you click **Save & Finish**, the **User mapping** page in the connector wizard in the Microsoft 365 compliance center is displayed.

Step 3: Map users and complete the connector setup

To map users and complete the connector setup, follow these steps:

1. On the **Map Yieldbroker users to Microsoft 365 users** page, enable automatic user mapping. The Yieldbroker items include a property called *Email*, which contains email addresses for users in your organization. If the connector can associate this address with a Microsoft 365 user, the items are imported to that user's mailbox.
2. Click **Next**, review your settings, and go to the **Data connectors** page to see the progress of the import process for the new connector.

Step 4: Monitor the Yieldbroker connector

After you create the Yieldbroker connector, you can view the connector status in the Microsoft 365 compliance center.

1. Go to https://compliance.microsoft.com and click **Data connectors** in the left nav.
2. Click the **Connectors** tab and then select the **Yieldbroker** connector to display the flyout page, which contains the properties and information about the connector.
3. Under **Connector status with source**, click the **Download log** link to open (or save) the status log for the connector. This log contains data that has been imported to the Microsoft cloud.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.

Set up a connector to archive Zoom Meetings data

2/18/2021 • 4 minutes to read • Edit Online

Use a Globanet connector in the Microsoft 365 compliance center to import and archive data from Zoom Meetings to user mailboxes in your Microsoft 365 organization. Globanet provides a Zoom Meetings connector that is configured to capture items from the third-party data source (on a regular basis) and import those items to Microsoft 365. The connector converts the content of the meetings (including chats, recorded files, and metadata) from the Zoom Meetings account to an email message format and then imports those items to user mailboxes in Microsoft 365.

After Zoom Meetings data is stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, eDiscovery, retention policies and retention labels, and communication compliance. Using a Zoom Meetings connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving Zoom Meetings data

The following overview explains the process of using a connector to archive Zoom Meetings data in Microsoft 365.

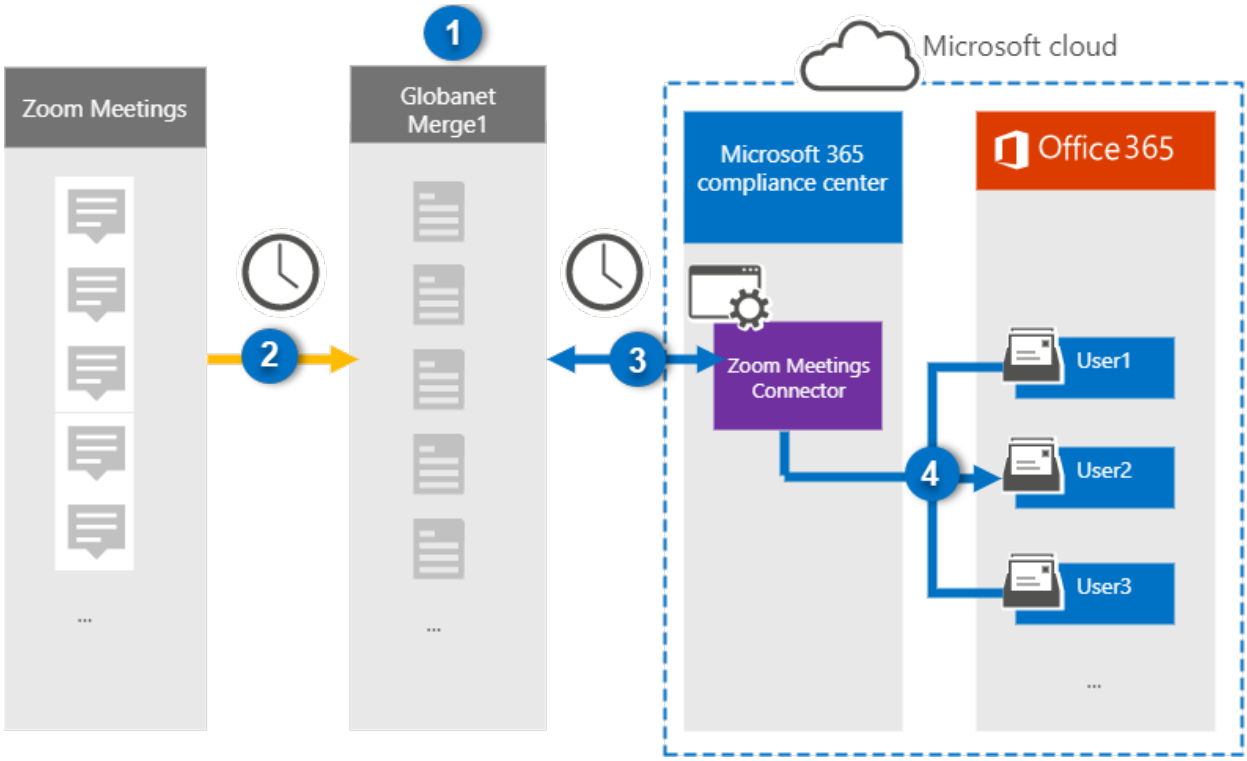

1. Your organization works with Zoom Meetings to set up and configure a Zoom Meetings site.
2. Once every 24 hours, meeting items from Zoom Meetings are copied to the Globanet Merge1 site. The connector also converts the content of the meetings to an email message format.
3. The Zoom Meetings connector that you create in the Microsoft 365 compliance center, connects to the Globanet Merge1 every day, and transfers the meeting messages to a secure Azure Storage location in the Microsoft cloud.
4. The connector imports the converted meeting items to the mailboxes of specific users using the value of the *Email* property and automatic user mapping, as described in Step 3. A new subfolder in the Inbox

folder named **Zoom Meetings** is created in user mailboxes, and the meeting items are imported to that folder. The connector does this by using the value of the *Email* property. Every meeting item contains this property, which is populated with the email address of every participant of the meeting.

Before you begin

- Create a Globanet Merge1 account for Microsoft connectors. To create this account, contact Globanet Customer Support. You will sign into this account when you create the connector in Step 1.
- Obtain the username and password for your organization's Zoom Business or Zoom Enterprise account. You'll need to sign into this account in Step 2 when you configure the Zoom Meetings connector.
- Create the following applications in the Zoom Marketplace:
  - OAuth application
  - JWT application

  After you create these applications, the Zoom platform generates a set of unique credentials used to generate the tokens. These tokens are used to authenticate the connector when it connects to your Zoom account and copies items to the Merge1 site. You will use these tokens when you configure the Zoom connector in Step 2.

  For step-by step instructions on how to create the OAuth and JWT applications, see Merge1 Third-Party Connectors User Guide.
- The user who creates the Zoom Meetings connector in Step 1 (and completes it in Step 3) must be assigned to the Mailbox Import Export role in Exchange Online. This role is required to add connectors on the **Data connectors** page in the Microsoft 365 compliance center. By default, this role is not assigned to a role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Step 1: Set up the Zoom Meetings connector

The first step is to access the **Data Connectors** in the Microsoft 365 compliance center and create a Zoom Meetings connector.

1. Go to https://compliance.microsoft.com and then click **Data connectors** > **Zoom Meetings**.
2. On the **Zoom Meetings** product description page, click **Add connector**.
3. On the **Terms of service** page, click **Accept**.
4. Enter a unique name that identifies the connector, and then click **Next**.
5. Sign in to your Merge1 account to configure the connector.

Step 2: Configure the Zoom Meetings connector

The second step is to configure the Zoom Meetings connector on the Merge1 site. For more information about how to configure the Zoom Meetings connector on the Globanet Merge1 site, see Merge1 Third-Party Connectors User Guide.

After you click **Save & Finish**, the **User mapping** page in the connector wizard in the Microsoft 365 compliance center is displayed.

Step 3: Map users and complete the connector setup

1. On the **Map external users to Microsoft 365 users** page, enable automatic user mapping.

   Zoom Meetings items include a property called *Email* that contains email addresses for users in your organization. If the connector can associate this address with a Microsoft 365 user, the items are imported to that user's mailbox

2. Click **Next**, review your settings, and go to the **Data connectors** page to see the progress of the import process for the new connector.

Step 4: Monitor the Zoom Meetings connector

After you create the Zoom Meetings connector, you can view the connector status in the Microsoft 365 compliance center.

1. Go to https://compliance.microsoft.com and click **Data connectors** in the left nav.
2. Click the **Connectors** tab and then select the **Zoom Meetings** connector to display the flyout page. This page contains the properties and information about the connector.
3. Under **Connector status with source**, click the **Download log** link to open (or save) the status log for the connector. This log contains information about the data that has been imported to the Microsoft cloud.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.
- For the Zoom Meetings connector to work, you must enable recordings when setting up Zoom Meetings.

Set up a connector to archive Android mobile data

2/18/2021 • 4 minutes to read • Edit Online

Use a TeleMessage connector in the Microsoft 365 compliance center to import and archive SMS, MMS, voice calls, and call logs from Android mobile phones. After you set up and configure a connector, it connects to your organization's TeleMessage account once every day, and imports the mobile communication of employees using the TeleMessage Android Archiver to mailboxes in Microsoft 365.

After data from Android mobile phones is stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, Content Search, and Microsoft 365 retention policies to Android Archiver data. For example, you can search Android Archiver mobile communication using Content Search or associate the mailbox that contains the Android Archiver connector data with a custodian in an Advanced eDiscovery case. Using an Android Archiver connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving Android mobile data

The following overview explains the process of using a connector to archive Android mobile data in Microsoft 365.

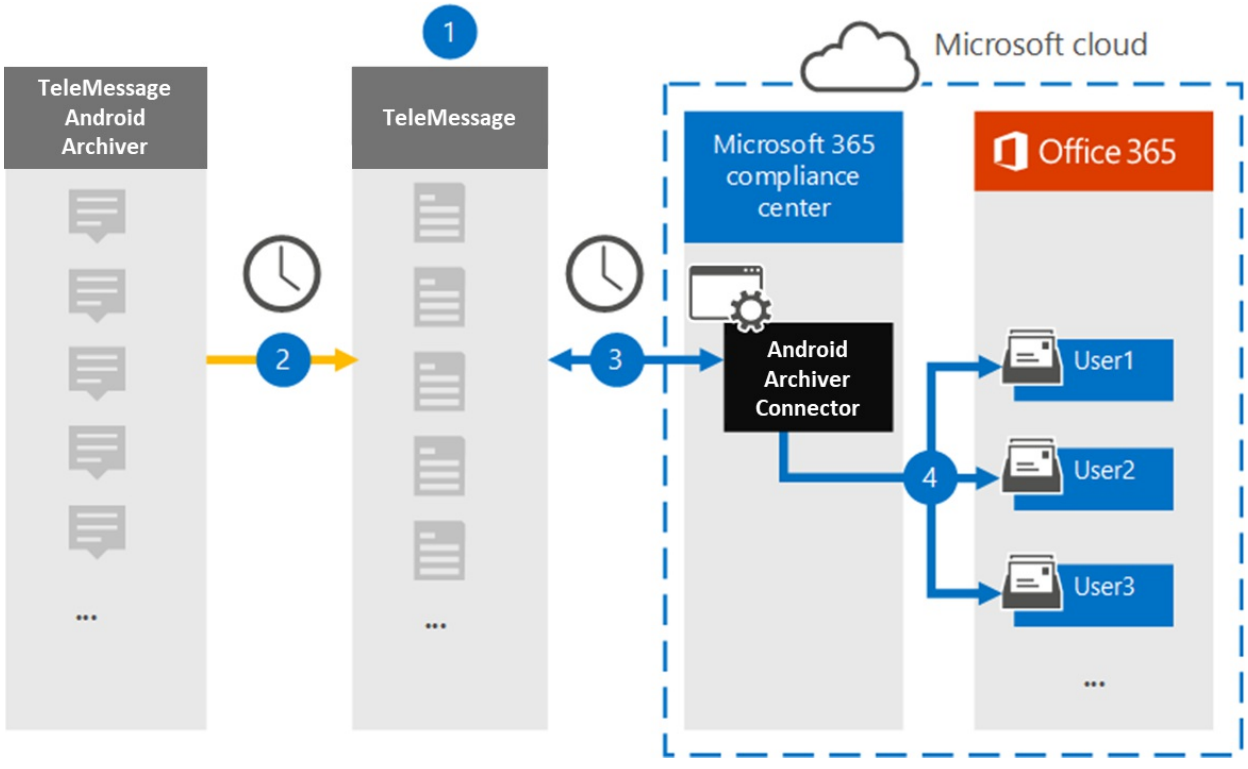

1. Your organization works with TeleMessage to set up an Android Archiver connector. For more information, see Android Archiver.

2. Once every 24 hours, SMS, MMS, voice calls, and call logs from your organization's Android mobile phones are copied to the TeleMessage site.

3. The Android Archiver connector that you create in the Microsoft 365 compliance center connects to the TeleMessage site every day and transfers the Android data from the previous 24 hours to a secure Azure Storage location in the Microsoft Cloud. The connector also converts the Android data to an email message format.

4. The connector imports the mobile communication items to the mailbox of a specific user. A new folder named Android Archiver is created in the specific user's mailbox and the items are imported to it. The connector does mapping by using the value of the *User's Email address* property. Every email message contains this property, which is populated with the email address of every participant of the email message. In addition to automatic user mapping using the value of the *User's Email address* property, you can also define a custom mapping by uploading a CSV mapping file. This mapping file should contain User's mobile Number and the corresponding Microsoft 365 mailbox address for each user. If you enable automatic user mapping and provide a custom mapping, for every email item the connector will first look at custom mapping file. If it doesn't find a valid Microsoft 365 user that corresponds to a user's mobile number, the connector will use the User 's email address property of the email item. If the connector doesn't find a valid Microsoft 365 user in either the custom mapping file or the *user's email address* property of the email item, the item won't be imported.

Before you begin

Some of the implementation steps required to archive Android communication data are external to Microsoft 365 and must be completed before you can create the connector in the compliance center.

- Order the Android Archiver service from TeleMessage and get a valid administration account for your organization. You'll need to sign into this account when you create the connector.
- Register all users that require the Android Archiver service in the TeleMessage account. When registering users, be sure to use the same email address that's used for their Microsoft 365 account.
- Install and activate the TeleMessage Android Archiver app on the mobile phones of your employees.
- The user who creates a Android Archiver connector must be assigned the Mailbox Import Export role in Exchange Online. This is required to add connectors in the **Data connectors** page in the Microsoft 365 compliance center. By default, this role isn't assigned to any role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Create an Android Archiver connector

The last step is to create an Android Archiver connector in the Microsoft 365 compliance center. The connector uses the information you provide to connect to the TeleMessage site and transfer Android communication to the corresponding user mailbox boxes in Microsoft 365.

1. Go to https://compliance.microsoft.com and click **Data connectors** > **Android Archiver**.
2. On the **Android Archiver** product description page, click **Add connector**.
3. On the **Terms of service** page, click **Accept**.
4. On the **Login to TeleMessage** page, under Step 3, enter the required information in the following boxes and then click **Next**.
   - **Username:** Your TeleMessage username.
   - **Password:** Your TeleMessage password.
5. After the connector is created, close the pop-up window and click **Next**.
6. On the **User mapping** page, enable automatic user mapping and click **Next**. In case you need custom mapping upload a CSV file, and click **Next**.

7. Review your settings, and then click **Finish** to create the connector.
8. Go to the Connectors tab in **Data connectors** page to see the progress of the import process for the new connector.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.

Set up a connector to archive AT&T SMS/MMS data

2/18/2021 • 4 minutes to read • Edit Online

Use a TeleMessage connector in the Microsoft 365 compliance center to import and archive SMS and MMS data from AT&T Mobile Network. After you set up and configure a connector, it connects to your organization's AT&T Network once every day, and imports SMS and MMS data to mailboxes in Microsoft 365.

After SMS and MMS messages are stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, Content Search, and Microsoft 365 retention policies to AT&T Network data. For example, you can search AT&T Network data using Content Search or associate the mailbox that contains the AT&T Network connector data with a custodian in an Advanced eDiscovery case. Using a AT&T Network connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving AT&T Network data

The following overview explains the process of using a connector to archive AT&T Network data in Microsoft 365.

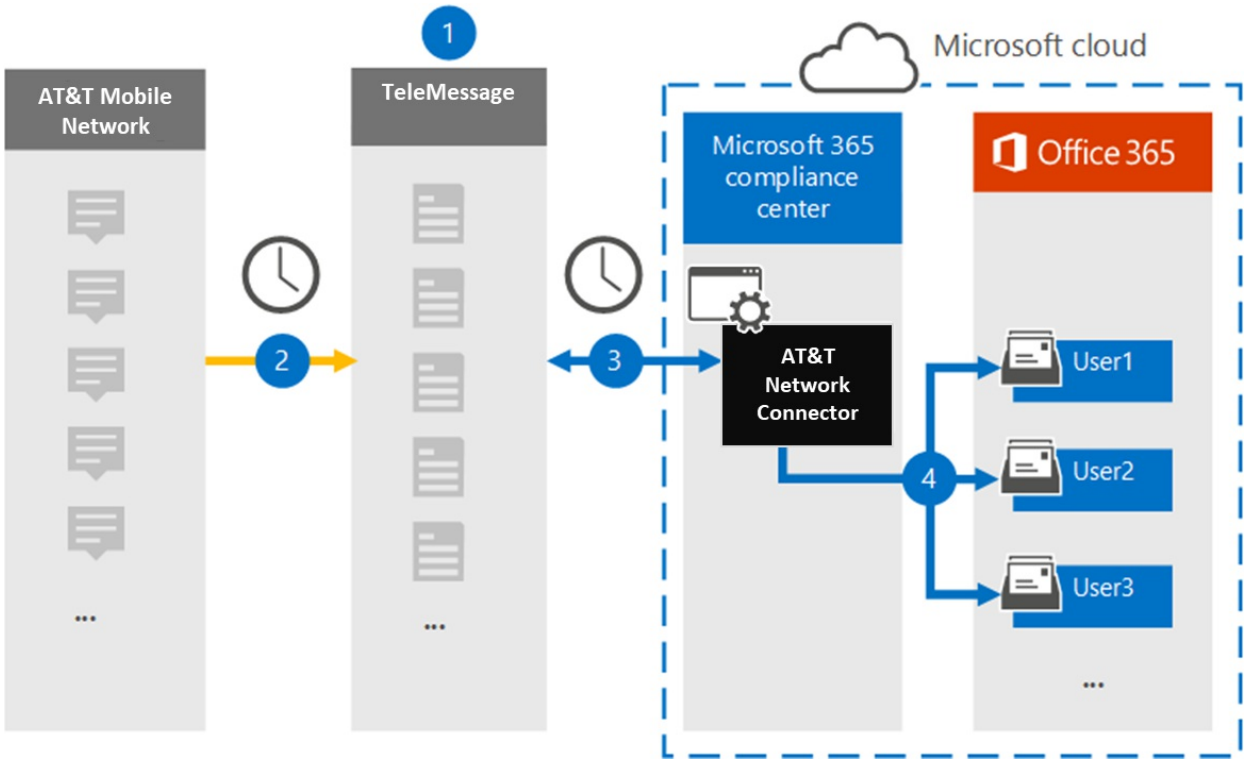

1. Your organization works with TeleMessage to set up an AT&T Network connector. For information, see AT&T Network Archiver.

2. Once every 24 hours, SMS and MMS messages from your organization's AT&T Network are copied to the TeleMessage site.

3. The AT&T Network connector that you create in the Microsoft 365 compliance center connects to the TeleMessage site every day and transfers the SMS and MMS messages from the previous 24 hours to a secure Azure Storage location in the Microsoft Cloud. The connector also converts the content of SMS and MMS messages to an email message format.

4. The connector imports the mobile communication items to the mailbox of specific users. A new folder named **AT&T SMS/MMS Network Archiver** is created in the user's mailbox and the items are imported to it. The connector does this mapping by using the value of the *User's Email address* property. Every SMS and MMS message contains this property, which is populated with the email address of every participant of the message.

   In addition to automatic user mapping using the value of the *User's Email address* property, you can also define a custom mapping by uploading a CSV mapping file. This mapping file contains the mobile phone number and corresponding Microsoft 365 email address for users in your organization. If you enable both automatic user mapping and custom mapping, for every email item the connector first looks at the custom mapping file. If it doesn't find a valid Microsoft 365 user that corresponds to a mobile phone number, the connector uses the values in the email address property of the item it's trying to import. If the connector doesn't find a valid Microsoft 365 user in either the custom mapping file or in the email address property of the email item, the item won't be imported.

Before you begin

Some of the implementation steps required to archive AT&T Network data are external to Microsoft 365 and must be completed before you can create the connector in the compliance center.

- Order the mobile archiver service from TeleMessage and get a valid administration account for your organization. You'll need to sign into this account when you create the connector in the compliance center.
- Obtain your AT&T account and billing contact details so you can fill-out the TeleMessage onboarding forms and order the message archiving service from AT&T.
- Register all users that require AT&T SMS/MMS Network archiving in the TeleMessage account. When registering users, be sure to use the same email address that's used for their Microsoft 365 account.
- Your employees must have corporate-owned and corporate-liable mobile phones on the AT&T mobile network. Archiving messages in Microsoft 365 isn't available for employee-owned or "Bring Your Own Devices (BYOD) devices.
- The user who creates a AT&T Network connector must be assigned the Mailbox Import Export role in Exchange Online. This is required to add connectors in the **Data connectors** page in the Microsoft 365 compliance center. By default, this role isn't assigned to any role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Create a AT&T Network connector

After you've completed the prerequisites described in the previous section, you can create an AT&T Network connector in the Microsoft 365 compliance center. The connector uses the information you provide to connect to the TeleMessage site and transfer SMS and MMS messages to the corresponding user mailbox boxes in Microsoft 365.

1. Go to https://compliance.microsoft.com and then click **Data connectors** \ **AT&T Network**.
2. On the **AT&T Network product** description page, click **Add connector**
3. On the **Terms of service** page, click **Accept**.
4. On the **Login to TeleMessage** page, under Step 3, enter the required information in the following boxes and then click **Next**.

- **Username:** Your TeleMessage username.
- **Password:** Your TeleMessage password.

5. After the connector is created, you can close the pop-up window and go to the next page.
6. On the **User mapping** page, enable automatic user mapping. To enable custom mapping, upload a CSV file that contains the user mapping information, and then click **Next**.
7. Review your settings, and then click **Finish** to create the connector.
8. Go to the **Connectors** tab on the **Data connectors** page in the compliance center to see the progress of the import process for the new connector.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.

Set up a connector to archive Bell Network data

2/18/2021 • 4 minutes to read • Edit Online

Use a TeleMessage connector in the Microsoft 365 compliance center to import and archive Short Messaging Service (SMS) and Multimedia Messaging Service (MMS) messages from the Bell Network. After you set up and configure a connector, it connects to your organization's Bell Network once every day, and imports SMS and MMS messages to mailboxes in Microsoft 365.

After the SMS and MMS messages are stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, Content Search, and Microsoft 365 retention policies to Bell Network data. For example, you can search Bell Network SMS/MMS using Content Search or associate the mailbox that contains the Bell Network connector data with a custodian in an Advanced eDiscovery case. Using a Bell Network connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving Bell Network data

The following overview explains the process of using a connector to archive Bell Network data in Microsoft 365.

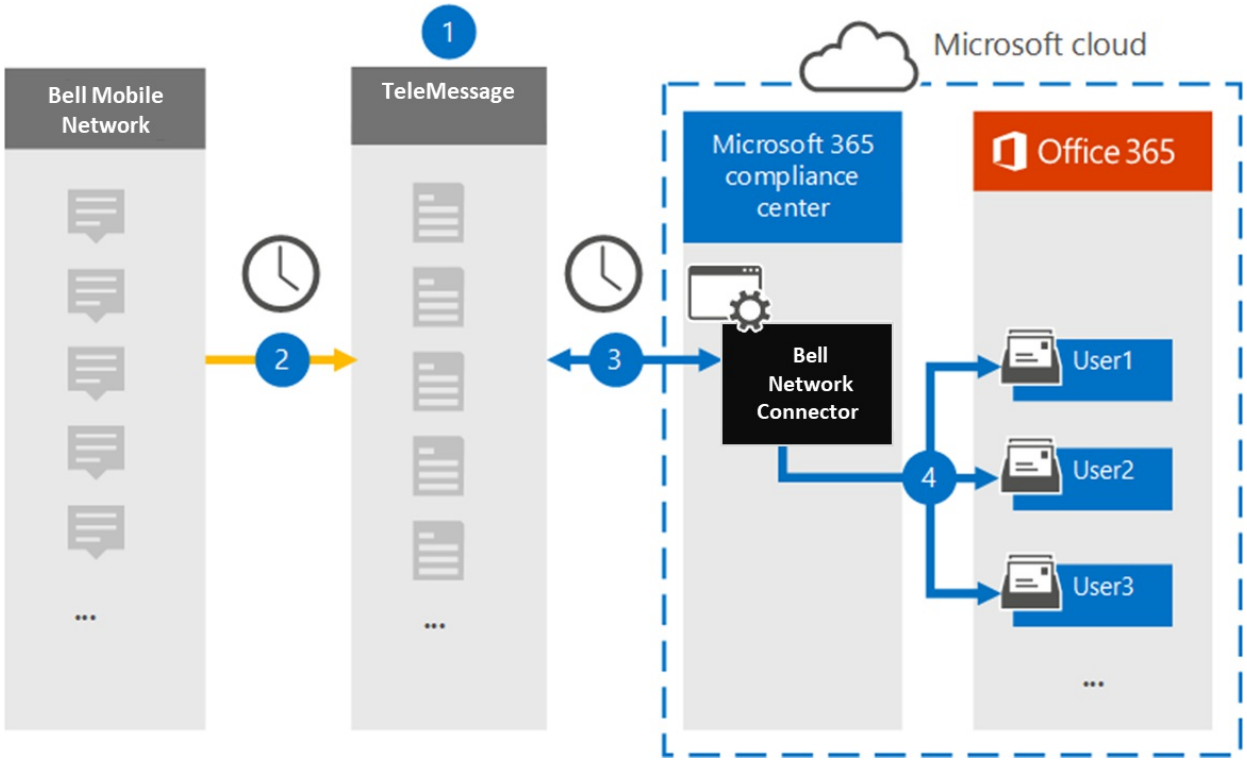

1. Your organization works with TeleMessage and Bell to set up a Bell Network connector. For more information, see Bell Network Archiver.

2. Once every 24 hours, SMS and MMS messages from your organization's Bell Network are copied to the TeleMessage site.

3. The Bell Network connector that you create in the Microsoft 365 compliance center connects to the TeleMessage site every day and transfers the SMS and MMS messages from the previous 24 hours to a secure Azure Storage location in the Microsoft Cloud. The connector also converts the content of SMS and MMS messages to an email message format.

4. The connector imports the mobile communication items to the mailbox of specific users. A new folder

named **Bell SMS/MMS Network Archiver** is created in a specific user's mailbox and the items are imported to it. The connector does this mapping by using the value of the *User's Email address* property. Every SMS and MMS message contains this property, which is populated with the email address of every participant of the message.

In addition to automatic user mapping using the value of the *User's Email address* property, you can also define a custom mapping by uploading a CSV mapping file. This mapping file contains the mobile phone number and corresponding Microsoft 365 email address for users in your organization. If you enable both automatic user mapping and custom mapping, for every Bell Network item the connector first looks at custom mapping file. If it doesn't find a valid Microsoft 365 user that corresponds to a user's mobile phone number, the connector will use the values in the email address property of the item it's trying to import. If the connector doesn't find a valid Microsoft 365 user in either the custom mapping file or in the email address property of the Bell Network item, the item won't be imported.

Before you begin

Some of the implementation steps required to archive Bell Network data are external to Microsoft 365 and must be completed before you can create a connector in the compliance center.

- Order the Bell Network Archiver service from TeleMessage and get a valid administration account for your organization. You'll need to sign into this account when you create the connector in the compliance center.
- Obtain your Bell Network account and billing contact details so you can fill-out the TeleMessage onboarding forms and order the message archiving service from Bell.
- Register all users that require Bell SMS/MMS Network archiving in the TeleMessage account. When registering users, be sure to use the same email address that's used for their Microsoft 365 account.
- Your employees must have corporate-owned and corporate-liable mobile phones on the Bell mobile network. Archiving messages in Microsoft 365 isn't available for employee-owned or "Bring Your Own Devices (BYOD) devices.
- The user who creates a Bell Network connector must be assigned the Mailbox Import Export role in Exchange Online. This is required to add connectors in the **Data connectors** page in the Microsoft 365 compliance center. By default, this role isn't assigned to any role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Create a Bell Network connector

The last step is to create a Bell Network connector in the Microsoft 365 compliance center. The connector uses the information you provide to connect to the TeleMessage site and transfer SMS/ MMS messages to the corresponding user mailbox boxes in Microsoft 365.

1. Go to https://compliance.microsoft.com and then click **Data connectors** > **Bell SMS/MMS Network Archiver**.
2. On the **Bell Network** product description page, click **Add connector**
3. On the **Terms of service** page, click **Accept**.
4. On the **Login to TeleMessage** page, under Step 3, enter the required information in the following boxes and then click **Next**.

- **Username:** Your TeleMessage username.
- **Password:** Your TeleMessage password.

5. After the connector is created, you can close the pop-up window and go to the next page.
6. On the **User mapping** page, enable automatic user mapping. To enable custom mapping, upload a CSV file that contains the user mapping information, and then click **Next**.
7. Review your settings, and then click **Finish** to create the connector.
8. Go to the **Connectors** tab on the **Data connectors** page in the compliance center to see the progress of the import process for the new connector.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.

Set up a connector to archive Enterprise Number data

2/18/2021 • 4 minutes to read • Edit Online

Use a TeleMessage connector in the Microsoft 365 compliance center to import and archive Short Messaging Service (SMS) and Multimedia Messaging Service (MMS) messages, chat messages, voice call recordings, and voice call logs from the Enterprise Number Archiver. After you set up and configure a connector, it connects to your organization's TeleMessage account once every day and imports the mobile communication data of employees using the TeleMessage Enterprise Number Archiver to mailboxes in Microsoft 365.

After the TeleMessage Enterprise Number Archiver connector data is stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, Content Search, In-Place Archiving, Auditing, Communication compliance, and Microsoft 365 retention policies to Enterprise Number Archiver data. For example, you can search the TeleMessage Enterprise Number Archiver SMS, MMS, and Voice Call using Content Search or associate the mailbox that contains the Enterprise Number Archiver connector data with a custodian in an Advanced eDiscovery case. Using an Enterprise Number Archiver connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving Enterprise Number data

The following overview explains the process of using a connector to archive Enterprise Network data in Microsoft 365.

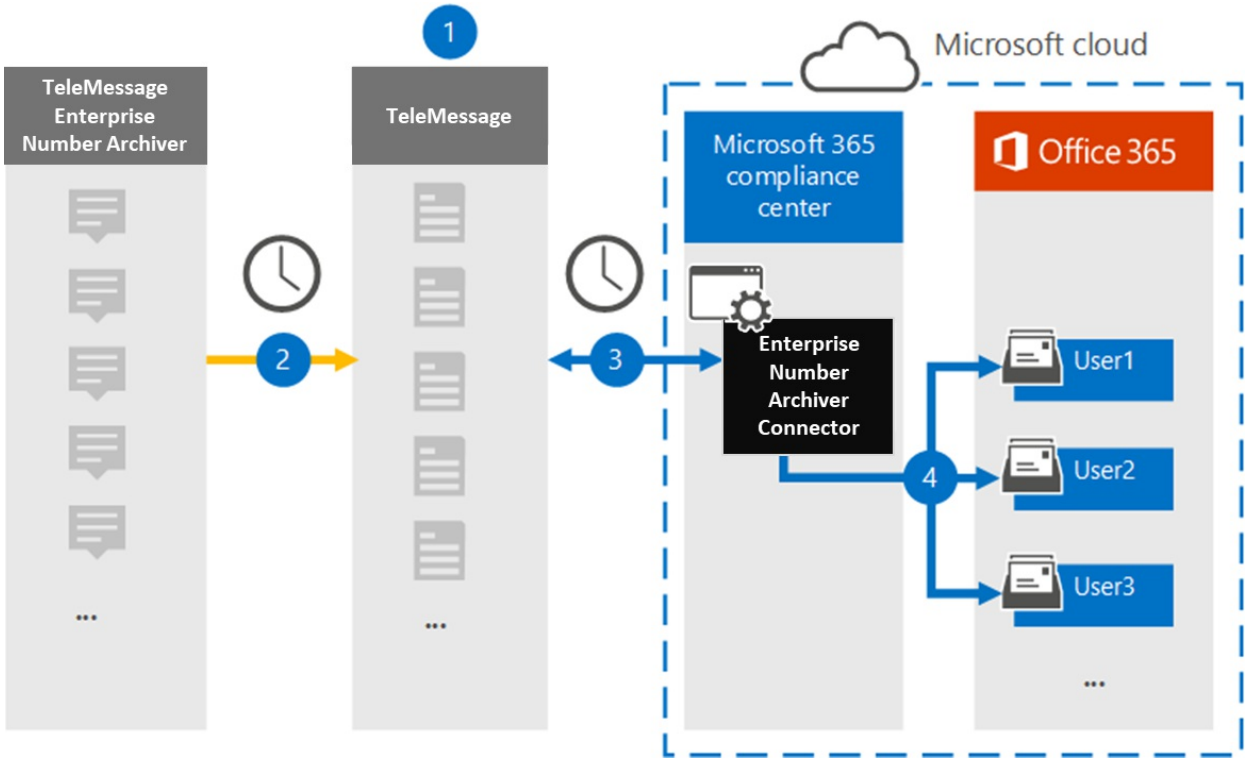

1. Your organization works with TeleMessage to set up an Enterprise Number Archiver connector. For more details refer to here.

2. The Enterprise Number Archiver connector that you create in the Microsoft 365 compliance center connects to the TeleMessage site every day and transfers the email messages from the previous 24 hours to a secure Azure Storage area in the Microsoft Cloud.

3. The connector imports the mobile communication items to the mailbox of a specific user. A new folder named Enterprise Number Archiver is created in the specific user's mailbox and the items are imported to it. The connector does mapping by using the value of the *User's Email address* property. Every email message contains this property, which is populated with the email address of every participant of the email message. In addition to automatic user mapping using the value of the *User's Email address* property, you can also define a custom mapping by uploading a CSV mapping file. This mapping file should contain User's mobile Number and the corresponding Microsoft 365 mailbox address for each user. If you enable automatic user mapping and provide a custom mapping, for every email item the connector will first look at custom mapping file. If it doesn't find a valid Microsoft 365 user that corresponds to a user's mobile number, the connector will use the User 's email address property of the email item. If the connector doesn't find a valid Microsoft 365 user in either the custom mapping file or the *user's email address* property of the email item, the item won't be imported.

Before you begin

Some of the implementation steps required to archive Enterprise Number Archiver data are external to Microsoft 365 and must be completed before you can create the connector in the compliance center.

- Order the Enterprise Number Archiver service from TeleMessage and get a valid administration account for your organization. You'll need to sign into this account when you create the connector in the compliance center.
- Register all users that require Enterprise Number SMS/MMS Network archiving in the TeleMessage account. When registering users, be sure to use the same email address that's used for their Microsoft 365 account.
- Install and activate the TeleMessage Enterprise Number Archiver app on the mobile phones of your employees.
- The user who creates a Enterprise Number Archiver connector must be assigned the Mailbox Import Export role in Exchange Online. This is required to add connectors in the **Data connectors** page in the Microsoft 365 compliance center. By default, this role isn't assigned to any role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Create an Enterprise Number Archiver connector

After you've completed the prerequisites described in the previous section, you can create an Enterprise Number Archiver connector in the Microsoft 365 compliance center. The connector uses the information you provide to connect to the TeleMessage site and transfer SMS, MMS, and voice call messages to the corresponding user mailbox boxes in Microsoft 365.

1. Go to https://compliance.microsoft.com and then click **Data connectors** > **Enterprise Number Archiver**.
2. On the **Enterprise Number Archiver** product description page, click **Add connector**
3. On the **Terms of service** page, click **Accept**.
4. On the **Login to TeleMessage** page, under Step 3, enter the required information in the following boxes and then click **Next**.
   - **Username:** Your TeleMessage username.
   - **Password:** Your TeleMessage password.

5. After the connector is created, you can close the pop-up window and go to the next page.
6. On the **User mapping** page, enable automatic user mapping. To enable custom mapping, upload a CSV file that contains the user mapping information, and then click **Next**.
7. Review your settings, and then click **Finish** to create the connector.
8. Go to the Connectors tab in **Data connectors** page to see the progress of the import process for the new connector.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.

Set up a connector to archive O2 Network data

2/18/2021 • 4 minutes to read • Edit Online

Use a TeleMessage connector in the Microsoft 365 compliance center to import and archive Short Messaging Service (SMS) messages and voice calls from the O2 mobile network. After you set up and configure a connector, it connects to your organization's O2 Network once every day, and imports SMS and voice calls to mailboxes in Microsoft 365.

After SMS messages and voice calls are stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, Content Search, and Microsoft 365 retention policies to O2 Network data. For example, you can search O2 Network SMS messages and voice calls using Content Search or associate the mailbox that contains O2 Network data with a custodian in an Advanced eDiscovery case. Using an O2 Network connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving O2 Network data

The following overview explains the process of using a connector to archive O2 Network data in Microsoft 365.

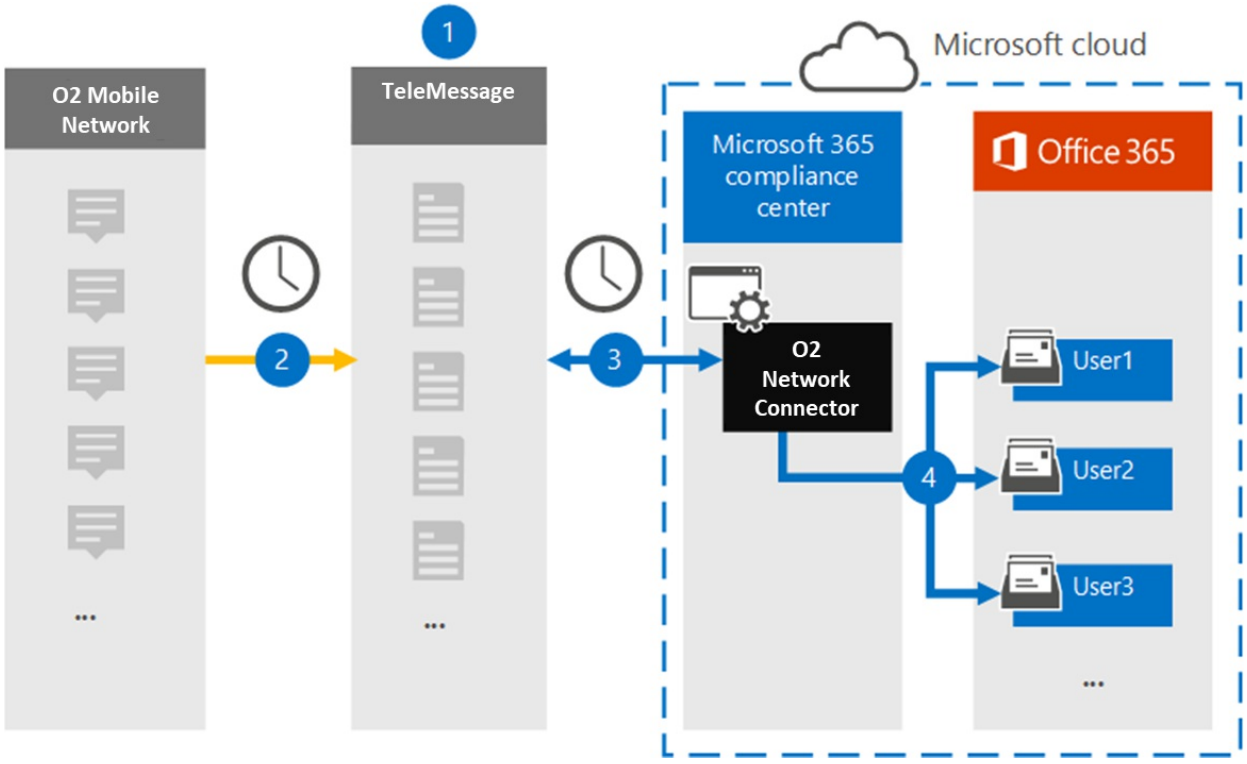

1. Your organization works with TeleMessage and O2 to set up an O2 Network connector. For more information, see O2 Network Archiver.

2. Once every 24 hours, SMS messages and voice calls from your organization's O2 Network are copied to the TeleMessage site.

3. The O2 Network connector that you create in the Microsoft 365 compliance center connects to the TeleMessage site every day and transfers the SMS messages and voice calls from the previous 24 hours to a secure Azure Storage location in the Microsoft Cloud. The connector also converts the content of SMS messages and voice calls to an email message format.

4. The connector imports the mobile communication items to the mailbox of specific users. A new folder

named **O2 SMS and Voice Network Archiver** is created in a specific user's mailbox and the items are imported to it. The connector does this mapping by using the value of the *User's Email address* property. Every SMS message and voice call contains this property, which is populated with the email address of every participant of the message.

In addition to automatic user mapping using the value of the *User's Email address* property, you can also define a custom mapping by uploading a CSV mapping file. This mapping file contains the mobile phone number and corresponding Microsoft 365 email address for users in your organization. If you enable both automatic user mapping and custom mapping, for every O2 item the connector first looks at custom mapping file. If it doesn't find a valid Microsoft 365 user that corresponds to a user's mobile phone number, the connector will use the values in the email address property of the item it's trying to import. If the connector doesn't find a valid Microsoft 365 user in either the custom mapping file or in the email address property of the O2 item, the item won't be imported.

Before you begin

Some of the implementation steps required to archive O2 Network data are external to Microsoft 365 and must be completed before you can create a connector in the compliance center.

- Order the O2 Network Archiver service from TeleMessage and get a valid administration account for your organization. You'll need to sign into this account when you create the connector in the compliance center.
- Obtain your O2 Network account and billing contact details so you can fill-out the TeleMessage onboarding forms and order the message archiving service from O2.
- Register all users that require O2 SMS and Voice Network archiving in the TeleMessage account. When registering users, be sure to use the same email address that's used for their Microsoft 365 account.
- Your employees must have corporate-owned and corporate-liable mobile phones on the O2 mobile network. Archiving messages in Microsoft 365 isn't available for employee-owned or "Bring Your Own Devices (BYOD) devices.
- The user who creates an O2 Network connector must be assigned the Mailbox Import Export role in Exchange Online. This is required to add connectors in the **Data connectors** page in the Microsoft 365 compliance center. By default, this role isn't assigned to any role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Create an O2 Network connector

After you've completed the prerequisites described in the previous section, you can create an O2 Network connector in the Microsoft 365 compliance center. The connector uses the information you provide to connect to the TeleMessage site and transfer SMS messages and voice calls to the corresponding user mailbox boxes in Microsoft 365.

1. Go to https://compliance.microsoft.com and then click **Data connectors** > **O2 Network**.
2. On the **O2 Network** product description page, click **Add connector**
3. On the **Terms of service** page, click **Accept**.
4. On the **Login to TeleMessage** page, under Step 3, enter the required information in the following boxes and then click **Next**.
   - **Username:** Your TeleMessage username.

- **Password:** Your TeleMessage password.

5. After the connector is created, you can close the pop-up window and go to the next page.
6. On the **User mapping** page, enable automatic user mapping and click **Next**. In case you need custom mapping upload a CSV file, and click **Next**.
7. Review your settings, and then click **Finish** to create the connector.
8. Go to the Connectors tab in **Data connectors** page to see the progress of the import process for the new connector.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.

Set up a connector to archive TELUS Network data

2/18/2021 • 4 minutes to read • Edit Online

Use the TeleMessage connector in the Microsoft 365 compliance center to import and archive Short Messaging Service (SMS) data from your organization's TELUS Network. After you set up and configure a connector, it connects to your organization's TELUS Network once every day, and imports SMS data to mailboxes in Microsoft 365.

After SMS messages are stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, Content Search, and Microsoft 365 retention policies to TELUS data. For example, you can search TELUS SMS messages using Content Search or associate the mailbox that contains the TELUS data with a custodian in an Advanced eDiscovery case. Using a TELUS Network connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving TELUS Network data

The following overview explains the process of using a connector to archive TELUS Network data in Microsoft 365.

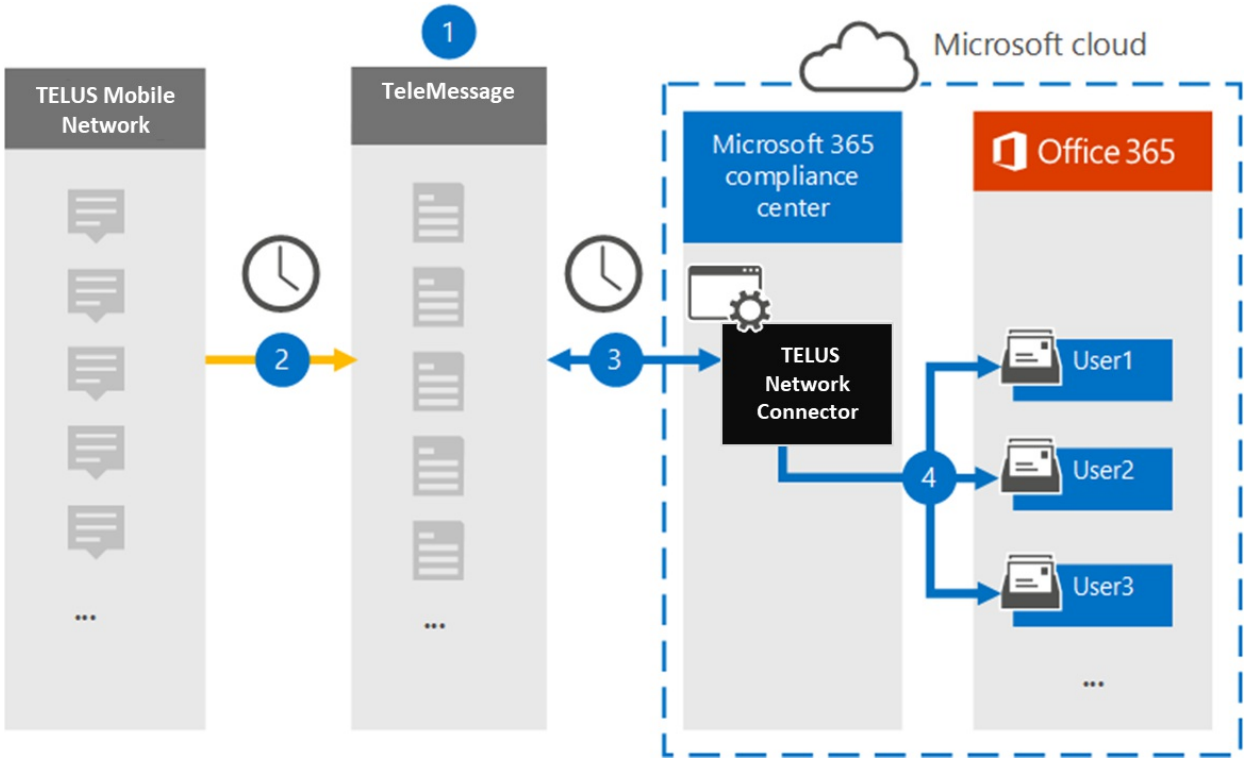

1. Your organization works with TeleMessage and TELUS to set up a TELUS Network connector. For more information, see TELUS Network Archiver.
2. Once every 24 hours, SMS messages from your organization's TELUS Network are copied to the TeleMessage site.
3. The TELUS Network connector that you create in the Microsoft 365 compliance center connects to the TeleMessage site every day and transfers the SMS messages from the previous 24 hours to a secure Azure Storage location in the Microsoft Cloud. The connector also converts the content of SMS messages to an email message format.
4. The connector imports the mobile communication items to the mailbox of a specific user. A new folder

named **TELUS SMS Network Archiver** is created in the specific user's mailbox and the items are imported to it. The connector does mapping by using the value of the *User's Email address* property. Every SMS message contains this property, which is populated with the email address of every participant of the SMS message.

In addition to automatic user mapping using the value of the *User's Email address* property, you can also implement custom mapping by uploading a CSV mapping file. This mapping file contains the mobile phone number and corresponding Microsoft 365 email address for users in your organization. If you enable both automatic user mapping and custom mapping, for every TELUS item the connector first looks at custom mapping file. If it doesn't find a valid Microsoft 365 user that corresponds to a user's mobile phone number, the connector will use the values in the email address property of the item it's trying to import. If the connector doesn't find a valid Microsoft 365 user in either the custom mapping file or in the email address property of the TELUS item, the item won't be imported.

Before you begin

Some of the implementation steps required to archive TELUS Network data are external to Microsoft 365 and must be completed before you can create a connector in the compliance center.

- Order the TELUS Network Archiver service from TeleMessage and get a valid administration account for your organization. You'll need to sign into this account when you create the connector in the compliance center.
- Obtain your TELUS Network account and billing contact details so you can fill-out the TeleMessage onboarding forms and order the message archiving service from TELUS.
- Register all users that require TELUS SMS Network archiving in the TeleMessage account. When registering users, be sure to use the same email address that's used for their Microsoft 365 account.
- Your employees must have corporate-owned and corporate-liable mobile phones on theTELUS mobile network. Archiving messages in Microsoft 365 isn't available for employee-owned or Bring Your Own Devices (BYOD) devices.
- The user who creates a TELUS Network connector must be assigned the Mailbox Import Export role in Exchange Online. This is required to add connectors in the **Data connectors** page in the Microsoft 365 compliance center. By default, this role isn't assigned to any role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Create a TELUS Network connector

After you've completed the prerequisites described in the previous section, you can create TELUS Network connector in the Microsoft 365 compliance center. The connector uses the information you provide to connect to the TeleMessage site and transfer SMS messages to the corresponding user mailbox boxes in Microsoft 365.

1. Go to https://compliance.microsoft.com and then click **Data connectors** > **TELUS Network**.
2. On the **TELUS Network** product description page, click **Add connector**
3. On the **Terms of service** page, click **Accept**.
4. On the **Login to TeleMessage** page, under Step 3, enter the required information in the following boxes and then click **Next**.
   - **Username:** Your TeleMessage username.

- **Password:** Your TeleMessage password.

5. After the connector is created, you can close the pop-up window and go to the next page.
6. On the **User mapping** page, enable automatic user mapping and click **Next**. In case you need custom mapping upload a CSV file, and click **Next**.
7. Review your settings, and then click **Finish** to create the connector.
8. Go to the Connectors tab in **Data connectors** page to see the progress of the import process for the new connector.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.

Set up a connector to archive Verizon Network data

2/18/2021 • 4 minutes to read • Edit Online

Use the TeleMessage connector in the Microsoft 365 compliance center to import and archive Short Messaging Service (SMS) and Multimedia Messaging Service (MMS) data from Verizon Network. After you set up and configure a connector, it connects to your organization's Verizon Network once every day and imports SMS and MMS data to mailboxes in Microsoft 365.

After Verizon Network connector data is stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, Content Search, and Microsoft 365 retention policies to Verizon data. For example, you can search Verizon SMS and MMS messages using Content Search or associate the mailbox that contains Verizon Network data with a custodian in an Advanced eDiscovery case. Using a Verizon Network connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving Verizon Network data

The following overview explains the process of using a connector to archive Verizon Network data in Microsoft 365.

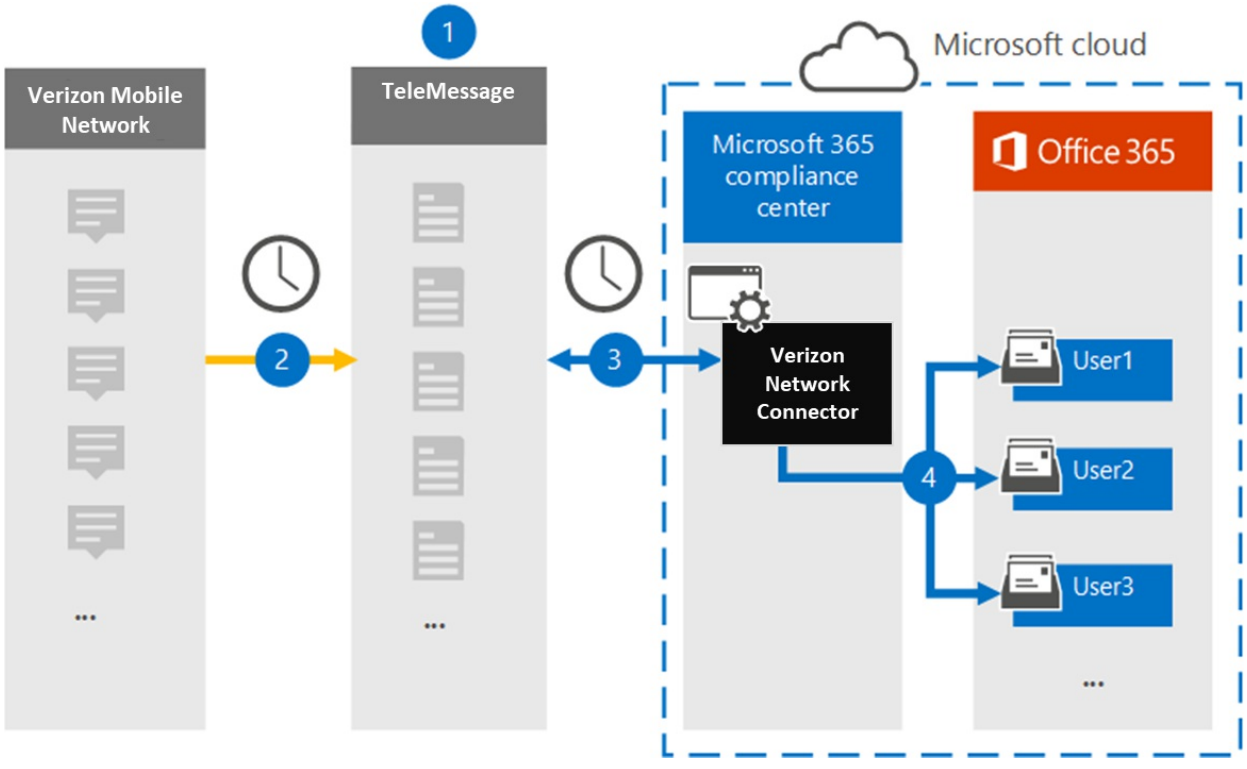

1. Your organization works with TeleMessage and Verizon to set up a Verizon Network connector. For more information, see Verizon Network Archiver.

2. Once every 24 hours, SMS and MMS messages from your organization's Verizon Network are copied to the TeleMessage site.

3. The Verizon Network connector that you create in the Microsoft 365 compliance center connects to the TeleMessage site every day and transfers the SMS and MMS messages from the previous 24 hours to a secure Azure Storage location in the Microsoft Cloud. The connector also converts the content of SMS and MMS messages to an email message format.

4. The connector imports the mobile communication items to the mailbox of a specific user. A new folder named **Verizon SMS/MMS Network Archiver** is created in the specific user's mailbox and the items are imported to it. The connector does this mapping by using the value of the *User's Email address* property. Every SMS and MMS message contains this property, which is populated with the email address of every participant of the message.

   In addition to automatic user mapping using the value of the *User's Email address* property, you can also implement custom mapping by uploading a CSV mapping file. This mapping file contains the mobile phone number and corresponding Microsoft 365 email address for users in your organization. If you enable both automatic user mapping and custom mapping, for every Verizon item the connector first looks at custom mapping file. If it doesn't find a valid Microsoft 365 user that corresponds to a user's mobile phone number, the connector will use the values in the email address property of the item it's trying to import. If the connector doesn't find a valid Microsoft 365 user in either the custom mapping file or in the email address property of the Verizon item, the item won't be imported.

Before you begin

Some of the implementation steps required to archive Verizon Network data are external to Microsoft 365 and must be completed before you can create a connector in the compliance center.

- Order the Verizon Network Archiver service from TeleMessage and get a valid administration account for your organization. You'll need to sign into this account when you create the connector in the compliance center.
- Obtain your Verizon Network account and billing contact details so you can fill-out the TeleMessage onboarding forms and order the message archiving service from Verizon.
- Register all users that require Verizon SMS and MMS archiving in the TeleMessage account. When registering users, be sure to use the same email address that's used for their Microsoft 365 account.
- Your employees must have corporate-owned and corporate-liable mobile phones on the Verizon mobile network. Archiving messages in Microsoft 365 isn't available for employee-owned or Bring Your Own Devices (BYOD) devices.
- The user who creates a Verizon Network connector must be assigned the Mailbox Import Export role in Exchange Online. This is required to add connectors in the **Data connectors** page in the Microsoft 365 compliance center. By default, this role isn't assigned to any role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Create a Verizon Network connector

After you've completed the prerequisites described in the previous section, you can create Verizon Network connector in the Microsoft 365 compliance center. The connector uses the information you provide to connect to the TeleMessage site and transfer SMS and MMS messages to the corresponding user mailbox boxes in Microsoft 365.

1. Go to https://compliance.microsoft.com and then click **Data connectors** > **Verizon Network**.
2. On the **Verizon Network** product description page, click **Add connector**
3. On the **Terms of service** page, click **Accept**.
4. On the **Login to TeleMessage** page, under Step 3, enter the required information in the following boxes and then click **Next**.

- **Username:** Your TeleMessage username.
- **Password:** Your TeleMessage password.

5. After the connector is created, you can close the pop-up window and go to the next page.
6. On the **User mapping** page, enable automatic user mapping and click **Next**. In case you need custom mapping upload a CSV file, and click **Next**.
7. Review your settings, and then click **Finish** to create the connector.
8. Go to the Connectors tab in **Data connectors** page to see the progress of the import process for the new connector.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.

Set up a connector to archive WhatsApp data

2/18/2021 • 4 minutes to read • Edit Online

Use the TeleMessage connector in the Microsoft 365 compliance center to import and archive WhatsApp calls, chats, attachments, files, and deleted messages. After you set up and configure a connector, it connects to your organization's TeleMessage account once every day, and imports the mobile communication of employees using the TeleMessage WhatsApp Phone Archiver or TeleMessage WhatsApp Cloud Archiver to mailboxes in Microsoft 365.

After WhatsApp data is stored in user mailboxes, you can apply Microsoft 365 compliance features such as Litigation Hold, Content Search, and Microsoft 365 retention policies to WhatsApp data. For example, you can search WhatsApp messages using Content Search or associate the mailbox that contains WhatsApp messages with a custodian in an Advanced eDiscovery case. Using a WhatsApp connector to import and archive data in Microsoft 365 can help your organization stay compliant with government and regulatory policies.

Overview of archiving WhatsApp data

The following overview explains the process of using a connector to archive WhatsApp data in Microsoft 365.

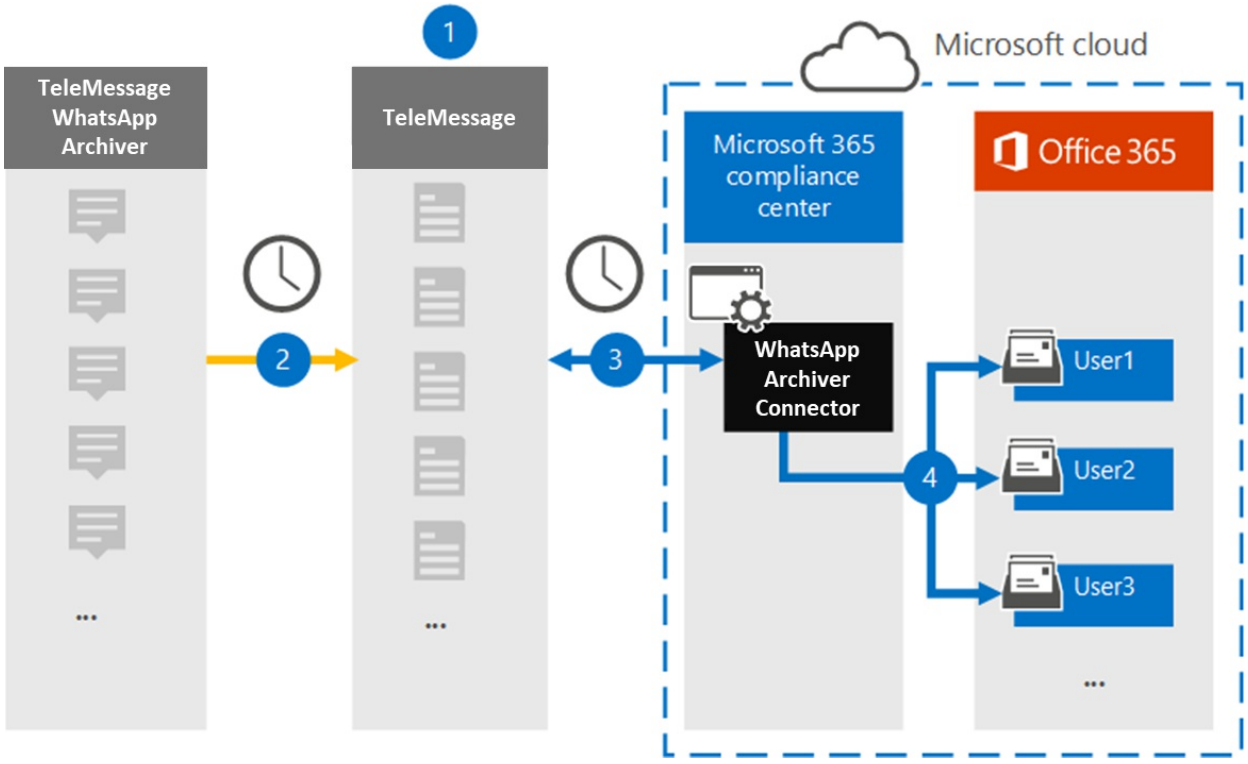

1. Your organization works with TeleMessage to set up a WhatsApp Archiver connector. For more information, see WhatsApp Archiver.

2. Once every 24 hours, your organization's WhatsApp data is copied to the TeleMessage site.

3. The WhatsApp connector that you create in the Microsoft 365 compliance center connects to the TeleMessage site every day and transfers WhatsApp data from the previous 24 hours to a secure Azure Storage location in the Microsoft Cloud. The connector also converts the content WhatsApp data to an email message format.

4. The connector imports WhatsApp data to the mailbox of a specific user. A new folder named **WhatsApp Archiver** is created in the specific user's mailbox and the items are imported to it. The connector does

this mapping by using the value of the *User's Email address* property. Every WhatsApp message contains this property, which is populated with the email address of every participant of the message.

In addition to automatic user mapping using the value of the *User's Email address* property, you can also implement custom mapping by uploading a CSV mapping file. This mapping file contains the mobile phone number and corresponding Microsoft 365 email address for users in your organization. If you enable both automatic user mapping and custom mapping, for every WhatsApp item the connector first looks at custom mapping file. If it doesn't find a valid Microsoft 365 user that corresponds to a user's mobile phone number, the connector will use the values in the email address property of the item it's trying to import. If the connector doesn't find a valid Microsoft 365 user in either the custom mapping file or in the email address property of the WhatsApp item, the item won't be imported.

Before you begin

Some of the implementation steps required to archive WhatsApp communication data are external to Microsoft 365 and must be completed before you can create the connector in the compliance center.

- Order the WhatsApp Archiver service from TeleMessage and get a valid administration account for your organization. You'll need to sign into this account when you create the connector in the compliance center.
- Register all users that require WhatsApp archiving in the TeleMessage account. When registering users, be sure to use the same email address that's used for their Microsoft 365 account.
- Install the TeleMessage WhatsApp Phone Archiver app on the mobile phones of your employees and activate it. Alternatively, you can install the regular WhatsApp or WhatsApp Business apps on the mobile phones of your employees and activate the WhatsApp Cloud Archiver service by scanning a QR code on the TeleMessage website. For more information, see WhatsApp Cloud Archiver.
- The user who creates a Verizon Network connector must be assigned the Mailbox Import Export role in Exchange Online. This is required to add connectors in the **Data connectors** page in the Microsoft 365 compliance center. By default, this role isn't assigned to any role group in Exchange Online. You can add the Mailbox Import Export role to the Organization Management role group in Exchange Online. Or you can create a role group, assign the Mailbox Import Export role, and then add the appropriate users as members. For more information, see the Create role groups or Modify role groups sections in the article "Manage role groups in Exchange Online".

Create a WhatsApp Archiver connector

After you've completed the prerequisites described in the previous section, you can create the WhatsApp connector in the Microsoft 365 compliance center. The connector uses the information you provide to connect to the TeleMessage site and transfer the WhatsApp data to the corresponding user mailbox boxes in Microsoft 365.

1. Go to https://compliance.microsoft.com and then click **Data connectors** > **WhatsApp Archiver**.
2. On the **WhatsApp Archiver** product description page, click **Add connector**
3. On the **Terms of service** page, click **Accept**.
4. On the **Login to TeleMessage** page, under Step 3, enter the required information in the following boxes and then click **Next**.
   - **Username:** Your TeleMessage username.
   - **Password:** Your TeleMessage password.
5. After the connector is created, you can close the pop-up window and go to the next page.
6. On the **User mapping** page, enable automatic user mapping and click **Next**. In case you need custom

mapping upload a CSV file, and click **Next**.

7. Review your settings, and then click **Finish** to create the connector.
8. Go to the Connectors tab in **Data connectors** page to see the progress of the import process for the new connector.

Known issues

- At this time, we don't support importing attachments or items that are larger than 10 MB. Support for larger items will be available at a later date.

Work with a partner to archive third-party data

2/18/2021 • 16 minutes to read • Edit Online

You can work with a Microsoft Partner to import and archive data from a third-party data source to Microsoft 365. A partner can provide you with a custom connector that is configured to extract items from the third-party data source (on a regular basis) and then import those items. The partner connector converts the content of an item from the data source to an email message format and then stores the items in mailboxes. After third-party data is imported, you can apply Microsoft 365 compliance features such as Litigation Hold, eDiscovery, In-Place Archiving, Auditing, and Microsoft 365 retention policies to this data.

> **IMPORTANT**
>
> The Communication compliance solution in Microsoft 365 can't be applied to the third-party data imported by partner connectors mentioned in this article.

Here's an overview of the process and the steps necessary to work with a Microsoft Partner to import third-party data.

Step 1: Find a third-party data partner

Step 2: Create and configure a third-party data mailbox

Step 3: Configure user mailboxes for third-party data

Step 4: Provide your partner with information

Step 5: Register the third-party data connector in Azure Active Directory

How the third-party data import process works

The following illustration and description explain how the third-party data import process works when working with a partner.

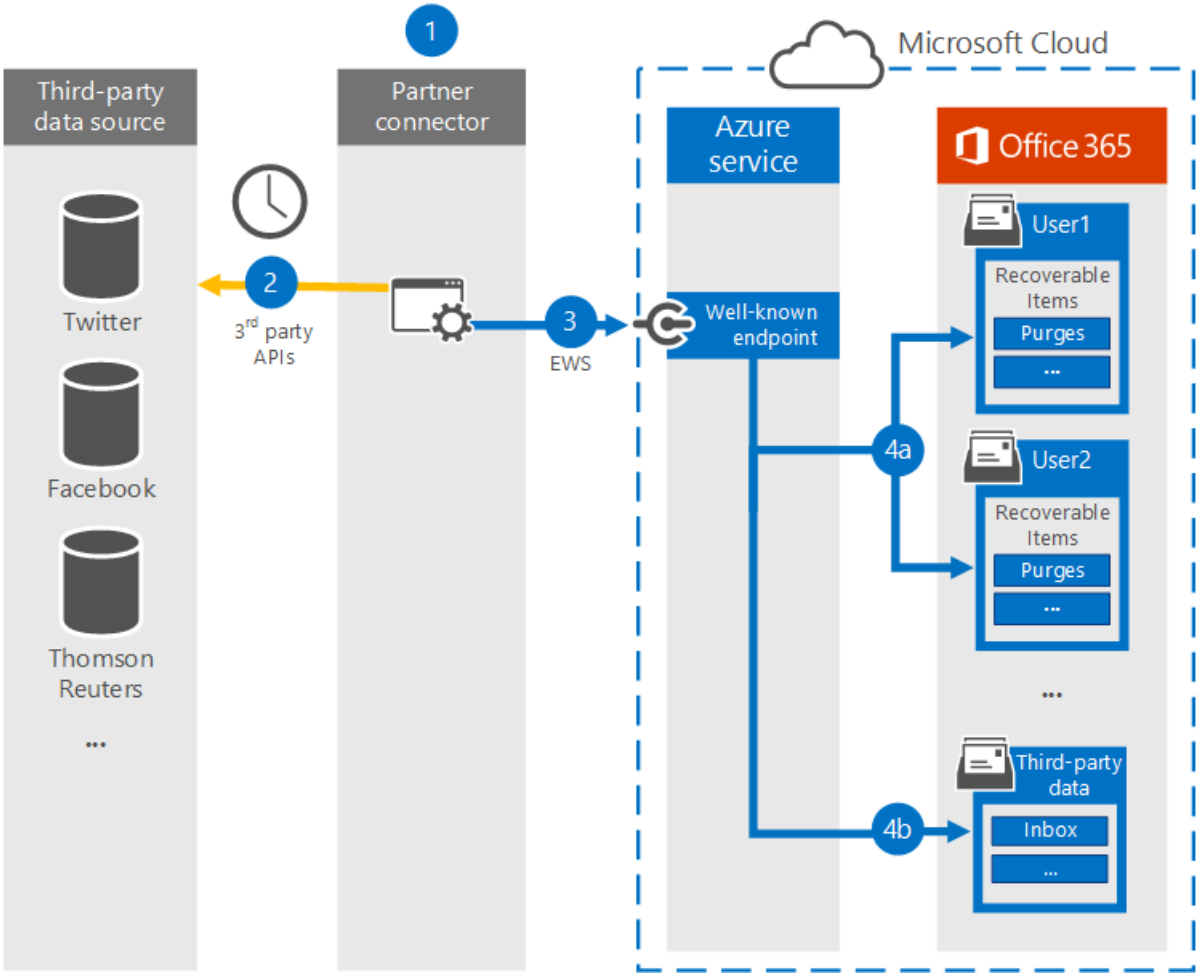

1. Customer works with their partner of choice to configure a connector that will extract items from the third-party data source and then import those items to Microsoft 365.
2. The partner connector connects to third-party data sources via a third-party API (on a scheduled or as-configured basis) and extracts items from the data source. The partner connector converts the content of an item to an email message format. See the More information section for a description of the message-format schema.
3. Partner connector connects to the Azure service in Microsoft 365 by using Exchange Web Service (EWS) via a well-known end point.
4. Items are imported into the mailbox of a specific user or into a "catch-all" third-party data mailbox. Whether an item is imported into a specific user mailbox or to the third-party data mailbox is based on the following criteria:
    a. **Items that have a user ID that corresponds to a user account:** If the partner connector can map the user ID of the item in the third-party data source to a specific user ID in Microsoft 365, the item is copied to the **Purges** folder in the user's Recoverable Items folder. Users can't access items in the Purges folder. However, you can use eDiscovery tools to search for items in the Purges folder.
    b. **Items that don't have a user ID that corresponds to a user account:** If the partner connector can't map the user ID of an item to a specific user ID, the item is copied to the **Inbox** folder of the third-party data mailbox. Importing items to the inbox allows you or someone in your organization to sign in to the third-party mailbox to view and manage these items, and see if any adjustments need to be made in the partner connector configuration.

Step 1: Find a third-party data partner

A key component for archiving third-party data in Microsoft 365 is finding and working with a Microsoft partner that specializes in capturing data from a third-party data source and importing it to Microsoft 365. After the data is imported, it can be archived and preserved along with your organization's other Microsoft data, such as email from Exchange and documents from SharePoint and OneDrive for Business. A partner creates a connector that extracts data from your organization's third-party data sources (such as BlackBerry, Facebook, Google+, Thomson Reuters, Twitter, and YouTube) and passes that data to a Microsoft 365 API that imports items to Exchange mailboxes as email messages.

The following sections list the Microsoft partners (and the third-party data sources they support) that are participating in the program for archiving third-party data in Microsoft 365.

17a-4 LLC

17a-4 LLC supports the following third-party data sources:

- BlackBerry
- Bloomberg Data Streams
- Cisco Jabber
- FactSet
- HipChat
- InvestEdge
- LivePerson
- MessageLabs Data Streams
- OpenText
- Oracle/ATG 'click-to-call' Live Help
- Pivot IMTRADER
- Microsoft SharePoint
- MindAlign
- Sitrion One (Newsgator)
- Skype for Business (Lync/OCS)
- Skype for Business Online (Lync Online)
- SQL Databases
- Squawker
- Thomson Reuters Eikon Messenger

ArchiveSocial

ArchiveSocial supports the following third-party data sources:

- Facebook
- Flickr
- Instagram
- LinkedIn
- Pinterest
- Twitter
- YouTube
- Vimeo

Globanet

Globanet supports the following third-party data sources:

- AOL with Pivot Client
- BlackBerry Call Logs (v5, v10, v12)
- BlackBerry Messenger (v5, v10, v12)
- BlackBerry PIN (v5, v10, v12)
- BlackBerry SMS (v5, v10, v12)
- Bloomberg Chat
- Bloomberg Mail
- Box
- CipherCloud for Salesforce Chatter
- Cisco IM & Presence Server (v10, v10.5.1 SU1, v11.0, v11.5 SU2)
- Cisco Webex Teams
- Citrix Workspace & ShareFile
- CrowdCompass
- Custom-delimited text files
- Custom XML files
- Facebook (Pages)
- Factset
- FXConnect
- ICE Chat/YellowJacket
- Jive
- Macgregor XIP
- Microsoft Exchange Server

- Microsoft OneDrive for Business
- Microsoft Teams
- Microsoft Yammer
- Mobile Guard
- Pivot
- Salesforce Chatter
- Skype for Business Online
- Skype for Business, versions 2007 R2 - 2016 (on-premises)
- Slack Enterprise Grid
- Symphony
- Thomson Reuters Eikon
- Thomson Reuters Messenger
- Thomson Reuters Dealings 3000 / FX Trading
- Twitter
- UBS Chat
- YouTube

OpenText

OpenText supports the following third-party data sources:

- Axs Encrypted
- Axs Exchange
- Axs Local Archive
- Axs PlaceHolder
- Axs Signed
- Bloomberg
- Thomson Reuters

Smarsh

Smarsh supports the following third-party data sources:

- AIM
- American Idol
- Apple Juice
- AOL with Pivot client
- Ares
- Bazaar Voice
- Bear Share

- Bit Torrent
- BlackBerry Call Logs (v5, v10, v12)
- BlackBerry Messenger (v5, v10, v12)
- BlackBerry PIN (v5, v10, v12)
- BlackBerry SMS (v5, v10, v12)
- Bloomberg Mail
- CellTrust
- Chat Import
- Chat Real Time Logging and Policy
- Chatter
- Cisco IM & Presence Server (v9.0.1, v9.1, v9.1.1 SU1, v10, v10.5.1 SU1)
- Cisco Unified Presence Server (v8.6.3, v8.6.4, v8.6.5)
- Collaboration Import
- Collaboration Real Time Logging
- Direct Connect
- Facebook
- FactSet
- FastTrack
- Gnutella
- Google+
- GoToMyPC
- Hopster
- HubConnex
- IBM Connections (v3.0.1, v4.0, v4.5, v4.5 CR3, v5)
- IBM Connections Chat Cloud
- IBM Connections Social Cloud
- IBM SameTime Advanced 8.5.2 IFR1
- IBM SameTime Communicate 9.0
- IBM SameTime Community (v8.0.2, v8.5.1 IFR2, v8.5.2 IFR1, v9.1)
- IBM SameTime Complete 9.0
- IBM SameTime Conference 9.0
- IBM SameTime Meeting 8.5.2 IFR1
- ICE/YellowJacket
- IM Import

- IM Real Time Logging and Policy
- Indii Messenger
- Instant Bloomberg
- IRC
- Jive
- Jive 6 Real Time Logging (v6, v7)
- Jive Import
- JXTA
- LinkedIn
- Microsoft Lync (2010, 2013)
- MFTP
- Microsoft Lync 2013 Voice
- Microsoft SharePoint (2010, 2013)
- Microsoft SharePoint Online
- Microsoft UC (Unified Communications)
- MindAlign
- Mobile Guard
- MSN
- My Space
- NEONetwork
- Microsoft 365 Lync Dedicated
- Microsoft 365 Shared IM
- Pinterest
- Pivot
- QQ
- Skype for Business 2015
- SoftEther
- Symphony
- Thomson Reuters Eikon
- Thomson Reuters Messenger
- Tor
- TTT
- Twitter

- WinMX
- Winny
- Yahoo
- Yammer
- YouTube

Verba

Verba supports the following third-party data sources:

- Avaya Aura Video
- Avaya Aura Voice
- Avtec Radio
- Bosch/Telex Radio
- BroadSoft Video
- BroadSoft Voice
- Centile Voice
- Cisco Jabber IM
- Cisco UC Video
- Cisco UC Voice
- Cisco UCCX/UCCE Video
- Cisco UCCX/UCCE Voice
- ESChat Radio
- Geoman Contact Expert
- IP Trade Voice
- Luware LUCS Contact Center
- Microsoft UC (Unified Communications)
- Mitel MiContact Center for Lync (prairieFyre)
- Oracle / Acme Packet Session Border Controller Video
- Oracle / Acme Packet Session Border Controller Voice
- Singtel Mobile Voice
- SIPREC Video
- SIPREC Voice
- Skype for Business / Lync IM
- Skype for Business / Lync Video
- Skype for Business / Lync Voice
- Speakerbus Voice

- Standard SIP/H.323 Video
- Standard SIP/H.323 Voice
- Truphone Voice
- TwistedPair Radio
- Windows Desktop Computer Screen

Step 2: Create and configure a third-party data mailbox in Microsoft 365

Here are the steps for creating and configuring a third-party data mailbox for importing data to Microsoft 365. As previous explained, items are imported to this mailbox if the partner connector can't map the user ID of the item to a user account.

Complete these tasks in the Microsoft 365 admin center

1. Create a user account and assign it an Exchange Online Plan 2 license; see Add users to Microsoft 365. A Plan 2 license is required to place the mailbox on Litigation Hold or enable an archive mailbox that has an unlimited storage quota.
2. Add the user account for the third-party data mailbox to the **Exchange administrator** admin role in Microsoft 365; see Assign admin roles in Microsoft 365.

   > **TIP**
   > Write down the credentials for this user account. You need to provide them to your partner, as described in Step 4.

Complete these tasks in the Exchange admin center

1. Hide the third-party data mailbox from the address book and other address lists in your organization; see Manage user mailboxes. Alternatively, you can run the following PowerShell command:

   ```
   Set-Mailbox -Identity <identity of third-party data mailbox> -HiddenFromAddressListsEnabled $true
   ```

2. Assign the **FullAccess** permission to the third-party data mailbox so that administrators or compliance officers can open the third-party data mailbox in the Outlook desktop client; see Manage permissions for recipients.
3. Enable the following compliance-related features for the third-party data mailbox:
   - Enable the archive mailbox; see Enable archive mailboxes and Enable unlimited archiving. This lets you free-up storage space in the primary mailbox by setting up an archive policy that moves third-party data items to the archive mailbox. This provides you with unlimited storage for third-party data.
   - Place the third-party data mailbox on Litigation Hold. You can also apply a Microsoft 365 retention policy in the security and compliance center. Placing this mailbox on hold retains third-party data items (indefinitely or for a specified duration) and prevent them from being purged from the mailbox. See one of the following topics:
     - Place a mailbox on Litigation Hold
     - Learn about retention policies and retention labels

- Enable mailbox audit logging for owner, delegate, and admin access to the third-party data mailbox; see Enable mailbox auditing. This allows you to audit all activity performed by any user who has access to the third-party data mailbox.

Step 3: Configure user mailboxes for third-party data

The next step is to configure user mailboxes to support third-party data. Complete these tasks by using the Exchange admin center or by using the corresponding Windows PowerShell cmdlets.

1. Enable the archive mailbox for each user; see Enable archive mailboxes and Enable unlimited archiving.
2. Place user mailboxes on Litigation Hold or apply a Microsoft 365 retention policy; see one of the following topics:

   - Place a mailbox on Litigation Hold
   - Learn about retention policies and retention labels

   As previously stated, when you place mailboxes on hold, you can set a duration for how long to hold items from the third-party data source or you can choose to hold items indefinitely.

Step 4: Provide your partner with information

The final step is to provide your partner with the following information so they can configure the connector to connect to your organization to import data to user mailboxes and to the third-party data mailbox.

- The endpoint used to connect to the Azure service in Microsoft 365:

  ```
  https://office365ingestionsvc.gble1.protection.outlook.com/service/ThirdPartyIngestionService.svc
  ```

- The sign-in credentials (Microsoft 365 user ID and password) of the third-party data mailbox that you created in Step 2. These credentials are required so that the partner connector can access and import items to user mailboxes and to the third-party data mailbox.

Step 5: Register the third-party data connector in Azure Active Directory

Starting September 30, 2018, the Azure service in Microsoft 365 will begin using modern authentication in Exchange Online to authenticate third-party data connectors that attempt to connect to your organization to import data. The reason for this change is that modern authentication provides more security than the current method, which was based on an allow list for third-party connectors that use the previously described endpoint to connect to the Azure service.

To enable a third-party data connector to connect to Microsoft 365 using the new modern authentication method, an administrator in your organization must consent to register the connector as a trusted service application in Azure Active Directory. This is done by accepting a permission request to allow the connector to access your organization's data in Azure Active Directory. After you accept this request, the third-party data connector is added as an enterprise application to Azure Active Directory and represented as a service principal. For more information the consent process, see Tenant Admin Consent.

Here are the steps to access and accept the request to register the connector:

1. Go to this page and sign in using the credentials of a global administrator.

   The following dialog box is displayed. You can expand the carets to review the permissions that will be assigned to the connector.

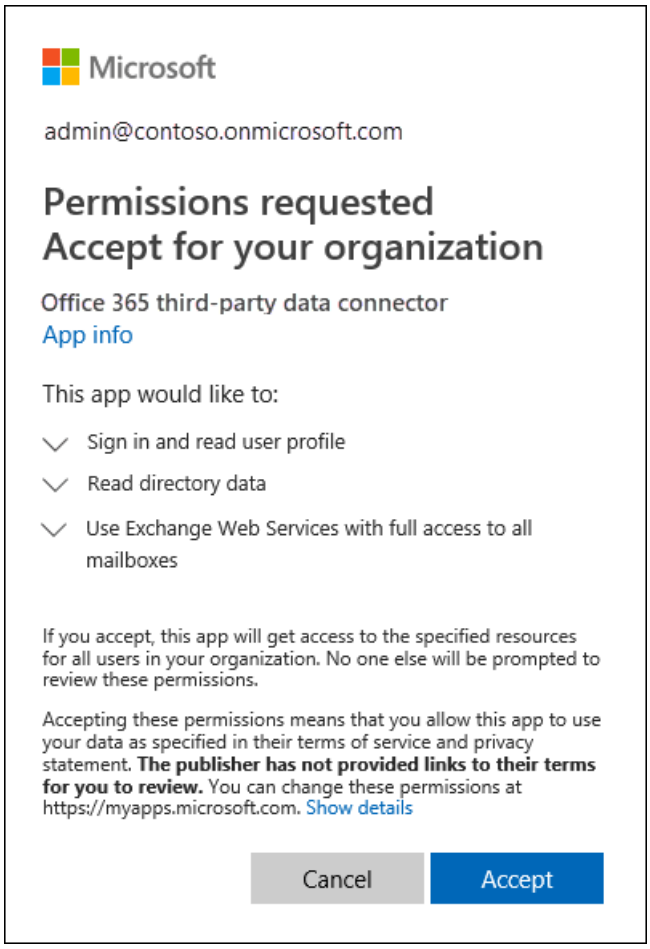

2. Click **Accept**.

After you accept the request, the Azure portal is displayed. To view the list of applications for your organization, click **Azure Active Directory** > **Enterprise applications**. The Microsoft 365 third-party data connector is listed on the **Enterprise applications** blade.

> **IMPORTANT**
> After September 30, 2018, third-party data will no longer be imported into mailboxes in your organization if you don't register a third-party data connector in Azure Active Directory. Note existing third-party data connectors (those created before September 30, 2018) must also be registered in Azure Active Directory by following the procedure in Step 5.

Revoking consent for a third-party data connector

After your organization consents to the permissions request to register a third-party data connector in Azure Active Directory, your organization can revoke that consent at any time. However, revoking the consent for a connector means that data from the third-party data source will no longer be imported into Microsoft 365.

To revoke consent for a third-party data connector, you can delete the application (by deleting the corresponding service principal) from Azure Active Directory using the **Enterprise applications** blade in the Azure portal, or by using the Remove-MsolServicePrincipal in Microsoft 365 PowerShell. You can also use the Remove-AzureADServicePrincipal cmdlet in Azure Active Directory PowerShell.

More information

- As previous explained, items from third-party data sources are imported to Exchange mailboxes as email messages. The partner connector imports the item using a schema required by the Microsoft 365 API. The following table describes the message properties of an item from a third-party data source after it's imported to an Exchange mailbox as an email message. The table also indicates if the message property is mandatory. Mandatory properties must be populated. If an item is missing a mandatory property, it

won't be imported to Microsoft 365. The import process returns an error message explaining why an item wasn't imported and which property is missing.

MESSAGE PROPERTY	MANDATORY?	DESCRIPTION	EXAMPLE VALUE
FROM	Yes	The user who originally created or sent the item in the third-party data source. The partner connector attempts to map the user ID from the source item (for example a Twitter handle) to a user account for all participants (users in the FROM and TO fields). A copy of the message will be imported to the mailbox of every participant. If none of the participants from the item can be mapped to a user account, the item will be imported to the third-party archiving mailbox in Microsoft 365. The participant who's identified as the sender of the item must have an active mailbox in the organization that the item is being imported to. If the sender doesn't have an active mailbox, the following error is returned: `One or more messages in the Request failed to be delivered to either From or Sender email address. You will need to resend your entire Request. Error: The request failed. The remote server returned an error: (401) Unauthorized.`	`bob@contoso.com`
TO	Yes	The user who received an item, if applicable for an item in the data source.	`bob@contoso.com`
SUBJECT	No	The subject from the source item.	`"Mega deals with Contoso coming your way! #ContosoHolidayDeals"`

MESSAGE PROPERTY	MANDATORY?	DESCRIPTION	EXAMPLE VALUE
DATE	Yes	The date the item was originally created or posted in the customer data source. For example, that date when a Twitter message was tweeted.	`01 NOV 2015`

BODY	No	The contents of the message or post. For some data sources, the contents of this property could be the same as the content for the **SUBJECT** property. During the import process, the partner connector attempts to maintain full fidelity from the content source as possible. If possible files, graphics, or other content from the body of the source item is included in this property. Otherwise, content from the source item is included in the **ATTACHMENT** property. The contents of this property depends on the partner connector and on the capability of the source platform.	`Author: bob@contoso.com`
			`Date: 10 DEC 2014`
			`Tweet: "Mega deals with Contoso coming your way! #ContosoHolidayDeals"`
			`Date: 01 NOV 2015`

ATTACHMENT	No	If an item in the data source (such as a tweet in Twitter or an instant messaging conversation) has an attached file or include images, the partner connect will first attempt to include attachments in the **BODY** property. If that isn't possible, then it's added to the ** ATTACHMENT ** property. Other examples of attachments include Likes in Facebook, metadata from the content source, and responses to a message or post.	`image.gif`

MESSAGE PROPERTY	MANDATORY?	DESCRIPTION	EXAMPLE VALUE
MESSAGECLASS	Yes	This is a multi-value property, which is created and populated by partner connector. The format of this property is `IPM.NOTE.Source.Event`. (This property must begin with `IPM.NOTE`. This format is similar to the one for the `IPM.NOTE.X` message class.) This property includes the following information: `Source`: Indicates the third-party data source; for example, Twitter, Facebook, or BlackBerry. `Event`: Indicates the type of activity that was performed in the third-party data source that produced the items; for example, a tweet in Twitter or a post in Facebook. Events are specific to the data source. One purpose of this property is to filter specific items based on the data source where an item originated or based on the type of event. For example, in an eDiscovery search you could create a search query to find all the tweets that were posted by a specific user.	`IPM.NOTE.Twitter.Tweet`

- When items are successfully imported to mailboxes in Microsoft 365, a unique identifier is returned back to the caller as part of the HTTP response. This identifier, called `x-IngestionCorrelationID`, can be used for subsequent troubleshooting purposes by partners for end-to-end tracking of items. It's recommended that partners capture this information and log it accordingly at their end. Here's an example of an HTTP response showing this identifier:

```
HTTP/1.1 200 OK
Content-Type: text/xml; charset=utf-8
Server: Microsoft-IIS/8.5
x-IngestionCorrelationID: 1ec7667d-f097-47fe-a9a2-bc7ab0a7552b
X-AspNet-Version: 4.0.30319
X-Powered-By: ASP.NET
Date: Tue, 02 Feb 2016 22:55:33 GMT
```

- You can use the Content Search tool in the security and compliance center to search for items that were imported to mailboxes from a third-party data source. To search specifically for these imported items,

you can use the following message property-value pairs in the keyword box for a Content Search.

- `kind:externaldata`: Use this property-value pair to search all third-party data types. For example, to search for items that were imported from a third-party data source and contained the word "contoso" in the Subject property of the imported item, you would use the keyword query `kind:externaldata AND subject:contoso`.
- `itemclass:ipm.externaldata.<third-party data type>`: Use this property-value pair to only search a specify type of third-party data. For example, to only search Facebook data that contains the word "contoso" in the Subject property, you would use the keyword query `itemclass:ipm.externaldata.Facebook* AND subject:contoso`.

For a complete list of values to use for third-party data types for the `itemclass` property, see Use Content Search to search third-party data that was imported to Microsoft 365.

For more information about using Content Search and creating keyword search queries, see:

- Content Search
- Keyword queries and search conditions for Content Search

Use Content Search to search third-party data imported by a custom partner connector

11/2/2020 • 3 minutes to read • Edit Online

You can use the Content Search eDiscovery tool in the Security & Compliance Center to search for items imported to mailboxes in Microsoft 365 from a third-party data source. You can create a query to search all imported third-party data items or you can create a query to search specific third-party data items. Also, you can also create a query-based retention policy or a query-based eDiscovery hold to preserve third-party data.

For more information about working with a partner to import third-party data and a list of the third-party data types that you can import to Microsoft 365, see Work with a partner to archive third-party data in Office 365.

> **IMPORTANT**
> The guidance in this article only applies to third-party data that was imported by a custom partner connector. This article doesn't apply to third-party data that is imported by using the third-party data connectors in the Microsoft compliance center.

Creating a query to search all third-party data

To search (or place on hold) any type of third-party data that you've imported to Office 365, you can use the `kind:externaldata` message property-value pair in the keyword box for a Content Search or when creating a query-based hold. For example, to search for items imported from any third-party data source and contain the word "contoso" in the Subject property of the imported item, you would use the following query:

```
kind:externaldata AND subject:contoso
```

The previous keyword query example includes the subject property. For a list of other properties for third-party data items that can include in a keyword query, see the "More information" section in Work with a partner to archive third-party data in Office 365.

When creating queries to search and hold third-party data, you can also use conditions to narrow the search results. For more information about creating Content Search queries, see Keyword queries and search conditions for Content Search.

Creating a query to search specific types of third-party data

Instead of searching all types of third-party data, you can create queries that only search for a specify type of third-party data by using the following message *property: value* pair in the keyword box for a Content Search:

```
itemclass:ipm.externaldata.<third-party data type>*
```

For example, to search Facebook data that contains the word "contoso" in the Subject property, you would use the following query:

```
itemclass:ipm.externaldata.Facebook* AND subject:contoso
```

The following table lists the third-party data types that you can search, and the value to use for the `itemclass:` message property to specifically search for that type of third-party data. The query syntax isn't case-sensitive.

THIRD-PARTY DATA TYPE	VALUE FOR `ITEMCLASS:` PROPERTY
AIM	`ipm.externaldata.AIM*`
American Idol	`ipm.externaldata.AmericanIdol*`
AOL with Pivot Client	`ipm.externaldata.Pivot.IM`
Apple Juice	`ipm.externaldata.AppleJuice*`
Ares	`ipm.externaldata.Ares*`
Axs Encrypted	`ipm.externaldata.AxsEncrypted*`
Axs Exchange	`ipm.externaldata.AxsExchange*`
Axs Local Archive	`ipm.externaldata.AxsLocalArchive*`
Axs Placeholder	`ipm.externaldata.AxsPlaceHolder*`
Axs Signed	`ipm.externaldata.AxsSigned*`
Bazaarvoice	`ipm.externaldata.Bazaarvoice*`
Bearshare	`ipm.externaldata.Bearshare*`
BitTorrent	`ipm.externaldata.BitTorrent*`
Blackberry	`ipm.externaldata.Blackberry*`
BlackBerry Call Logs	`ipm.externaldata.BlackBerryCall*`
BlackBerry Messenger	`ipm.externaldata.BlackBerryMessenger*`
BlackBerry PIN	`ipm.externaldata.BlackBerryPIN*`
BlackBerry SMS	`ipm.externaldata.BlackBerrySMS*`
Bloomberg	`ipm.externaldata.Bloomberg*`
Bloomberg Message	`ipm.externaldata.conversation.Bloomberg Message*`
Bloomberg Messaging	`ipm.externaldata.BloombergMessaging*`
Box	`ipm.externaldata.Box*`
Cisco IM & Presence Server	`ipm.externaldata.Jabber.IM`

THIRD-PARTY DATA TYPE	VALUE FOR `ITEMCLASS:` PROPERTY
Cisco Jabber	`ipm.externaldata.Jabber*`
CipherCloud for Salesforce Chatter	`ipm.externaldata.Chatter.Post` `ipm.externaldata.Chatter.Comment`
Direct Connect	`ipm.externaldata.DirectConnect*`
Facebook	`ipm.externaldata.Facebook*`
FastTrack	`ipm.externaldata.FastTrack*`
FXConnect	`ipm.externaldata.FXConnect.chat`
Flickr	`ipm.externaldata.Flickr*`
Gnutella	`ipm.externaldata.Gnutella*`
Google+	`ipm.externaldata.GooglePlus*`
Google Talk	`ipm.externaldata.GoogleTalk*`
GoToMyPC	`ipm.externaldata.GoToMyPC*`
HipChat	`ipm.externaldata.HipChat*`
Hopster	`ipm.externaldata.Hopster*`
HubConnex	`ipm.externaldata.HubConnex*`
IBM Connections	`ipm.externaldata.Connections*`
IBM SameTime	`ipm.externaldata.Sametime*`
ICE Chat	`ipm.externaldata.conversation.Ice Chat*`
Indii Messenger	`ipm.externaldata.Indii*`
Instagram	`ipm.externaldata.Instagram*`
Instant Bloomberg	`ipm.externaldata.InstantBloomberg*`
InvestEdge	`ipm.externaldata.InvestEdge*`
IRC	`ipm.externaldata.IRC*`
Jive	`ipm.externaldata.Jive*`
JiveApiRetention	`ipm.externaldata.JiveApiRetention*`

THIRD-PARTY DATA TYPE	VALUE FOR `ITEMCLASS:` PROPERTY
JXTA	`ipm.externaldata.JXTA*`
LinkedIn	`ipm.externaldata.LinkedIn*`
MFTP	`ipm.externaldata.MFTP*`
Microsoft UC	`ipm.externaldata.MicrosoftUC*`
Mind Align	`ipm.externaldata.MindAlign*`
Mobile Guard	`ipm.externaldata.MobileGuard*`
MSN	`ipm.externaldata.MSN*`
MySpace	`ipm.externaldata.MySpace*`
NEONetwork	`ipm.externaldata.NEONetwork*`
OpenNap	`ipm.externaldata.OpenNap*`
Pinterest	`ipm.externaldata.Pinterest*`
Pivot	`ipm.externaldata.Pivot*`
QQ	`ipm.externaldata.QQ*`
Microsoft SharePoint	`ipm.externaldata.SharePoint*`
Salesforce Chatter	`ipm.externaldata.Chatter*`
Skype for Business	`ipm.externaldata.Skype*`
Slack Enterprise Grid	`ipm.externaldata.Slack.IM`
SoftEther	`ipm.externaldata.SoftEther*`
Squawker	`ipm.externaldata.Squawker*`
Symphony	`ipm.externaldata.Symphony*`
Thomson Reuters	`ipm.externaldata.Reuters*`
Thomson Reuters Eikon Messenger	`ipm.externaldata.ReutersEikon*`
Tor	`ipm.externaldata.Tor*`
TTT	`ipm.externaldata.TTT*`

THIRD-PARTY DATA TYPE	VALUE FOR `ITEMCLASS:` PROPERTY
Twitter	`ipm.externaldata.Twitter*`
UBS Chat	`ipm.externaldata.UBS*`
Vimeo	`ipm.externaldata.Vimeo*`
WinMX	`ipm.externaldata.WinMX*`
Winny	`ipm.externaldata.Winny*`
Yahoo!	`ipm.externaldata.Yahoo!*`
Yammer	`ipm.externaldata.Yammer*`
YellowJacket	`ipm.externaldata.YellowJacket*`
YouTube	`ipm.externaldata.YouTube*`

Enable archive mailboxes in the compliance center

2/18/2021 • 6 minutes to read • Edit Online

Archiving in Microsoft 365 (also called *In-Place Archiving*) provides users with additional mailbox storage space. After you turn on archive mailboxes, users can access and store messages in their archive mailboxes by using Microsoft Outlook and Outlook on the web (formerly known as Outlook Web App). Users can also move or copy messages between their primary mailbox and their archive mailbox. They can also recover deleted items from the Recoverable Items folder in their archive mailbox by using the Recover Deleted Items tool.

> **NOTE**
> The auto-expanding archiving feature in Microsoft 365 provides additional storage in archive mailboxes. When auto-expanding archiving is turned on, and then the initial storage quota in a user's archive mailbox is reached, Microsoft 365 automatically adds additional storage space. This means that users won't run out of mailbox storage space and you won't have to manage anything after you initially enable the archive mailbox and turn on auto-expanding archiving for your organization. For more information, see Overview of unlimited archiving.

Get the necessary permissions

You have to be assigned the Mail Recipients role in Exchange Online to enable or disable archive mailboxes. By default, this role is assigned to the Recipient Management and Organization Management role groups on the **Permissions** page in the Exchange admin center. If you don't see the **Archive** page in the Security & Compliance Center, ask your administrator to assign you the necessary permissions.

Enable an archive mailbox

1. Go to https://protection.office.com.
2. Sign in using your work or school account.
3. In the left pane of the Security & Compliance Center, click **Information governance** > **Archive**.

   The **Archive** page is displayed. The **Archive mailbox** column indicates whether an archive mailbox is enabled or disabled for each user.

   > **NOTE**
   > The **Archive** page shows a maximum of 500 users.

4. In the list of mailboxes, select the user that you want to enable the archive mailbox for.

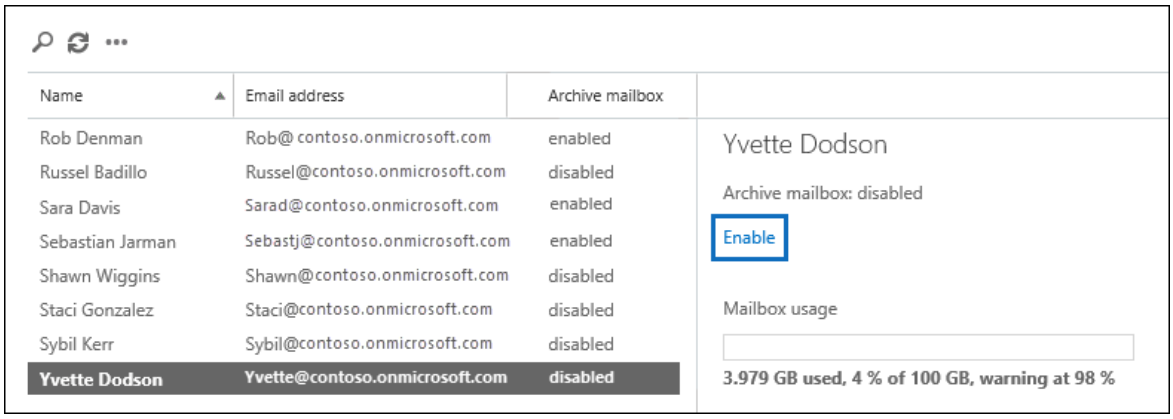

5. In the details pane for the selected user, click **Enable**.

   A warning is displayed saying that if you enable the archive mailbox, items in the user's mailbox that are older than the archiving policy assigned to the mailbox will be moved to the new archive mailbox. The default archive policy that is part of the retention policy assigned to Exchange Online mailboxes moves items to the archive mailbox two years after the date the item was delivered to the mailbox or created by the user. For more information, see the **More info** section in this article.

6. Click **Yes** to enable the archive mailbox.

   It might take a few moments to create the archive mailbox. When it's created, **Archive mailbox: enabled** is displayed in the details pane for the selected user. You might have to click **Refresh** to update the information in the details pane.

> **TIP**
> You can also bulk-enable archive mailboxes by selecting multiple users with disabled archive mailboxes (use the Shift or Ctrl keys). After selecting multiple mailboxes, click **Enable** in the details pane.

Disable an archive mailbox

You can also use the **Archive** page in the Security & Compliance Center to disable a user's archive mailbox. After you disable an archive mailbox, you can reconnect it to the user's primary mailbox within 30 days of disabling it. In this case, the original contents of the archive mailbox are restored. After 30 days, the contents of the original archive mailbox are permanently deleted and can't be recovered. So if you re-enable the archive more than 30 days after disabling it, a new archive mailbox is created.

The default archive policy assigned to users' mailboxes moves items to the archive mailbox two years after the date the item is delivered. If you disable a user's archive mailbox, no action will be taken on mailbox items and they will remain in the user's primary mailbox.

To disable an archive mailbox:

1. Go to https://protection.office.com.
2. Sign in using your work or school account.
3. In the left pane of the Security & Compliance Center, click **Information governance** > **Archive**.

   The **Archive** page is displayed. The **Archive mailbox** column indicates whether an archive mailbox is enabled or disabled for each user.

> **NOTE**
>
> The **Archive** page shows a maximum of 500 users.

4. In the list of mailboxes, select the user that you want to disable the archive mailbox for.
5. In the details pane, click **Disable**.

   A warning message is displayed saying that you'll have 30 days to re-enable the archive mailbox, and that after 30 days, all information in the archive will be permanently deleted.
6. Click **Yes** to disable the archive mailbox.

   It might take a few moments to disable the archive mailbox. When it's disabled, **Archive mailbox: disabled** is displayed in the details pane for the selected user. You might have to click **Refresh** to update the information in the details pane.

> **TIP**
>
> You can also bulk-disable archive mailboxes by selecting multiple users with enabled archive mailboxes (use the Shift or Ctrl keys). After selecting multiple mailboxes, click **Disable** in the details pane.

Use Exchange Online PowerShell to enable or disable archive mailboxes

You can also use Exchange Online PowerShell to enable archive mailboxes. The primary reason to use PowerShell is that you can quickly enable the archive mailbox for all users in your organization.

The first step is to connect to Exchange Online PowerShell. For instructions, see Connect to Exchange Online PowerShell.

After you're connected to Exchange Online, you can run the commands in the following sections to enable or disable archive mailboxes.

Enable archive mailboxes

Run the following command to enable the archive mailbox for a single user.

```
Enable-Mailbox -Identity <username> -Archive
```

Run the following command to enable the archive mailbox for all users in your organization (whose archive mailbox is currently not enabled).

```
Get-Mailbox -Filter {ArchiveGuid -Eq "00000000-0000-0000-0000-000000000000" -AND RecipientTypeDetails -Eq
"UserMailbox"} | Enable-Mailbox -Archive
```

Disable archive mailboxes

Run the following command to disable the archive mailbox for a single user.

```
Disable-Mailbox -Identity <username> -Archive
```

Run the following command to disable the archive mailbox for all users in your organization (whose archive mailbox is currently enabled).

```
Get-Mailbox -Filter {ArchiveGuid -Ne "00000000-0000-0000-0000-000000000000" -AND RecipientTypeDetails -Eq
"UserMailbox"} | Disable-Mailbox -Archive
```

More information

- When an archive mailbox is enabled, users can store messages in their archive mailbox. Users can access their archive mailboxes by using Microsoft Outlook and Outlook on the web. Using either of these client applications, users can view messages in their archive mailbox and move or copy messages between their primary mailbox and their archive mailbox. Users can also recover deleted items from the Recoverable Items folder in their archive mailbox by using the Recover Deleted Items tool.

  For a list of Outlook licenses that support In-Place Archiving, see Outlook license requirements for Exchange features.

- Archive mailboxes help you and your users to meet your organization's retention, eDiscovery, and hold requirements. For example, you can use your organization's Exchange retention policy to move mailbox content to users' archive mailbox. When you use the Content Search tool in the Security & Compliance Center to search a user's mailbox for specific content, the user's archive mailbox will also be searched. And, when you place a Litigation Hold or apply a retention policy to a user's mailbox, items in the archive mailbox are also retained.

- After archive mailboxes are enabled, your organization can take advantage of the default Exchange retention policy (also called Messaging Records Management or MRM policy) that is automatically assigned to every mailbox. When an archive mailbox is enabled, the default Exchange retention policy automatically does the following:

  - Moves items that are two years or older from a user's primary mailbox to their archive mailbox.
  - Moves items that are 14 days or older from the Recoverable Items folder in the user's primary mailbox to the Recoverable Items folder in their archive mailbox.

- For more information about archive mailboxes and Exchange retention policies, see:

  - Retention tags and retention policies in Exchange Online
  - Default Retention Policy in Exchange Online
  - Set up an archive and deletion policy for mailboxes in your organization

Overview of unlimited archiving

2/18/2021 • 5 minutes to read • Edit Online

In Office 365, archive mailboxes provide users with additional mailbox storage space. After a user's archive mailbox is enabled, up to 100 GB of additional storage is available. In the past, when the 100-GB storage quota was reached, organizations had to contact Microsoft to request additional storage space for an archive mailbox. That's no longer the case.

The unlimited archiving feature in Microsoft 365 (called *auto-expanding archiving*) provides additional storage in archive mailboxes. When the storage quota in the archive mailbox is reached, Microsoft 365 automatically increases the size of the archive, which means that users won't run out of mailbox storage space and administrators won't have to request additional storage for archive mailboxes.

For step-by-step instructions for turning on auto-expanding archiving, see Enable unlimited archiving.

> **NOTE**
> Auto-expanding archiving also supports shared mailboxes. To enable the archive for a shared mailbox, an Exchange Online Plan 2 license or an Exchange Online Plan 1 license with an Exchange Online Archiving license is required.

How auto-expanding archiving works

As previously explained, additional mailbox storage space is created when a user's archive mailbox is enabled. When auto-expanding archiving is enabled, Microsoft 365 periodically checks the size of the archive mailbox. When an archive mailbox gets close to its storage limit, Microsoft 365 automatically creates additional storage space for the archive. If the user runs out of this additional storage space, Microsoft 365 adds more storage space to the user's archive. This process happens automatically, which means administrators don't have to request additional archive storage or manage auto-expanding archiving.

Here's a quick overview of the process.

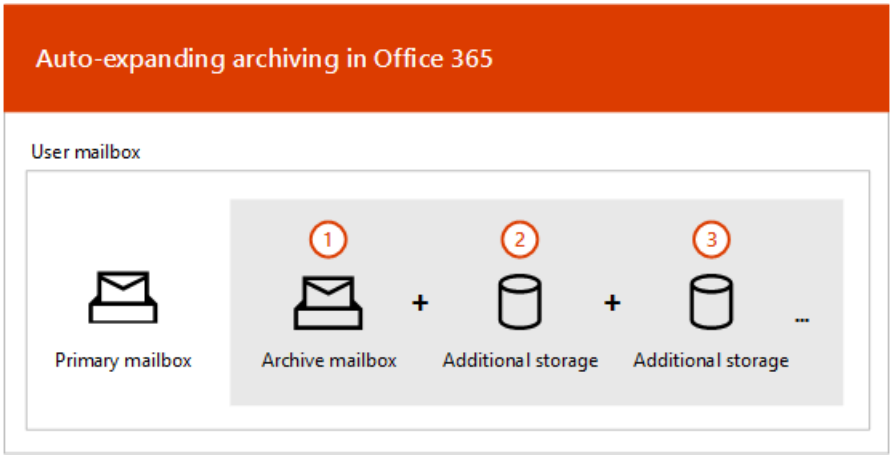

1. Archiving is enabled for a user mailbox or a shared mailbox. An archive mailbox with 100 GB of storage space is created, and the warning quota for the archive mailbox is set to 90 GB.
2. An administrator enables auto-expanding archiving for the mailbox. When the archive mailbox (including the Recoverable Items folder) reaches 90 GB, it's converted to an auto-expanding archive, and Microsoft 365 adds storage space to the archive. It can take up to 30 days for the additional storage space to be provisioned.

> **NOTE**
> If a mailbox is placed on hold or assigned to a retention policy, the storage quota for the archive mailbox is increased to 110 GB when auto-expanding archiving is enabled. Similarly, the archive warning quota is increased to 100 GB.

3. Microsoft 365 automatically adds more storage space when necessary.

> **IMPORTANT**
> Auto-expanding archive is only supported for mailboxes used for individual users (or shared mailboxes) with a growth rate that doesn't exceed 1 GB per day. A user's archive mailbox is intended for just that user. Using journaling, transport rules, or auto-forwarding rules to copy messages to an archive mailbox is not permitted. Microsoft reserves the right to deny unlimited archiving in instances where a user's archive mailbox is used to store archive data for other users or in other cases of the inappropriate use.

What gets moved to the additional archive storage space?

To make efficient use of auto-expanding archive storage, folders may get moved. Microsoft 365 determines which folders get moved when additional storage is added to the archive. Sometimes when a folder is moved, one or more subfolders are automatically created and items from the original folder are distributed to these folders to facilitate the moving process. When viewing the archive portion of the folder list in Outlook, these subfolders are displayed under the original folder. The naming convention that Microsoft 365 uses to name these subfolders is **<folder name>_yyyy (Created on mmm dd, yyyy h_mm)**, where:

- **yyyy** is the year the messages in the folder were received.
- **mmm dd, yyyy h_m** is the date and time that the subfolder was created by Office 365, in UTC format, based on the user's time zone and regional settings in Outlook.

The following screenshots show a folder list before and after messages are moved to an auto-expanded archive.

Before additional storage is added

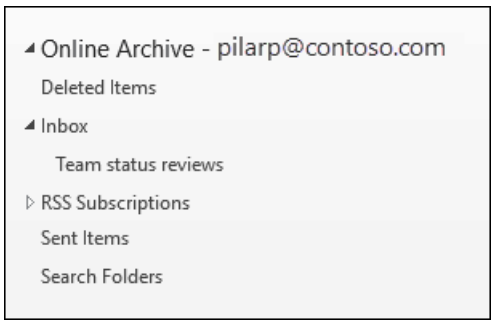

After additional storage is added

NOTE

As previously described, Microsoft 365 moves items to subfolders (and names them using the naming convention described above) to help distribute content to an auxiliary archive. But moving items to subfolders may not always be the case. Sometimes an entire folder may be moved to an auxiliary archive. In this case, the folder will retain its original name. It won't be apparent in the folder list in Outlook that the folder was moved to an auxiliary archive.

Outlook requirements for accessing items in an auto-expanded archive

To access messages that are stored in an auto-expanded archive, users have to use one of the following Outlook clients:

- Outlook 2016 or Outlook 2019 for Windows
- Outlook on the web
- Outlook 2016 or Outlook 2019 for Mac

Here are some things to consider when using Outlook or Outlook on the web to access messages stored in an auto-expanded archive.

- You can access any folder in your archive mailbox, including ones that were moved to the auto-expanded storage area.
- Search for auto-expanded archiving is available in Outlook for the web. Similar to Online Archive, you can search for items that were moved to an additional storage area only by searching the current folder itself. This means that you must select the archive folder in the folder list, and then select a single folder as your search scope. Similarly, if a folder in an auto-expanded storage area contains subfolders, you must search each subfolder separately.
- Auto-expanded archive search is available in Outlook Desktop in Current Channel (Preview). Within this preview, the Current Mailbox scope is available, thus allowing you to search the auto-expanded archive. For more information about this and other Microsoft Search support features, see How Outlook for Windows connected to Exchange Online utilizes Microsoft Search.
- Item counts in Outlook and Read/Unread counts (in Outlook and Outlook on the web) in an auto-expanded archive might not be accurate.
- You can delete items in a subfolder that points to an auto-expanded storage area, but the folder itself can't be deleted.
- You can't use the Recover Deleted Items feature to recover an item that was deleted from an auto-expanded storage area.

Auto-expanding archiving and other compliance features

This section explains the functionality between auto-expanding archiving and other compliance and data governance features.

- **eDiscovery:** When you use an eDiscovery tool, such as Content Search or In-Place eDiscovery, the additional storage areas in an auto-expanded archive are also searched.
- **Retention:** When you put a mailbox on hold by using tools such as Litigation Hold in Exchange Online or eDiscovery case holds and retention policies in the security and compliance center, content located in an auto-expanded archive is also placed on hold.

- **Messaging records management (MRM):** If you use MRM deletion policies in Exchange Online to permanently delete expired mailbox items, expired items located in the auto-expanded archive will also be deleted.
- **Import service:** You can use the Office 365 Import service to import PST files to a user's auto-expanded archive. You can import up to 100 GB of data from PST files to the user's archive mailbox.

More information

For more technical details about auto-expanding archiving, see Microsoft 365: Auto-Expanding Archives FAQ.

Enable unlimited archiving - Admin Help

2/18/2021 • 7 minutes to read • Edit Online

You can use the Exchange Online auto-expanding archiving feature to enable unlimited storage space for archive mailboxes. When auto-expanding archiving is turned on, additional storage space is automatically added to a user's archive mailbox when it approaches the storage limit. The result is unlimited mailbox storage capacity. You can turn on auto-expanding archiving for everyone in your organization or just for specific users. For more information about auto-expanding archiving, see Overview of unlimited archiving in Office 365.

Before you enable auto-expanding archiving

- You have to be a global administrator in your organization or a member of the Organization Management role group in your Exchange Online organization to enable auto-expanding archiving for your entire organization or for specific users. Alternately, you have to be a member of a role group that's assigned the Mail Recipients role to enable auto-expanding archiving for specific users.
- A user's archive mailbox has to be enabled before you can enable auto-expanding archiving. A user must be assigned an Exchange Online Plan 2 license to enable the archive mailbox. If a user is assigned an Exchange Online Plan 1 license, you would have to assign them a separate Exchange Online Archiving license to enable their archive mailbox. See Enable archive mailboxes in the Security & Compliance Center.
- You can also use PowerShell to enable archive mailboxes. See the More information section for an example of the PowerShell command that you can use to enable archive mailboxes for all users in your organization.
- Auto-expanding archiving also supports shared mailboxes. To enable the archive for a shared mailbox, an Exchange Online Plan 2 license or an Exchange Online Plan 1 license with an Exchange Online Archiving license is required.
- Auto-expanding archiving prevents you from recovering or restoring an inactive mailbox. That means if you enable auto-expanding archiving for a mailbox and the mailbox is made inactive at a later date, you won't be able to recover the inactive mailbox (by converting it to an active mailbox) or restore it (by merging the contents to an existing mailbox). If auto-expanding archiving is enabled on an inactive mailbox, the only way to recover data is by using the Content search tool in the Microsoft 365 compliance center to export the data from the mailbox and import to another mailbox. For more information, see the "Inactive mailboxes and auto-expanding archives" section in Overview of inactive mailboxes.
- You can't use the Exchange admin center or the Security & Compliance Center to enable auto-expanding archiving. You have to use Exchange Online PowerShell. To connect to your Exchange Online organization using remote PowerShell, see Connect to Exchange Online PowerShell.

Enable auto-expanding archiving for your entire organization

You can enable auto-expanding archiving for your entire organization. After you turn it on, auto-expanding archiving will be enabled for existing user mailboxes and for new user mailboxes that are created. When you create user mailboxes, be sure to enable the user's main archive mailbox so the auto-expanding archiving feature works for the new user mailbox.

1. Connect to Exchange Online PowerShell
2. Run the following command in Exchange Online PowerShell to enable auto-expanding archiving for your

entire organization.

```
Set-OrganizationConfig -AutoExpandingArchive
```

Enable auto-expanding archiving for specific users

Instead of enabling auto-expanding archiving for every user in your organization, you can enable it only for specific users. You might do this because only some users might have a need for a large archive storage capacity.

When you enable auto-expanding archiving for a specific user and the user's mailbox in on hold or assigned to a retention policy, the following two configurations changes are made:

- The storage quota for the user's primary archive mailbox is increased by 10 GB (from 100 GB to 110 GB). The archive warning quota is also increased by 10 GB (from 90 GB to 100 GB).
- The storage quota for the Recoverable Items folder in the user's primary mailbox is increased by 10 GB (also from 100 GB to 110 GB). The Recoverable Items warning quota is also increased by 10 GB (from 90 GB to 100 GB). These changes are applicable only if the mailbox in on hold or assigned to a retention policy.

This additional space is added to prevent any storage issues that may occur before the auto-expanding archive is provisioned. Additional storage space *is not* added when you enable auto-expanding archiving for your entire organization, as described in the previous section.

1. Connect to Exchange Online PowerShell
2. Run the following command in Exchange Online PowerShell to enable auto-expanding archiving for a specific user. As previously explained, the user's archive mailbox (main archive) must be enabled before you can turn on auto-expanding archiving for that user.

   ```
   Enable-Mailbox <user mailbox> -AutoExpandingArchive
   ```

> **IMPORTANT**
> In an Exchange hybrid deployment, you can't use the **Enable-Mailbox -AutoExpandingArchive** command to enable auto-expanding archiving for a specific user whose primary mailbox is on-premises and whose archive mailbox is cloud-based. To enable auto-expanding archiving for cloud-based archive mailboxes in an Exchange hybrid deployment, you have to run the **Set-OrganizationConfig -AutoExpandingArchive** command in Exchange Online PowerShell to enable auto-expanding archiving for the entire organization. If a user's primary and archive mailboxes are both cloud-based, then you can use the **Enable-Mailbox -AutoExpandingArchive** command to enable auto-expanding archiving for that specific user.

Verify that auto-expanding archiving is enabled

To verify that auto-expanding archiving is enabled for your organization, run the following command in Exchange Online PowerShell.

```
Get-OrganizationConfig | FL AutoExpandingArchiveEnabled
```

A value of `True` indicates that auto-expanding archiving is enabled for the organization.

To verify that auto-expanding archiving is enabled for a specific user, run the following command in Exchange

Online PowerShell.

```
Get-Mailbox <user mailbox> | FL AutoExpandingArchiveEnabled
```

A value of `True` indicates that auto-expanding archiving is enabled for the user.

To determine if auto-expanding archiving is enabled for inactive mailboxes, run the following command in Exchange Online PowerShell.

```
Get-Mailbox -InactiveMailboxOnly | FL UserPrincipalName,AutoExpandingArchiveEnabled
```

A value of `True` indicates that auto-expanding archiving is enabled for the inactive mailbox. A value of `False` indicates that auto-expanding archiving isn't enabled.

Keep the following things in mind after you enable auto-expanding archiving:

- If you run the **Set-OrganizationConfig -AutoExpandingArchive** command to enable auto-expanding archiving for your organization, you don't have to run the **Enable-Mailbox -AutoExpandingArchive** on individual mailboxes. Running the **Set-OrganizationConfig** cmdlet to enable auto-expanding archiving for your organization doesn't change the *AutoExpandingArchiveEnabled* property on user mailboxes to `True`.
- Similarly, the values for the *ArchiveQuota* and *ArchiveWarningQuota* mailbox properties aren't changed when you enable auto-expanding archiving. In fact, when you enable auto-expanding archiving for a user mailbox and the *AutoExpandingArchiveEnabled* property is set to `True`, the *ArchiveQuota* and *ArchiveWarningQuota* properties are ignored. Here's an example of these mailbox properties after auto-expanding archiving is enabled for a user's mailbox.

  ```
  PS C:\> Get-Mailbox sarad | FL Archive*Quota,AutoExpandingArchiveEnabled

  ArchiveQuota                : 100 GB (107,374,182,400 bytes)
  ArchiveWarningQuota         : 90 GB (96,636,764,160 bytes)
  AutoExpandingArchiveEnabled : True
  ```

More information

- You can also use PowerShell to enable archive mailboxes. For example, you can run the following command in Exchange Online PowerShell to enable archive mailboxes for all users whose archive mailbox isn't already enabled.

  ```
  Get-Mailbox -Filter {ArchiveStatus -Eq "None" -AND RecipientTypeDetails -eq "UserMailbox"} | Enable-Mailbox -Archive
  ```

- After you turn on auto-expanding archiving for your organization or for a specific user, an archive mailbox is converted to an auto-expanding archive when the archive mailbox (including the Recoverable Items folder) reaches 90 GB. It can take up to 30 days for the additional storage space to be provisioned.
- After you turn on auto-expanding archiving, it can't be turned off. Additionally, administrators can't adjust the storage quota for auto-expanding archiving.
- Auto-expanding archiving is supported for cloud-based archive mailboxes in an Exchange hybrid deployment for users who have an on-premises primary mailbox. However, after auto-expanding archiving is enabled for a cloud-based archive mailbox, you can't off-board that archive mailbox back to the on-premises Exchange organization. Auto-expanding archiving isn't supported for on-premises mailboxes in any version of Exchange Server.

- For a list of Outlook clients that users can use to access items in the additional storage area in their archive mailbox, see the "Outlook requirements for accessing items in an auto-expanded archive" section in Overview of unlimited archiving.
- As previously explained, 10 GB is added to the storage quota of the user's primary archive mailbox (and to the Recoverable Items folder if the mailbox is on hold) when you run the **Enable-Mailbox -AutoExpandingArchive** command. This provides additional storage until the auto-expanded storage space is provisioned (which can take up to 30 days). This additional storage space isn't added when you run the **Set-OrganizationConfig -AutoExpandingArchive** to enable auto-expanding archiving for all mailboxes in your organization. If you enabled auto-expanding archiving for the entire organization, but need to add the additional 10 GB of storage space for a specific user, you can run the **Enable-Mailbox -AutoExpandingArchive** command on that mailbox. You will receive an error saying that auto-expanding archiving has already been enabled, but the additional storage space will be added to the mailbox.

> **IMPORTANT**
>
> Auto-expanding archiving is only supported for mailboxes used for individual users or shared mailboxes with a growth rate that doesn't exceed 1 GB per day. Using journaling, transport rules, or auto-forwarding rules to copy messages to an archive mailbox for the purposes of archiving is not permitted. A user's archive mailbox is intended for just that user. Microsoft reserves the right to deny unlimited archiving in instances where a user's archive mailbox is used to store archive data for other users or in other cases of inappropriate use.

Set up an archive and deletion policy for mailboxes in your organization

2/18/2021 • 16 minutes to read • Edit Online

In Microsoft 365, admins can create an archiving and deletion policy that automatically moves items to a user's archive mailbox and automatically deletes items from the mailbox. The admin does this by creating a retention policy that's assigned to mailboxes, and moves items to a user's archive mailbox after a certain period of time and that also deletes items from the mailbox after they reach a certain age limit. The actual rules that determine what items are moved or deleted and when that happens are called retention tags. Retention tags are linked to a retention policy, that in turn is assigned to a user's mailbox. A retention tag applies retention settings to individual messages and folders in a user's mailbox. It defines how long a message remains in the mailbox and what action is taken when the message reaches the specified retention age. When a message reaches its retention age, it's either moved to the user's archive mailbox or it's deleted.

The steps in this article will set up an archiving and retention policy for a fictitious organization named Alpine House. Setting up this policy includes the following tasks:

- Enabling an archive mailbox for every user in the organization. This gives users addition mailbox storage, and is required so that a retention policy can move items to the archive mailbox. It also lets a user store archival information by moving items to their archive mailbox.
- Creating three custom retention tags that do the following:
  - Automatically moves items that are 3 years old to the user's archive mailbox. Moving items to the archive mailbox frees up space in a user's primary mailbox.
  - Automatically deletes items that are 5 years old from the Deleted Items folder. This also frees up space in the user's primary mailbox. User's will have the opportunity to recover these items if necessary. See the footnote in the More information section for more details.
  - Automatically (and permanently) deletes items that are 7 years old from both the primary and archive mailbox. Because of compliance regulations, some organization's are required to retain email for a certain period of time. After this time period expires, an organization might want to permanently remove these items user mailboxes.
- Creating a new retention policy and adding the new custom retention tags to it. Additionally, you'll also add built-in retention tags to the new retention policy. This includes personal tags that users can assign to items in their mailbox. You'll also add a retention tag that moves items from the Recoverable Items folder in the user's primary mailbox to the Recoverable Items folder in their archive mailbox. This helps free up space in a user's Recoverable Items folder when their mailbox is placed on hold.

You can follow some or all of the steps in this article to set up an archive and deletion policy for mailboxes in your own organization. We recommend that you test this process on a few mailboxes before implementing it on all mailboxes in your organization.

Before you set up an archive and deletion policy

- You have to be a global administrator in your organization to perform the steps in this topic.
- When you create a new user account and assign the user an Exchange Online license, a mailbox is automatically created for the user. When the mailbox is created, it's automatically assigned a default retention policy, named Default MRM Policy. In this article, you will create a new retention policy and then

assign it to user mailboxes, replacing the Default MRM policy. A mailbox can have only one retention policy assigned to it at any one time.

- To learn more about retention tags and retention policies in Exchange Online, see Retention tags and retention policies.

Step 1: Enable archive mailboxes for users

The first step is to enable the archive mailbox for each user in your organization. A user's archive mailbox has to be enabled so that a retention tag with a "Move to Archive" retention action can move the item after the retention age expires.

> **NOTE**
> You can enable archive mailboxes any time during this process, just as long as they're enabled at some point before you complete the process. If an archive mailbox isn't enabled, no action is taken on any items that have an archive or deletion policy assigned to it.

1. Go to https://protection.office.com.
2. Sign in using your global administrator account.
3. In the Security & Compliance Center, go to **Information governance** > **Archive**.

   A list of the mailboxes in your organization is displayed and whether the corresponding archive mailbox is enabled or disabled.
4. Select all the mailboxes by clicking on the first one in the list, holding down the **Shift** key, and then clicking the last one in the list.

   > **TIP**
   > This step assumes that no archive mailboxes are enabled. If you have any mailboxes with the archive enabled, hold down the **Ctrl** key and click each mailbox that has a disabled archive mailbox. Or you can click the **Archive mailbox** column header to sort the rows based on whether the archive mailbox is enabled or disabled to make it easier to select mailboxes.
5. In the details pane, under **Bulk Edit**, click **Enable**.

   A warning is displayed saying that items that are older than two years will be moved to the new archive mailbox. This is because the default retention policy that's assigned a new user mailbox when it's created has an archive default policy tag that has a retention age of 2 years. The custom archive default policy tag that you'll create in Step 2 has a retention age of 3 years. That means items that are 3 years or older will be moved to the archive mailbox.
6. Click **Yes** to close the warning message and start the process to enable the archive mailbox for each selected mailbox.
7. When the process is complete, click **Refresh** to update the list on the **Archive** page.

   The archive mailbox is enabled for all user's in your organization.

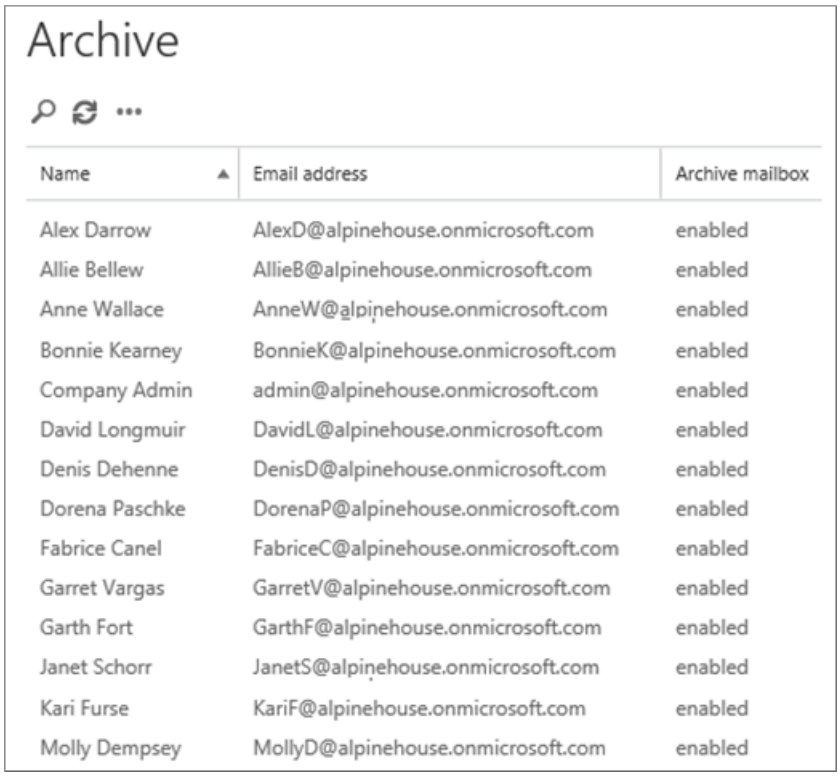

Step 2: Create new retention tags for the archive and deletion policies

In this step, you'll create the three custom retention tags that were previously described.

- Alpine House 3 Year Move to Archive (custom archive policy)
- Alpine House 7 Year Permanently Delete (custom deletion policy)
- Alpine House Deleted Items 5 Years Delete and Allow Recovery (custom tag for the Deleted Items folder)

To create new retention tags, you'll use the Exchange admin center (EAC) in your Exchange Online organization. Be sure to use the classic version of the EAC.

1. Go to https://admin.protection.outlook.com/ecp/ and sign in using your credentials.
2. In the EAC, go to **Compliance management** > **Retention tags**

   A list of the retention tags for your organization is displayed.

Create a custom archive default policy tag

First, you'll create a custom archive default policy tag (DPT) that will move items to the archive mailbox after 3 years.

1. On the **Retention tags** page, click **New tag+**, and then select **applied automatically to entire mailbox (default)**.
2. On the **New tag applied automatically to entire mailbox (default)** page, complete the following fields:

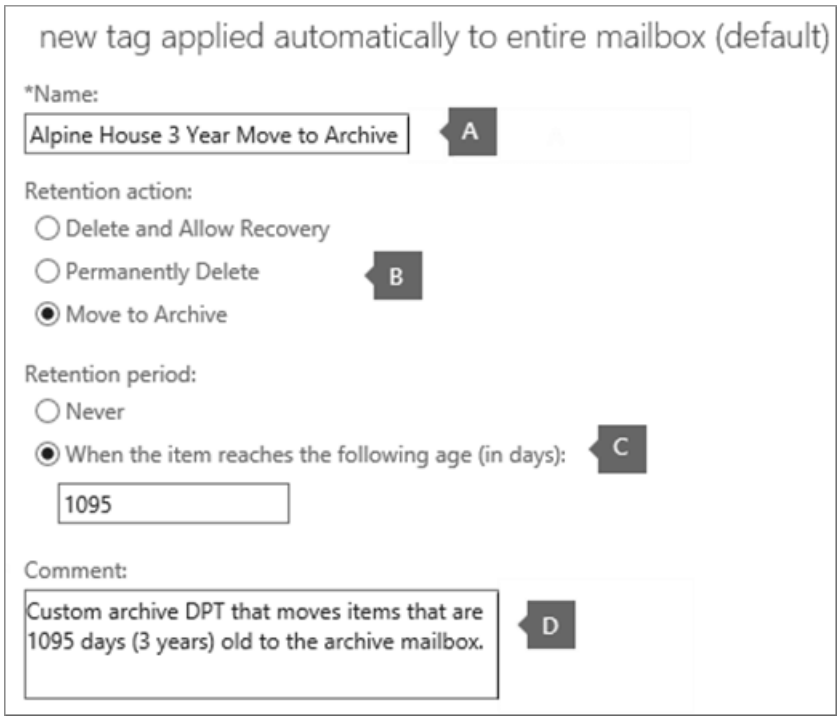

a. **Name** Type a name for the new retention tag.

b. **Retention action** Select **Move to Archive** to move items to the archive mailbox when the retention period expires.

c. **Retention period** Select **When the item reaches the following age (in days)**, and then enter the duration of the retention period. For this scenario, items will be moved to the archive mailbox after 1095 days (3 years).

d. **Comment** (Optional) Type a comment that explains the purpose of the custom retention tag.

3. Click **Save** to create the custom archive DPT.

   The new archive DPT is displayed in the list of retention tags.

Create a custom deletion default policy tag

Next, you'll create another custom DPT but this one will be a deletion policy that permanently deletes items after 7 years.

1. On the **Retention tags** page, click **New tag+**, and then select **applied automatically to entire mailbox (default)**.

2. On the **New tag applied automatically to entire mailbox (default)** page, complete the following fields:

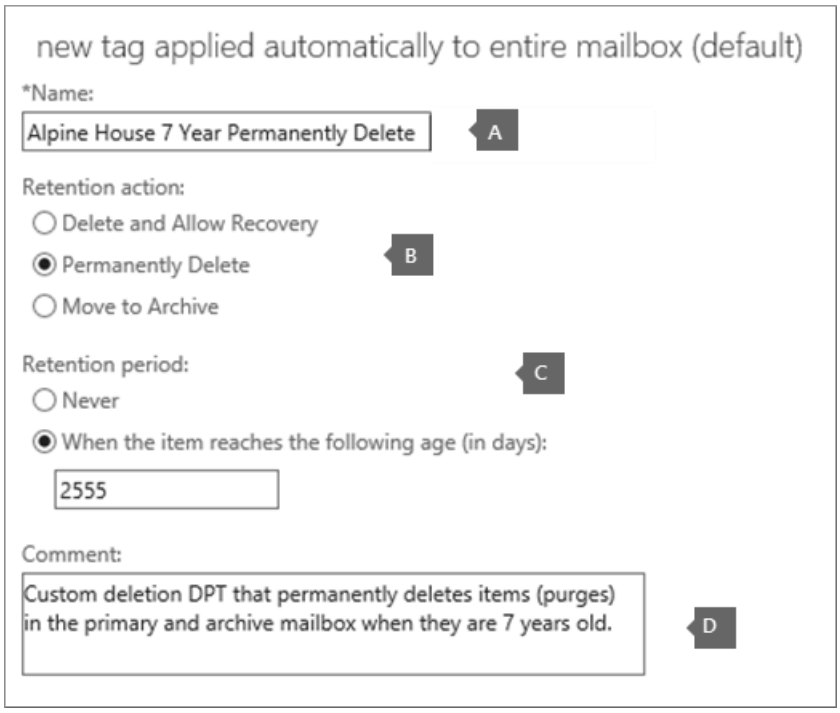

a. **Name** Type a name for the new retention tag.

b. **Retention action** Select **Permanently Delete** to purge items from the mailbox when the retention period expires.

c. **Retention period** Select **When the item reaches the following age (in days)**, and then enter the duration of the retention period. For this scenario, items will be purged after 2555 days (7 years).

d. **Comment** (Optional) Type a comment that explains the purpose of the custom retention tag.

3. Click **Save** to create the custom deletion DPT.

   The new deletion DPT is displayed in the list of retention tags.

Create a custom retention policy tag for the Deleted Items folder

The last retention tag that you'll create is a custom retention policy tag (RPT) for the Deleted Items folder. This tag will delete items in the Deleted Items folder after 5 years, and provides a recovery period when users can use the Recover Deleted Items tool to recover an item.

1. On the **Retention tags** page, click **New tag +**, and then select **applied automatically to a default folder**.

2. On the **New tag applied automatically to a default folder** page, complete the following fields:

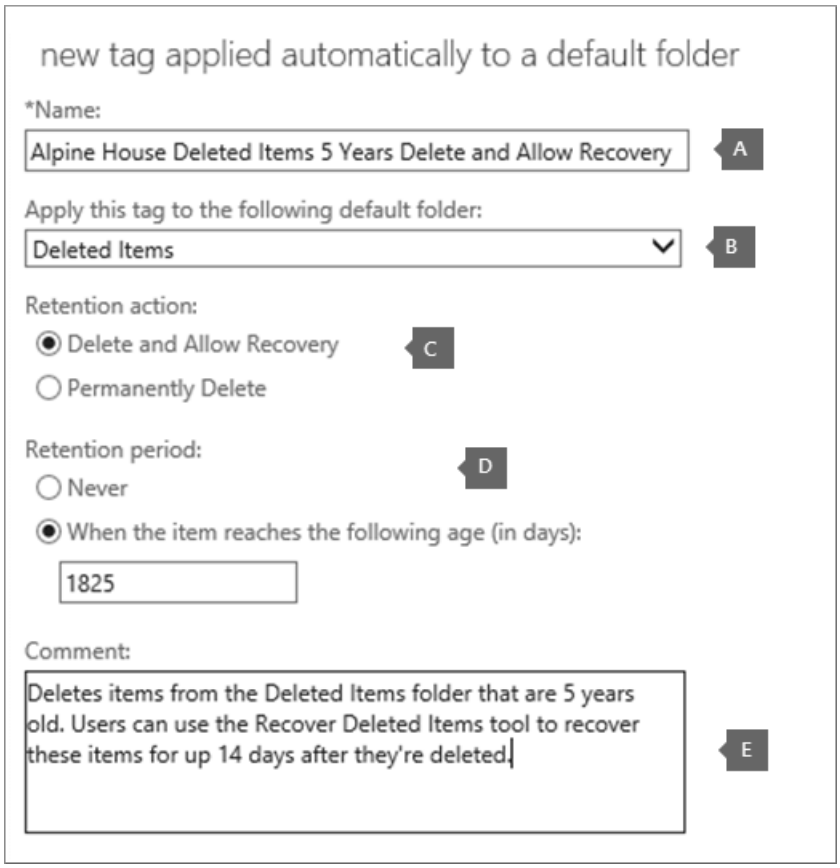

a. **Name** Type a name for the new retention tag.

b. **Apply this tag to the following default folder** In the drop-down list, select **Deleted Items**.

c. **Retention action** Select **Delete and Allow Recovery** to delete items when the retention period expires, but allow users to recover a deleted item within the deleted item retention period (which by default is 14 days).

d. **Retention period** Select **When the item reaches the following age (in days)**, and then enter the duration of the retention period. For this scenario, items will be deleted after 1825 days (5 years).

e. **Comment** (Optional) Type a comment that explains the purpose of the custom retention tag.

3. Click **Save** to create the custom RPT for the Deleted Items folder.

   The new RPT is displayed in the list of retention tags.

Step 3: Create a new retention policy

After you create the custom retention tags, the next step is to create a new retention policy and add the retention tags. You'll add the three custom retention tags that you created in Step 2, and the built-in tags that were mentioned in the first section. In Step 4, you'll assign this new retention policy to user mailboxes.

1. In the EAC, go to **Compliance management** > **Retention policies**.
2. On the **Retention policies** page, click **New +**.
3. In the **Name** box, type a name for the new retention policy; for example, **Alpine House Archive and Deletion Policy**.
4. Under **Retention tags**, click **Add +**.

   A list of the retention tags in your organization is displayed. Note the custom tags that you created in Step 2 are displayed.

5. Add the 9 retention tags that are highlighted in the following screenshot (these tags are described in more detail in the More information section). To add a retention tag, select it and then click **Add**.

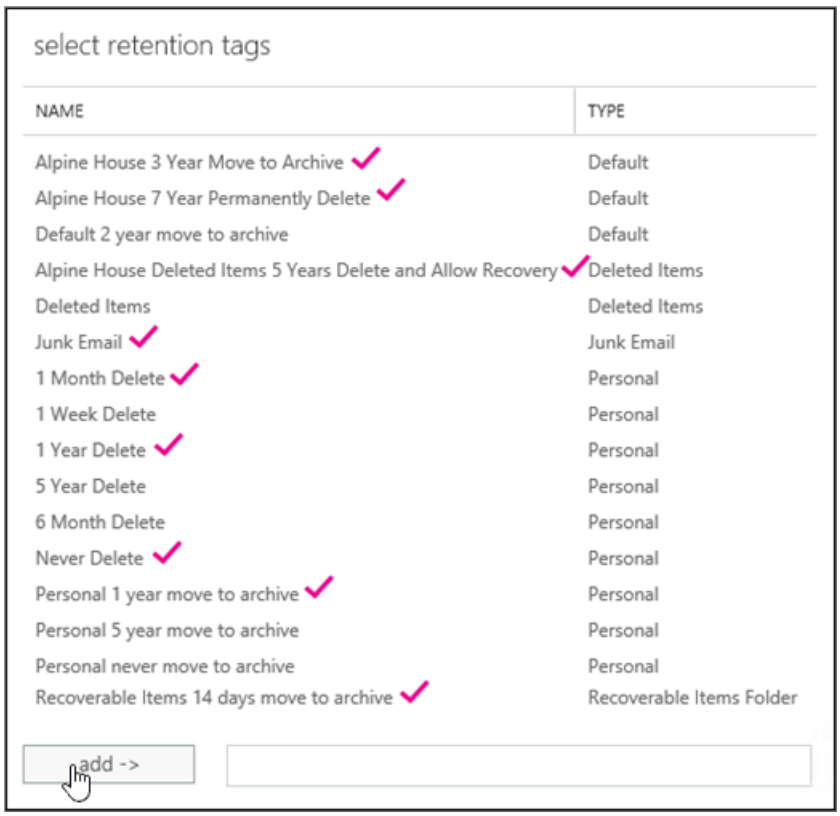

> **TIP**
> You can select multiple retention tags by holding down the **Ctrl** key and then clicking each tag.

6. After you've added the retention tags, click **OK**.

7. On the **New retention policy** page, click **Save** to create the new policy.

   The new retention policy is displayed in the list. Select it to display the retention tags linked to it in the details pane.

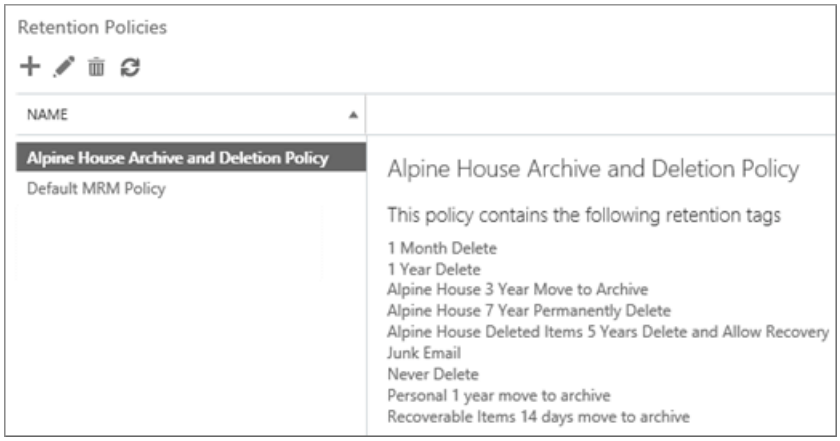

Step 4: Assign the new retention policy to user mailboxes

When a new mailbox is created, a retention policy named Default MRM policy is assigned to it by default. In this step, you'll replace this retention policy (because a mailbox can have only one retention policy assigned to it) by assigning the new retention policy that you created in Step 3 to the user mailboxes in your organization. This step assumes that you'll assign the new policy to all mailboxes in your organization.

1. In the EAC, go to **Recipients** > **Mailboxes**.

   A list of all user mailboxes in your organization is displayed.

2. Select all the mailboxes by clicking on the first one in the list, holding down the **Shift** key, and then clicking the last one in the list.

3. In the details pane on the right side of the EAC, under **Bulk Edit**, click **More options**.

4. Under **Retention Policy**, click **Update**.

5. On the **Bulk assign retention policy** page, in the **Select the retention policy** drop-down list, select the retention policy that you created in Step 3; for example, **Alpine House Archive and Retention Policy**.

6. Click **Save** to save the new retention policy assignment.

7. To verify that the new retention policy was assigned to mailboxes, you can do the following:

   a. Select a mailbox on the **Mailboxes** page, and then click **Edit** ✎.

   b. On the mailbox properties page for the selected user, click **Mailbox features**.

   The name of the new policy assigned to the mailbox is displayed in the **Retention policy** drop-down list.

(Optional) Step 5: Run the Managed Folder Assistant to apply the new settings

After you apply the new retention policy to mailboxes in Step 4, it can take up to 7 days in Exchange Online for the new retention settings to be applied to the mailboxes. This is because a process called the *Managed Folder Assistant* processes mailboxes at least once every 7 days. Instead of waiting for the Managed Folder Assistant to run, you can force this to happen by running the **Start-ManagedFolderAssistant** cmdlet in Exchange Online PowerShell.

What happens when you run the Managed Folder Assistant? It applies the settings in the retention policy by inspecting items in the mailbox and determining whether they're subject to retention. It then stamps items subject to retention with the appropriate retention tag, and then takes the specified retention action on items past their retention age.

Here are the steps to connect to Exchange Online PowerShell, and then run the Managed Folder Assistant on every mailbox in your organization.

1. Connect to Exchange Online PowerShell.

2. Run the following two commands to start the Managed Folder Assistant for all user mailboxes in your organization.

```
$Mailboxes = Get-Mailbox -ResultSize Unlimited -Filter {RecipientTypeDetails -eq "UserMailbox"}
```

```
$Mailboxes.Identity | Start-ManagedFolderAssistant
```

That's it! You've set up an archive and deletion policy for the Alpine House organization.

NOTE

As previously stated, the Managed Folder Assistant processes mailboxes at least once every 7 days. So it's possible that a mailbox can be processed by the Managed Folder Assistant more frequently. Also, admins can't predict the next time a mailbox is processed by the Managed Folder Assistant, which is one reason why you may want to run it manually. However, if you want to temporarily prevent the Managed Folder Assistant from applying the new retention settings to a mailbox, you can run the `Set-Mailbox -ElcProcessingDisabled $true` command to temporarily disable the the Managed Folder Assistant from processing a mailbox. To re-enable the Managed Folder Assistant for a mailbox, run the `Set-Mailbox -ElcProcessingDisabled $false` command. Finally, if a mailbox user has a disabled account, we will not process the move items to archive action for that mailbox.

(Optional) Step 6: Make the new retention policy the default for your organization

In Step 4, you have to assign the new retention policy to existing mailboxes. But you can configure Exchange Online so that the new retention policy is assigned to new mailboxes that are created in the future. You do this by using Exchange Online PowerShell to update your organization's default mailbox plan. A *mailbox plan* is a template that automatically configures properties on new mailboxes. In this optional step, you can replace the current retention policy that's assigned to the mailbox plan (by default, the Default MRM Policy) with the retention policy that you created in Step 3. After you update the mailbox plan, the new retention policy will be assigned to new mailboxes.

1. Connect to Exchange Online PowerShell.
2. Run the following command to display information about the mailbox plans in your organization.

   ```
   Get-MailboxPlan | Format-Table DisplayName,RetentionPolicy,IsDefault
   ```

   Note the mailbox plan that's set as the default.
3. Run the following command to assign the new retention policy that you created in Step 3 (for example, **Alpine House Archive and Retention Policy**) to the default mailbox plan. This example assumes the name of the default mailbox plan is **ExchangeOnlineEnterprise**.

   ```
   Set-MailboxPlan "ExchangeOnlineEnterprise" -RetentionPolicy "Alpine House Archive and Retention
   Policy"
   ```
4. You can rerun the command in step 2 to verify that the retention policy assigned to the default mailbox plan was changed.

More information

- How is retention age calculated? The retention age of mailbox items is calculated from the date of delivery or the date of creation for items such as draft messages that aren't sent but are created by the user. When the Managed Folder Assistant processes items in a mailbox, it stamps a start date and an expiration date for all items that have retention tags with the Delete and Allow Recovery or Permanently Delete retention action. Items that have an archive tag are stamped with a move date.
- The following table provides more information about each retention tag that is added to the custom retention policy that was created by following the steps in this topic.

RETENTION TAG	WHAT THIS TAG DOES	BUILT-IN OR CUSTOM?	TYPE
Alpine House 3 Year Move to Archive	Moves items that are 1095 days (3 years) old to the archive mailbox.	Custom (See Step 2: Create new retention tags for the archive and deletion policies)	Default Policy Tag (archive); this tag is automatically applied to the entire mailbox.
Alpine House 7 Year Permanently Delete	Permanently deletes items in the primary mailbox or the archive mailbox when they are 7 years old.	Custom (See Step 2: Create new retention tags for the archive and deletion policies)	Default Policy Tag (deletion); this tag is automatically applied to the entire mailbox.
Alpine House Deleted Items 5 Years Delete and Allow Recovery	Deletes items from the Deleted Items folder that are 5 years old. Users can recover these items for up 14 days after they're deleted.*	Custom (See Step 2: Create new retention tags for the archive and deletion policies)	Retention Policy Tag (Deleted Items); this tag is automatically applied to items in the Deleted items folder.
Recoverable Items 14 days Move to Archive	Moves items that have been in the Recoverable Items folder for 14 days to the Recoverable Items folder in the archive mailbox.	Built-in	Retention Policy Tag (Recoverable Items); this tag is automatically applied to items in the Recoverable Items folder.
Junk Email	Permanently deletes items that have been in the Junk Email folder for 30 days. Users can recover these items for up 14 days after they're deleted.*	Built-in	Retention Policy Tag (Junk Email); this tag is automatically applied to items in Junk Email folder.
1 Month Delete	Permanently deletes items that are 30 days old. Users can recover these items for up 14 days after they're deleted.*	Built-in	Personal; this tag can be applied by users.
1 Year Delete	Permanently deletes items that are 365 days old. Users can recover these items for up 14 days after they're deleted.*	Built-in	Personal; this tag can be applied by users.
Never Delete	This tag prevents items from being deleted by a retention policy.	Built-in	Personal; this tag can be applied by users.
Personal 1 year move to archive	Moves items to the archive mailbox after 1 year.	Built-in	Personal; this tag can be applied by users.

* Users can use the Recover Deleted Items tool in Outlook and Outlook on the web (formerly known as Outlook Web App) to recover a deleted item within the deleted item retention period, which by default is 14 days in Exchange Online. An administrator can use Windows PowerShell to increase the deleted item retention period to a maximum of 30 days. For more information, see: Recover deleted items in Outlook for Windows and Change the deleted item retention period for a mailbox in

> Exchange Online

- Using the **Recoverable Items 14 days Move to Archive** retention tag helps free up storage space in the Recoverable Items folder in the user's primary mailbox. This is useful when a user's mailbox is placed on hold, which means nothing is ever permanently deleted the user's mailbox. Without moving items to the archive mailbox, it's possible the storage quota for the Recoverable Items folder in the primary mailbox will be reached. For more information about this and how to avoid it, see Increase the Recoverable Items quota for mailboxes on hold.

Overview of inactive mailboxes

11/2/2020 • 12 minutes to read • Edit Online

Your organization might need to retain former employees' email after they leave the organization. Depending on your organization's retention requirements, you might need to retain mailbox content for a few months or years after employment ends, or you might need to retain mailbox content indefinitely. Regardless of how long you need to retain email, you can create inactive mailboxes to retain the mailbox of former employees.

What are inactive mailboxes?

When an employee leaves your organization (or goes on an extended leave of absence), you can remove their Microsoft 365 account. The employee's mailbox data is retained for 30 days after the account is removed. During this period, you can still recover the mailbox data by undeleting the account. After 30 days, the data is permanently removed.

But if your organization needs to retain mailbox content for former employees, you can turn the mailbox into an inactive mailbox by placing the mailbox on Litigation Hold or applying a Microsoft 365 retention policy to the mailbox in the Security & Compliance Center and then removing the corresponding Microsoft 365 account. The contents of an inactive mailbox are retained for the duration of the Litigation Hold placed on the mailbox or the retention period of the retention policy applied to it before the mailbox was deleted. You can still recover the corresponding user account for a 30-day period. However, after 30 days, the inactive mailbox is retained in Microsoft 365 until the hold or retention policy is removed.

> **IMPORTANT**
> As we continue to invest in different ways to preserve mailbox content, we're announcing the retirement of In-Place Holds in the Exchange admin center. That means you should use Litigation Holds and Microsoft 365 retention policies to create an inactive mailbox. Starting July 1, 2020 you won't be able to create new In-Place Holds in Exchange Online. But you'll still be able to change the hold duration of an In-Place Hold placed on an inactive mailbox. However, starting October 1, 2020, you won't be able to change the hold duration. You'll only be able to delete an inactive mailbox by removing the In-Place Hold. Existing inactive mailboxes that are on In-Place Hold will still be preserved until the hold is removed. For more information about when In-Place Holds will be retired, see Retirement of legacy eDiscovery tools.

Inactive mailboxes and Microsoft 365 retention policies

In addition to Litigation Hold, using the new Microsoft 365 retention policy feature in the Security & Compliance Center is another way to make a mailbox inactive. To use a retention policy to make an inactive mailbox:

- It has to be configured to retain content or retain and then delete content. If a retention policy is configured to only delete content, a mailbox that the policy is applied to won't become inactive when the mailbox is deleted.
- It has to be applied to Exchange mailboxes or Skype for Business locations (because Skype-related content is stored in the user's mailbox).
- It can be query-based so that it retains only items that match a search query.

For more information about retention policies, see Learn about retention policies and retention labels.

If you use a retention policy to make an inactive mailbox, Microsoft 365 continues to process the retention policy on the inactive mailbox. This means if the retention policy is configured to retain and then delete content,

items will be moved to the Recoverable Items folder when the retention duration expires, and then eventually purged from the inactive mailbox. If retention policy isn't configured to deleted items, then items that haven't been permanently deleted by the user (before the mailbox was made inactive) won't be moved to the Recoverable Items folder and will be retained indefinitely after the mailbox becomes inactive.

You might consider creating a Microsoft 365 retention policy specifically for inactive mailboxes. Here are some reasons for doing this and things to keep in mind.

- You can configure the retention policy to retain mailbox content only as long as necessary to meet your organization's requirement for former employees.
- It's a good way to identify inactive mailboxes because the retention policy will only be applied to inactive mailboxes.
- You are able to quickly identify the retention policy that's assigned to inactive mailboxes in your organization. This makes it easier to change the retention (or deletion) settings if necessary. It will also make it easier to permanently delete an inactive mailbox because you can remove it from the policy by using the Security & Compliance Center. Otherwise, you have to use Exchange Online PowerShell to remove a Litigation Hold from an inactive mailbox or use Security & Compliance Center PowerShell to exclude an inactive mailbox from an organization-wide Microsoft 365 retention policy.
- If you create a Microsoft 365 retention policy specifically for inactive mailboxes, you can add a maximum of 1,000 mailboxes to the policy. If you're a large organization, you might have to create more than one Microsoft 365 retention policy to use for inactive mailboxes.

Caution

If you use a retention policy to make a mailbox inactive, do not change or remove the user principal name (UPN) for the mailbox before you delete the corresponding user account. Additionally, do not change the primary SMTP address (that's derived from the UPN) or remove this email address from the list of secondary SMTP addresses associated with the mailbox before making the mailbox inactive. If you change the UPN or email addresses (that were assigned to the mailbox at the time the retention policy was applied to it) and then delete the user account to make the mailbox inactive, you won't be able to delete the inactive mailbox when you no longer need to retain it. That's because you can't remove the inactive mailbox from the retention policy using a UPN or email address (to identify the inactive mailbox) that's different than the ones that existed when the retention policy was initially applied to the mailbox. For more information about deleting inactive mailboxes, see Delete an inactive mailbox in Office 365.

Inactive mailboxes and eDiscovery case holds

If a hold that's associated with an eDiscovery case in the Security & Compliance Center is placed on a mailbox and then the mailbox or the user's account is deleted, the mailbox becomes an inactive mailbox. However, we don't recommend using eDiscovery case holds to make a mailbox inactive. That's because eDiscovery cases are intended for specific, time-bound cases related to a legal issue. At some point, a legal case will probably end and the holds associated with the case will be removed and the eDiscovery case will be closed. In fact, if a hold that's placed on an inactive mailbox is associated with an eDiscovery case, and then the hold is released or the eDiscovery case is closed (or deleted), the inactive mailbox will be permanently deleted. Also, you can't create a time-based eDiscovery hold. That's means that content in an inactive mailbox is retained forever or until the hold is removed and the inactive mailbox is deleted. Therefore, we recommend using a Litigation Hold or a retention policy for inactive mailboxes.

For more information about eDiscovery cases and holds, see eDiscovery cases.

Inactive mailboxes and labels

Retention labels help you classify email data in your organization for governance, and enforce retention rules based on that classification. A retention label can be applied to an email item either manually by users or

automatically by administrators, and an email item can only have single label assigned to it. If a single email item in a user's mailbox has a label assigned to it (and it's configured to retain or retain and then delete the item) and the mailbox or the user's account is deleted, the mailbox becomes an inactive mailbox. Similar to eDiscovery case holds, we don't recommend using retention labels to make a mailbox inactive. Instead, we recommend that you use a Litigation Hold or a retention policy. In the case of retention labels, you might not realize that a retention label has been applied to an email item and then inadvertently make an inactive mailbox when you delete the user's account.

For more information about retention policies and retention labels, see Learn about retention policies and retention labels in Office 365.

Inactive mailboxes and auto-expanding archives

An inactive mailbox that's configured with an auto-expanding archive can't be recovered or restored. In situations where it's necessary to recover data from an inactive mailbox with an auto-expanding archive, we recommended that you use the content search tool to export the data from the mailbox and then import to another mailbox. For step-by-step instructions to search an inactive mailbox and export the search results, see:

- Content search
- Export content search results

Inactive mailboxes and Exchange MRM retention policies

If an Exchange retention policy (the Messaging Records Management, or MRM, feature in Exchange Online) was applied to mailbox when it was made inactive, any deletion policies (which are retention tags configured with a **Delete** retention action) will continue to be processed on the inactive mailbox. That means items that are tagged with a deletion policy will be moved to the Recoverable Items folder when the retention period expires. Those items are purged from the inactive mailbox when the hold duration expires. If a hold duration isn't specified for the inactive mailbox, items in the Recover Items folder will be retained indefinitely.

Conversely, any archive policies (which are retention tags configured with a **MoveToArchive** retention action) that are included in the retention policy assigned to an inactive mailbox are ignored. That means items in an inactive mailbox that are tagged with an archive policy remain in the primary mailbox when the retention period expires. They're not moved to the archive mailbox or to the Recoverable Items folder in the archive mailbox. They will be retained indefinitely.

Creating an inactive mailbox

To make a mailbox inactive, it must be assigned an Exchange Online Plan 2 license (or an Exchange Online Plan 1 license with an Exchange Online Archiving add-on license) so that a Litigation Hold or Microsoft 365 retention policy can be applied to the mailbox before it's deleted. After the mailbox is deleted, any Exchange Online license associated with it will be available to assign to a new user.

The following table summarizes the process of making an inactive mailbox for different retention scenarios. For more information, see Manage inactive mailboxes.

TO...	DO THIS...	RESULT

TO...	DO THIS...	RESULT
Retain mailbox content indefinitely after an employee leaves the organization	Place the mailbox on Litigation Hold or apply a Microsoft 365 retention policy (that's configured to retain content) to the mailbox. Don't specify a hold duration for the Litigation Hold or don't configure the retention policy to delete items. Alternatively you can use a retention policy that retains items forever. Remove the user's Microsoft 365 account.	All content in the inactive mailbox, including items in the Recoverable Items folder, is retained indefinitely.
Retain mailbox content for a specific period after an employee leaves the organization and then delete it	Apply a Microsoft 365 retention policy to the mailbox. Configure the retention policy to retain and then delete items when the retention period expires. Remove the user's Microsoft 365 account.	When the retention period for a mailbox item expires, the item is moved to the Recoverable Items folder and then it's permanently deleted (purged) from the inactive mailbox when the deleted item retention period (for Exchange mailboxes) expires. The retention period of the Microsoft 365 retention policy can be configured based on the original date a mailbox item was received or created, or when it was last modified.

NOTE: If a Litigation Hold is already placed on a mailbox, or if a Microsoft 365 retention policy (that's configured to retain or retain and then delete content) is already applied to the mailbox, then all you have to do is delete the corresponding user account to create an inactive mailbox.

Managing inactive mailboxes

After you make a mailbox inactive, you can perform various management tasks on inactive mailboxes.

- **Change the hold duration for an inactive mailbox.** After a mailbox is made inactive, you can change the hold duration of the Litigation Hold or Microsoft 365 retention policy applied to the inactive mailbox. For step-by-step procedures, see Change the hold duration for an inactive mailbox.

  > **NOTE**
  > You can't apply other retention policies to an inactive mailbox. You can only change the retention duration of an existing retention policy applied to the inactive mailbox.

- **Recover an inactive mailbox.** If a former employee (or an employee on a leave of absence) returns to your organization, or if a new employee is hired to take on the job responsibilities of the former employee, you can recover the contents of the inactive mailbox. When you recover an inactive mailbox, the mailbox is converted to a new mailbox, the contents and the folder structure of the inactive mailbox are retained, and the mailbox is linked to a new user account. After it's recovered, the inactive mailbox no longer exists. For step-by-step procedures and information about what happens when you recover an inactive mailbox, see Recover an inactive mailbox.

> **NOTE**
>
> If you recover an inactive mailbox that was assigned to a retention policy with Preservation Lock (called a *locked retention policy*), the recovered mailbox is assigned to the same locked retention policy. If you recover an inactive mailbox that was assigned to a retention policy without Preservation Lock, the recovered mailbox is removed from the unlocked retention policy. However, Litigation Hold is enabled on the recovered mailbox to prevent the deletion of mailbox content based on any organization-wide retention policies that delete content older than a specific age.

- **Restore an inactive mailbox.** If another employee takes on the job responsibilities of a former employee, or if another person needs access to the contents of the inactive mailbox, you can restore (or merge) the contents of the inactive mailbox to an existing mailbox. When you restore an inactive mailbox, the contents are copied to another mailbox. The inactive mailbox is retained and remains an inactive mailbox. The inactive mailbox can still be searched using eDiscovery tools, its contents can be restored to another mailbox, and it can be recovered or deleted later. For step-by-step procedures, see Restore an inactive mailbox.
- **Delete an inactive mailbox.** When you no longer need to retain the contents of an inactive mailbox, you can permanently delete it by removing all holds or Microsoft 365 retention policies applied to the inactive mailbox. If a mailbox was made inactive more than 30 days ago, it will be marked for permanent deletion after you remove the hold. If the mailbox was made inactive within the last 30 days, you can make it active again after removing the hold or retention policy. For step-by-step procedures, see Delete an inactive mailbox.

Create and manage inactive mailboxes

11/2/2020 • 9 minutes to read • Edit Online

Microsoft 365 makes it possible for you to retain the contents of deleted mailboxes. This feature is called inactive mailboxes. Inactive mailboxes allow you to retain former employees' email after they leave your organization. A mailbox becomes inactive when a Litigation Hold or a retention policy (created in the security and compliance center in Office 365 or Microsoft 365) is applied to the mailbox before the corresponding user account is deleted. The contents of an inactive mailbox are retained for the duration of the hold that was placed on the mailbox before it was made inactive. This allows administrators, compliance officers, and records managers to use Content Search to search and export the contents of an inactive mailbox. Inactive mailboxes can't receive email and aren't displayed in your organization's shared address book or other lists.

> **IMPORTANT**
>
> As we continue to invest in different ways to preserve mailbox content, we're announcing the retirement of In-Place Holds in the Exchange admin center. That means you should use Litigation Holds and retention policies to create an inactive mailbox. Starting July 1, 2020 you won't be able to create new In-Place Holds in Exchange Online. But you'll still be able to change the hold duration of an In-Place Hold placed on an inactive mailbox. However, starting October 1, 2020, you won't be able to change the hold duration. You'll only be able to delete an inactive mailbox by removing the In-Place Hold. Existing inactive mailboxes that are on In-Place Hold will still be preserved until the hold is removed. For more information about the retirement of In-Place Holds, see Retirement of legacy eDiscovery tools.

Preparations before creating an inactive mailbox

- To make a mailbox inactive, it must be assigned an Exchange Online Plan 2 license so that a Litigation Hold or a retention policy can be applied to the mailbox before it's deleted. Exchange Online Plan 2 licenses are part of an Office 365 Enterprise E3 and E5 subscription. If a mailbox is assigned an Exchange Online Plan 1 or Exchange Online Kiosk license (which are part of an Office 365 E1 and F1 subscription respectively), you would have to assign it a separate Exchange Online Archiving license so that a hold can be applied to the mailbox before it's deleted. For more information, see Exchange Online Archiving.
- The licenses associated with the deleted Exchange Online mailbox will be available after you delete the corresponding user account. You can then assign those licenses to another user.
- If a Litigation Hold or a retention policy (that's configured to retain or retain and then delete content) isn't applied to a mailbox before it's deleted, the contents of the mailbox won't be retained or discoverable. However, the deleted mailbox can be recovered within 30 days of deletion, but the mailbox and its contents will be permanently deleted after 30 days if it isn't recovered.
- For more information about Litigation Hold, see In-Place Hold and Litigation Hold. For more information about retention policies, see Learn about retention policies and retention labels.

Create an inactive mailbox

Making a mailbox inactive involves two steps: 1) placing the mailbox on Litigation Hold or applying a retention policy to it, and 2) deleting the mailbox or corresponding user account. After the mailbox is inactive, its contents are retained until the hold or retention policy is removed.

Step 1: Place a mailbox on Litigation Hold or apply a retention policy

Placing a mailbox on Litigation Hold or applying a retention policy (that's configured to retain or retain and then

delete content) retains the contents in the mailbox before it's deleted. Both types of holds will retain all mailbox content, including deleted items and original versions of modified items. Deleted and modified items are retained in the inactive mailbox for a specified period, or until you permanently delete the inactive mailbox by removing the hold or retention policy that's applied to the inactive mailbox.

If a hold is already placed on a mailbox, or if a retention policy is already applied to a mailbox, then all you have to do is delete the corresponding user account as explained in Step 2.

For step-by-step procedures for placing a mailbox on Litigation Hold or applying a retention policy, see:

- Place a mailbox on Litigation Hold
- Learn about retention policies and retention labels in Office 365

> **NOTE**
> For Litigation Holds and retention policies, you can create an indefinite hold or on a time-based hold. In an indefinite hold, the contents of the inactive mailbox will be retained forever, or until the hold is removed or until the hold duration is changed. After the hold or retention policy is removed (assuming that the mailbox was deleted more than 30 days ago), the inactive mailbox will be marked for permanent deletion and the contents of the mailbox will no longer be retained or discoverable. In a time-based hold or retention policy, you specify the duration of the hold. This duration is on a per-item basis and is calculated from the date a mailbox item was received or created. After the hold expires for a mailbox item, and that item moved to or is located in the Recoverable Items folder in the inactive mailbox, the item is permanently deleted (purged) from the inactive mailbox after the deleted item retention period expires.

Step 2: Delete the mailbox

After the mailbox is placed on hold or a retention policy is applied to it, the next step is to delete the mailbox. The best way to delete a mailbox is to delete the corresponding user account in the Microsoft 365 admin center. For information about deleting user accounts, see Delete a user from your organization.

> **NOTE**
> You can also delete the mailbox by using the **Remove-Mailbox** cmdlet in Exchange Online PowerShell. For more information, see Delete or restore user mailboxes in Exchange Online.

View a list of inactive mailboxes

To view a list of the inactive mailboxes in your organization:

1. Go to https://protection.office.com and sign in using the credentials for an administrator account in your organization.
2. Click **Information governance** > **Retention**.
3. On the **Retention** page, click **More**, and then click **Inactive mailboxes**.

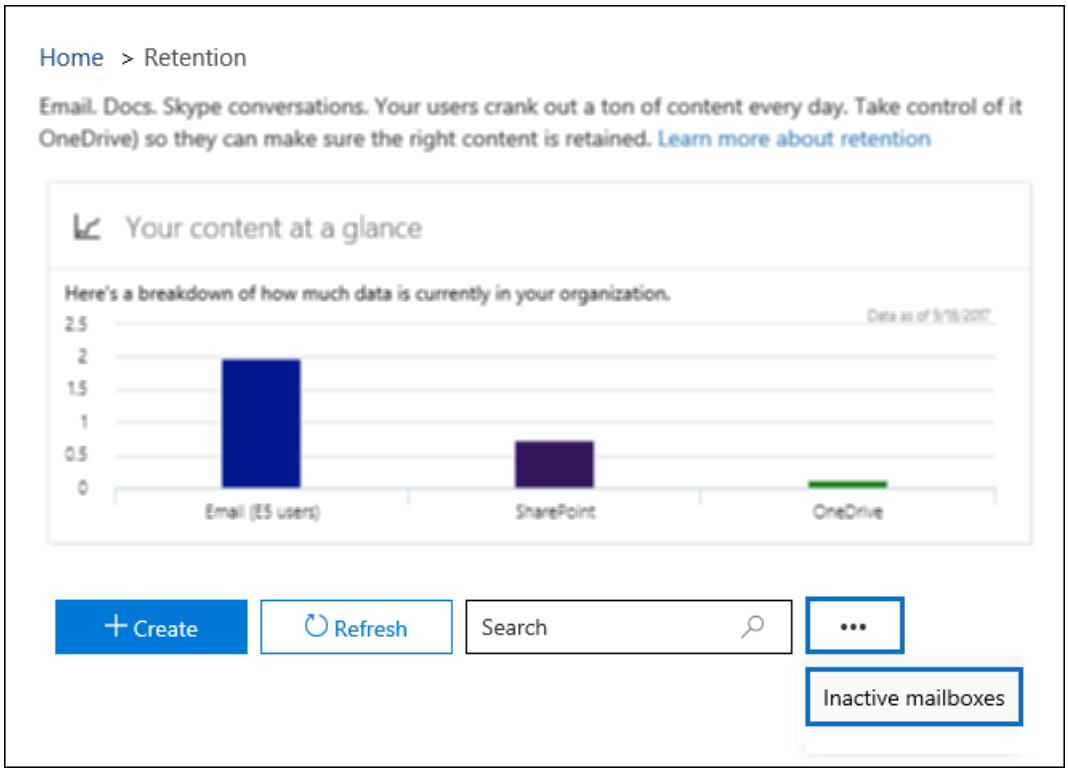

The **Inactive mailboxes** page is displayed. Note the total number of inactive mailboxes in your organization is displayed.

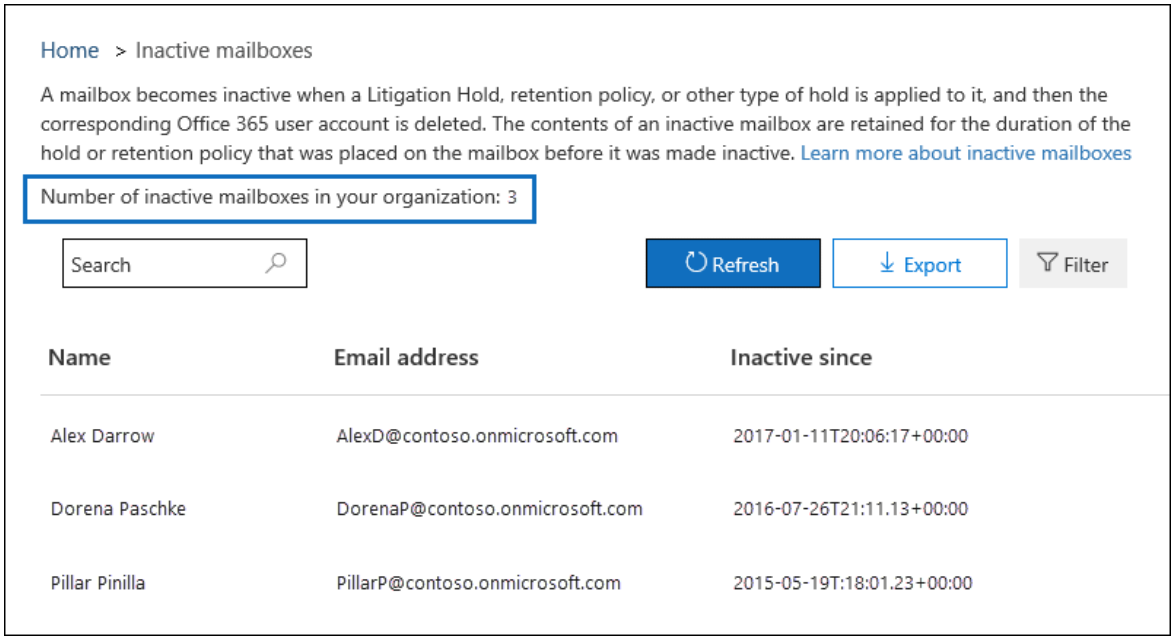

Alternatively, you can run the following command in Exchange Online PowerShell to display the list of inactive mailboxes.

```
Get-Mailbox -InactiveMailboxOnly | FT DisplayName,PrimarySMTPAddress,WhenSoftDeleted
```

You can click **Export** to view or download a CSV file that contains additional information about the inactive mailboxes in your organization.

You can also run the following command to export the list of inactive mailboxes and other information to a CSV file. In this example, the CSV file is created in the current directory.

```
Get-Mailbox -InactiveMailboxOnly | Select
Displayname,PrimarySMTPAddress,DistinguishedName,ExchangeGuid,WhenSoftDeleted | Export-Csv
InactiveMailboxes.csv -NoType
```

NOTE

It's possible that an inactive mailbox may have the same SMTP address as an active user mailbox. In this case, the value of the **DistinguishedName** or **ExchangeGuid** property can be used to uniquely identify an inactive mailbox.

Search and export the contents of an inactive mailbox

You can access the contents of the inactive mailbox by using the Content Search tool in the Security & Compliance Center. When you search an inactive mailbox, you can create a keyword search query to search for specific items or you can return the entire contents of the inactive mailbox. You can preview the search results or export the search results to an Outlook Data (PST) file or as individual email messages. For step-by-step procedures for searching mailboxes and exporting search results, see the following topics:

- Content Search in Office 365
- Export Content Search results

Here are a few things to keep in mind when searching inactive mailboxes.

- If a content search includes a user mailbox and that mailbox is made inactive, the content search will continue to search the inactive mailbox when you rerun the search after it becomes inactive.
- In some cases, a user may have an active mailbox and an inactive mailbox that have the same SMTP address. In this case, only the specific mailbox that you select as a location for a content search will be searched. In other words, if you add a user's mailbox to a search, you can't assume that both their active and inactive mailboxes will be searched; only the mailbox that you explicitly add to the search will be searched.
- We strongly recommend that you avoid having an active mailbox and inactive mailbox with the same SMTP address. If you need to reuse the SMTP address that is currently assigned to an inactive mailbox, we recommend that you recover the inactive mailbox or restore the contents of an inactive mailbox to an active mailbox (or the archive of an active mailbox), and then delete the inactive mailbox.

Change the hold duration for an inactive mailbox

After a mailbox is made inactive, you can change the duration of the hold or the retention policy applied to the inactive mailbox. For step-by-step procedures, see Change the hold duration for an inactive mailbox in Office 365.

Recover an inactive mailbox

If a former employee returns to your organization, or if a new employee is hired to take on the job responsibilities of the departed employee, you can recover the contents of the inactive mailbox. When you recover an inactive mailbox, the mailbox is converted to a new mailbox, the contents and folder structure of the inactive mailbox are retained, and the mailbox is linked to a new user account. After it's recovered, the inactive mailbox no longer exists. For step-by-step procedures and more information about happens when you recover an inactive mailbox, see Recover an inactive mailbox in Office 365.

Restore the contents of an inactive mailbox to another mailbox

If another employee takes on the job responsibilities of a former employee, or if another person needs access to the contents of the inactive mailbox, you can restore (or merge) the contents of the inactive mailbox to an existing mailbox. When you restore an inactive mailbox, the contents are copied to another mailbox. The inactive mailbox is retained and remains an inactive mailbox. The inactive mailbox can still be searched using eDiscovery, its contents can be restored to another mailbox, or it can be recovered or deleted at a later date. For step-by-step procedures, see Restore an inactive mailbox in Office 365.

Delete an inactive mailbox

If you no longer need to retain the contents of an inactive mailbox, you can permanently delete the inactive mailbox by removing the hold or removing the retention policy applied to the inactive mailbox. If the mailbox was deleted more than 30 days ago, the mailbox will be marked for permanent deletion after you remove the hold, and the mailbox will become non-recoverable. If the mailbox was deleted within the last 30 days, you can still recover the mailbox after removing the hold or retention policy. For step-by-step procedures for removing a hold or a retention policy to permanently delete an inactive mailbox, see Delete an inactive mailbox.

Change the hold duration for an inactive mailbox

11/2/2020 • 10 minutes to read • Edit Online

An inactive mailbox is used to retain a former employee's email after he or she leaves your organization. A mailbox becomes inactive when a Litigation Hold, an In-Place Hold, a Microsoft 365 retention policy, or a hold that's associated with an eDiscovery case is placed on the mailbox, and the corresponding user account is deleted. The contents of an inactive mailbox are retained for the duration of the hold that was placed on the mailbox before it was made inactive. The hold duration defines how long items in the Recoverable Items folder are held. When the hold duration expires for an item in the Recoverable Items folder, the item is permanently deleted (purged) from the inactive mailbox. After a mailbox is made inactive, you can change the duration of the hold or Microsoft 365 retention policy assigned to the inactive mailbox.

> **IMPORTANT**
>
> As we continue to invest in different ways to preserve mailbox content, we're announcing the retirement of In-Place Holds in the Exchange admin center. That means you should use Litigation Holds and Microsoft 365 retention policies to create an inactive mailbox. Starting April 1, 2020 you won't be able to create new In-Place Holds in Exchange Online. But you'll still be able to change the hold duration of an In-Place Hold placed on an inactive mailbox. However, starting July 1, 2020, you won't be able to change the hold duration. You'll only be able to delete an inactive mailbox by removing the In-Place Hold. Existing inactive mailboxes that are on In-Place Hold will still be preserved until the hold is removed. For more information about the retirement of In-Place Holds, see Retirement of legacy eDiscovery tools.

Connect to PowerShell

- You have to use Exchange Online PowerShell to change the hold duration for a Litigation Hold on an inactive mailbox. You can't use the Exchange admin center (EAC). But you can use Exchange Online PowerShell or the EAC to change the hold duration for an In-Place Hold. You can use the security and compliance center or the Security & Compliance Center PowerShell to change the hold duration for a Microsoft 365 retention policy.
- To connect to Exchange Online PowerShell or Security & Compliance Center PowerShell, see one of the following topics:
  - Connect to Exchange Online PowerShell
  - Connect to Security & Compliance Center PowerShell
- Holds associated with eDiscovery cases are infinite holds, which means there's no hold duration that can be changed. Items are held forever or until the hold is removed and the inactive mailbox is deleted.
- For more information about inactive mailboxes, see Inactive mailboxes in Microsoft 365.

Step 1: Identify the holds on an inactive mailbox

Because different types of holds or one or more Microsoft 365 retention policies might be placed on an inactive mailbox, the first step is to identify the holds on an inactive mailbox.

Run the following command in Exchange Online PowerShell to display the hold information for all inactive mailboxes in your organization.

```
Get-Mailbox -InactiveMailboxOnly | FL
DisplayName,Name,IsInactiveMailbox,LitigationHoldEnabled,LitigationHoldDuration,InPlaceHolds
```

The value of **True** for the **LitigationHoldEnabled** property indicates that the inactive mailbox is on Litigation Hold. If an In-Place Hold, eDiscovery hold, or Microsoft 365 retention policy is placed on an inactive mailbox, a GUID for the hold or retention policy is displayed as the value for the **InPlaceHolds** property. For example, the following shows results for five inactive mailboxes.

```
DisplayName           : Ann Beebe
Name                  : annb
IsInactiveMailbox     : True
LitigationHoldEnabled : True
LitigationHoldDuration: 365.00:00:00
InPlaceHolds          : {}
...
DisplayName           : Pilar Pinilla
Name                  : pilarp
IsInactiveMailbox     : True
LitigationHoldEnabled : False
LitigationHoldDuration: Unlimited
InPlaceHolds          : {c0ba3ce811b6432a8751430937152491}
...
DisplayName           : Mario Necaise
Name                  : marion
IsInactiveMailbox     : True
LitigationHoldEnabled : False
LitigationHoldDuration: Unlimited
InPlaceHolds          : {}
...
DisplayName           : Carol Olson
Name                  : carolo
IsInactiveMailbox     : True
LitigationHoldEnabled : False
LitigationHoldDuration: Unlimited
InPlaceHolds          : {mbxcdbbb86ce60342489bff371876e7f224}
...
DisplayName           : Abraham McMahon
Name                  : abrahamm
IsInactiveMailbox     : True
LitigationHoldEnabled : False
LitigationHoldDuration: Unlimited
InPlaceHolds          : {UniH7d895d48-7e23-4a8d-8346-533c3beac15d}
```

The following table identifies the five different hold types that were used to make each mailbox inactive.

INACTIVE MAILBOX	HOLD TYPE	HOW TO IDENTIFY THE HOLD ON THE INACTIVE MAILBOX
Ann Beebe	Litigation Hold	The *LitigationHoldEnabled* property is set to `True`.

INACTIVE MAILBOX	HOLD TYPE	HOW TO IDENTIFY THE HOLD ON THE INACTIVE MAILBOX
Pilar Pinilla	In-Place Hold	The *InPlaceHolds* property contains the GUID of the In-Place Hold that's placed on the inactive mailbox. You can tell this is an In-Place Hold because the ID doesn't start with a prefix. You can use the `Get-MailboxSearch -InPlaceHoldIdentity <hold GUID> | FL` command in Exchange Online PowerShell to get information about the In-Place Hold on the inactive mailbox.

Mario Necaise	Organization-wide Microsoft 365 retention policy in the Security & Compliance Center	The *InPlaceHolds* property is empty. This indicates that one or more organization-wide or (Exchange-wide) Microsoft 365 retention policy is applied to the inactive mailbox. In this case, you can run the
		`Get-OrganizationConfig | Select-Object -ExpandProperty InPlaceHolds`
		command in Exchange Online PowerShell to get a list of the GUIDs for organization-wide Microsoft 365 retention policies. The GUID for organization-wide retention policies that are applied to Exchange mailboxes start with the `mbx` prefix; for example, `mbxa3056bb15562480fadb46ce523ff7b02`.
		
		To identity the Microsoft 365 retention policy that's applied to the inactive mailbox, run the following command in Security & Compliance Center PowerShell.
		
		`Get-RetentionCompliancePolicy <retention policy GUID without prefix> | FL Name`

INACTIVE MAILBOX	HOLD TYPE	HOW TO IDENTIFY THE HOLD ON THE INACTIVE MAILBOX
Carol Olson	Microsoft 365 retention policy in the Security & Compliance Center applied to specific mailboxes	The *InPlaceHolds* property contains the GUID of the Microsoft 365 retention policy that's applied to the inactive mailbox. You can tell this is a retention policy that applied to specific mailboxes because the GUID starts with the `mbx` prefix. If the GUID of the retention policy applied to the inactive mailbox started with the `skp` prefix, it would indicate that the retention policy is applied to Skype for Business conversations. To identity the Microsoft 365 retention policy that's applied to the inactive mailbox, run the following command in Security & Compliance Center PowerShell. `Get-RetentionCompliancePolicy <retention policy GUID without prefix> | FL Name` Be sure to remove the `mbx` or `skp` prefix when you run this command.

INACTIVE MAILBOX	HOLD TYPE	HOW TO IDENTIFY THE HOLD ON THE INACTIVE MAILBOX

Abraham McMahon	eDiscovery case hold in the Security & Compliance Center	The *InPlaceHolds* property contains the GUID of the eDiscovery case hold that's placed on the inactive mailbox. You can tell this is an eDiscovery case hold because the GUID starts with the `UniH` prefix.
		You can use the `Get-CaseHoldPolicy` cmdlet in Security & Compliance Center PowerShell to get information about the eDiscovery case that the hold on the inactive mailbox is associated with. For example, you can run the command `Get-CaseHoldPolicy <hold GUID without prefix> | FL Name` to display the name of the case hold that's on the inactive mailbox. Be sure to remove the `UniH` prefix when you run this command.
		
		To identity the eDiscovery case that the hold on the inactive mailbox is associated with, run the following commands.
		
		`$CaseHold = Get-CaseHoldPolicy <hold GUID without prefix>`
		
		`Get-ComplianceCase $CaseHold.CaseId | FL Name`
		
		Note: We don't recommend using eDiscovery holds for inactive mailboxes. That's because eDiscovery cases are intended for specific, time-bound cases related to a legal issue. At some point, a legal case will probably end and the holds associated with the case will be removed and the eDiscovery case will be closed (or deleted). In fact, if a hold that's placed on an inactive mailbox is associated with an eDiscovery case, and the hold is released or the eDiscovery case is closed or deleted, the inactive mailbox will be permanently deleted.

For more information about Microsoft 365 retention policies, see Learn about retention policies and retention labels.

Step 2: Change the hold duration for an inactive mailbox

After you identify what type of hold is placed on the inactive mailbox (and whether there are multiple holds), the next step is to change the duration for the hold.

Change the duration for a Litigation Hold

Here's how to use Exchange Online PowerShell to change the hold duration for a Litigation Hold that is placed on an inactive mailbox. You can't use the EAC. Run the following command to change the hold duration. In this

example, the hold duration is changed to an unlimited period of time.

```
Set-Mailbox -InactiveMailbox -Identity <identity of inactive mailbox> -LitigationHoldDuration unlimited
```

The result is that items in the inactive mailbox are retained indefinitely or until the hold is removed or the hold duration is changed to a different value.

> **TIP**
> The best way to identify an inactive mailbox is by using its **Distinguished Name** or **Exchange GUID** value. Using one of these values helps prevent accidentally specifying the wrong mailbox.

Change the duration for an In-Place Hold

You can use the EAC or Exchange Online PowerShell to change the hold duration for an In-Place Hold.

Use the EAC to change the hold duration

1. If you know the name of the In-Place Hold that you want to change, go to the next step. Otherwise, run the following command to get the name of the In-Place Hold that is placed on the inactive mailbox. Use the In-Place Hold GUID that you obtained in Step 1.

   ```
   Get-MailboxSearch -InPlaceHoldIdentity <In-Place Hold GUID> | FL Name
   ```

2. In the EAC, go to **Compliance management** > **In-Place eDiscovery & Hold**.
3. Select the In-Place Hold you want to change, and then select **Edit** ✎.
4. On the **In-Place eDiscovery & Hold** properties page, select **In-Place Hold**.
5. Do one of the following based on the current hold duration:
   a. Select **Hold indefinitely** to hold items for an unlimited period of time.
   b. Select **Specify number of days to hold items relative to their received date** to hold items for a specific period. Type the number of days that you want to hold items for.

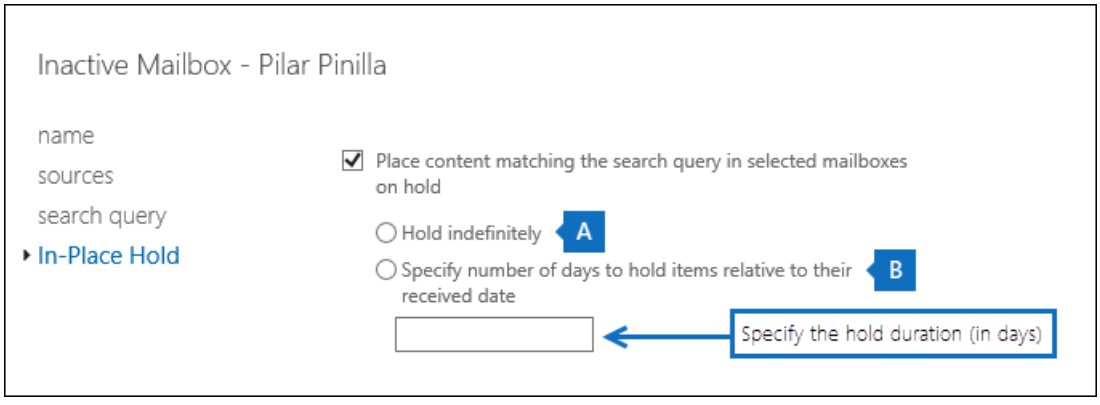

6. Select **Save**.

Use Exchange Online PowerShell to change the hold duration

1. If you know the name of the In-Place Hold that you want to change, go to the next step. Otherwise, run the following command to get the name of the In-Place Hold that is placed on the inactive mailbox. Use the In-Place Hold GUID that you obtained in Step 1.

   ```
   Get-MailboxSearch -InPlaceHoldIdentity <In-Place Hold GUID> | FL Name
   ```

2. Run the following command to change the hold duration. In this example, the hold duration is changed to 2,555 days (approximately seven years).

```
Set-MailboxSearch <identity of In-Place Hold> -ItemHoldPeriod 2555
```

   To change the hold duration to an unlimited period of time, use *-ItemHoldPeriod unlimited*.

More information

- **How is the hold duration calculated for an item in an inactive mailbox?** The duration is calculated from the original date a mailbox item was received or created.
- **What happens when the hold duration expires?** When the hold duration expires for a mailbox item in the Recoverable Items folder, the item is permanently deleted (purged) from the inactive mailbox. If there's no duration specified for the hold placed on the inactive mailbox, items in the Recoverable Items folder are never purged (unless the hold duration for the inactive mailbox is changed).
- **Is an Exchange retention policy still processed on inactive mailboxes?** If an Exchange retention policy (the messaging records management, or MRM, feature in Exchange Online) was applied to a mailbox when it was made inactive, the deletion policies (which are retention tags configured with a **Delete** retention action) will continue to be processed on the inactive mailbox. That means items that are tagged with a deletion policy are moved to the Recoverable Items folder when the retention period expires. Those items are then purged from the inactive mailbox when the hold duration for an item expires.

  Conversely, any archive policies (which are retention tags configured with a **MoveToArchive** retention action) that are included in the retention policy assigned to an inactive mailbox are ignored. That means items in an inactive mailbox that are tagged with an archive policy remain in the primary mailbox when the retention period expires. They're not moved to the archive mailbox or to the Recoverable Items folder in the archive mailbox. Because a user can't sign in to an inactive mailbox, there's no reason to consume datacenter resources to process archive policies.
- **To check the new hold duration, run one of the following commands.** The first command is for Litigation Hold; the second is for In-Place Hold.

```
Get-Mailbox -InactiveMailboxOnly -Identity <identity of inactive mailbox> | FL LitigationHoldDuration
```

```
Get-MailboxSearch <identity of In-Place Hold> | FL ItemHoldPeriod
```

- **Like regular mailboxes, the Managed Folder Assistant (MFA) also processes inactive mailboxes.** In Exchange Online, the MFA processes mailboxes approximately once every seven days. After you change the hold duration for an inactive mailbox, you can use the **Start-ManagedFolderAssistant** cmdlet to immediately start processing the new hold duration for the inactive mailbox. Run the following command.

```
Start-ManagedFolderAssistant -InactiveMailbox <identity of inactive mailbox>
```

- **If many holds are placed on an inactive mailbox, not all of the hold GUIDs will be displayed.** You can run the following command to display the GUIDs for all holds (except Litigation Holds) that are placed on an inactive mailbox.

```
Get-Mailbox -InactiveMailboxOnly -Identity <identity of inactive mailbox> | Select-Object -
ExpandProperty InPlaceHolds
```

Recover an inactive mailbox

2/18/2021 • 7 minutes to read • Edit Online

An inactive mailbox (which is a type of soft-deleted mailbox) is used to preserve a former employee's email after he or she leaves your organization. If that employee returns to your organization or if another employee takes on the job responsibilities of the former employee, there are two ways that you can make the contents of the inactive mailbox available to a user:

- **Recover an inactive mailbox.** If the former employee returns to your organization, or if a new employee is hired to take on the job responsibilities of the former employee, you can recover the contents of the inactive mailbox. This method converts the inactive mailbox to a new, active mailbox that contains the contents of the inactive mailbox. After it's recovered, the inactive mailbox no longer exists. The procedures in this topic describe this method.
- **Restore an inactive mailbox.** If another employee takes on the job responsibilities of the former employee, or if another user needs access to the contents of the inactive mailbox, you can restore (or merge) the contents of the inactive mailbox to an existing mailbox. You can also restore the archive from an inactive mailbox. For the procedures for this method, see Restore an inactive mailbox in Office 365.

See the More information section for more details about the differences between recovering and restoring an inactive mailbox, and for a description of what happens when an inactive mailbox is recovered.

> **NOTE**
> You can't recover or restore an inactive mailbox that's configured with an auto-expanding archive. If you need to recover data from an inactive mailbox with an auto-expanding archive, use content search to export the data from the mailbox and then import to another mailbox. For instructions, see following topics:
>
> - Content search
> - Export content search results

Requirements to recover an inactive mailbox

- You have to use Exchange Online PowerShell to recover an inactive mailbox. You can't use the Exchange admin center (EAC). For step-by-step instructions, see Connect to Exchange Online PowerShell.
- Run the following command to get identity information for the inactive mailboxes in your organization.

  ```
  Get-Mailbox -InactiveMailboxOnly | Format-List Name,DistinguishedName,ExchangeGuid,PrimarySmtpAddress
  ```

  Use the information returned by this command to recover a specific inactive mailbox.

Recover inactive mailboxes

Use the **New-Mailbox** cmdlet with the *InactiveMailbox* parameter to recover an inactive mailbox.

1. Create a variable that contains the properties of the inactive mailbox.

   ```
   $InactiveMailbox = Get-Mailbox -InactiveMailboxOnly -Identity <identity of inactive mailbox>
   ```

> **IMPORTANT**
> In the previous command, use the value of the **DistinguishedName** or **ExchangeGUID** property to identify the inactive mailbox. These properties are unique for each mailbox in your organization, whereas it's possible that an active and an inactive mailbox might have the same primary SMTP address.

2. This example uses the properties obtained in the previous command and recovers the inactive mailbox to an active mailbox for the user Ann Beebe. Be sure that the values specified for the *Name* and *MicrosoftOnlineServicesID* parameters are unique within your organization.

```
New-Mailbox -InactiveMailbox $InactiveMailbox.DistinguishedName -Name annbeebe -FirstName Ann -
LastName Beebe -DisplayName "Ann Beebe" -MicrosoftOnlineServicesID Ann.Beebe@contoso.com -Password
(ConvertTo-SecureString -String 'P@ssw0rd' -AsPlainText -Force) -ResetPasswordOnNextLogon $true
```

   The primary SMTP address for the recovered inactive mailbox will have the same value as the one specified by the *MicrosoftOnlineServicesID* parameter.

After you recover an inactive mailbox, a new user account is also created. You need to activate this user account by assigning a license. To assign a license in the Microsoft 365 admin center, see Add users and assign licenses at the same time.

More information

- **What's the main difference between recovering and restoring an inactive mailbox?** When you recover an inactive mailbox, the mailbox is converted to a new mailbox, the contents and folder structure of the inactive mailbox are retained, and the mailbox is linked to a new user account. After it's recovered, the inactive mailbox no longer exists, and any changes made to the content in the new mailbox will affect the content that was originally on hold in the inactive mailbox. Conversely, when you restore an inactive mailbox, the contents are merely copied to another mailbox. The inactive mailbox is preserved and remains an inactive mailbox. Any changes made to the content in the target mailbox won't affect the original content held in the inactive mailbox. The inactive mailbox can still be searched by using In-Place eDiscovery, its contents can be restored to another mailbox, or it can be recovered or deleted at a later date.
- **What happens when you recover an inactive mailbox?** When you recover an inactive mailbox, the following things occur:
  - The hold that was applied to an inactive mailbox is changed or removed based on the type of hold that was applied to the inactive mailbox before it was recovered.
    - **Litigation Hold.** If Litigation Hold was enabled for the inactive mailbox, it's removed from the recovered mailbox.
    - **In-Place Hold** In-Place Holds are removed from the recovered mailbox. This means the recovered mailbox is removed as a source mailbox from any In-Place Hold or In-Place eDiscovery search.
    - **Microsoft 365 retention policy with Preservation Lock.** If the inactive mailbox was assigned to a retention policy with Preservation Lock (called a *locked retention policy*), the recovered mailbox is assigned to the same locked retention policy. For more information about locked retention policies, see [Use Preservation Lock to restrict changes to retention policies and retention label policies.
    - **Microsoft 365 retention policy without Preservation Lock.** The inactive mailbox is removed from any unlocked Microsoft 365 retention policy that was applied to it. However,

Litigation Hold is enabled on the recovered mailbox to prevent the deletion of mailbox content based on any organization-wide retention policies that delete content older than a specific age. You can keep the Litigation Hold or remove it. For more information, see Create a Litigation Hold.

- The single item recovery period (which is defined by the **RetainDeletedItemsFor** mailbox property) is set to 30 days. Typically, when a new mailbox is created in Exchange Online, this retention period is set to 14 days. Setting this to the maximum value of 30 days gives you more time to recover any data that's been permanently deleted (or purged) from the inactive mailbox. You can also disable single item recovery or set the single item recovery period back to the default of 14 days. For more information, see Enable or disable single item recovery for a mailbox.
- Retention hold is enabled, and the retention hold duration is set to 30 days. This means that the default Exchange retention policy and any organization-wide or Exchange-wide Microsoft 365 retention policies that are assigned to the new mailbox won't be processed for 30 days. This gives the returning employee or the new owner of the recovered inactive mailbox time to manage the old messages. Otherwise, the Exchange or Microsoft 365 retention policy might delete old mailbox items (or move items to the archive mailbox, if it's enabled) that have expired based on the settings configured for the Exchange or Microsoft 365 retention policies. After 30 days, the retention hold expires, the **RetentionHoldEnabled** mailbox property is set to **False**, and the Managed Folder Assistant starts processing the policies assigned to the mailbox. If you don't need this additional time, you can just remove the retention hold. Alternatively, you can increase the duration of the retention hold by using the **Set-Mailbox -EndDateForRetentionHold** command. For more information, see Place a mailbox on retention hold.

- **Put a hold on the recovered mailbox if you need to preserve the original state of the inactive mailbox.** To prevent the new mailbox owner or retention policy from permanently deleting any messages from the recovered inactive mailbox, you can place the mailbox on Litigation Hold. For more information, see Create a Litigation Hold.
- **What user ID can you use when recovering an inactive mailbox?** When you recover an inactive mailbox, the value that you specify for the *MicrosoftOnlineServicesID* parameter can be different from the original one that was associated with the inactive mailbox. You can also use the original user ID. But as previously stated, make sure that the values used for *Name* and *MicrosoftOnlineServicesID* are unique within your organization when you recover the inactive mailbox.
- **What if the mailbox retention period for the inactive mailbox hasn't expired?** If an inactive mailbox was soft-deleted less than 30 days ago, you can't use the **New-Mailbox -InactiveMailbox** command to recover it. You need to recover it by restoring the corresponding user account. For more information, see Delete a user from your organization.
- **How do you know if the soft-deleted mailbox retention period for an inactive mailbox has expired?** Run the following command.

  ```
  Get-Mailbox -InactiveMailboxOnly <identity of inactive mailbox> | Format-List
  ExternalDirectoryObjectId
  ```

  If there isn't a value for the **ExternalDirectoryObjectId** property, the mailbox retention period has expired, and you can recover the inactive mailbox by running the **New-Mailbox -InactiveMailbox** command. If there is a value for the **ExternalDirectoryObjectId** property, the soft-deleted mailbox retention period hasn't expired and you have to recover the mailbox by restoring the user account. See Delete a user from your organization.
- **Consider enabling the archive mailbox after you recover an inactive mailbox.** This lets the returning user or new employee move old messages to the archive mailbox. And when the retention hold

expires, the archive policy that is part of the default Exchange retention policy assigned to Exchange Online mailboxes will move items that are two years or older to the archive mailbox. If you don't enable the archive mailbox, items older than two years will remain in the user's primary mailbox. For more information, see Enable archive mailboxes.

Restore an inactive mailbox

11/2/2020 • 7 minutes to read • Edit Online

An inactive mailbox (which is a type of soft-deleted mailbox) is used to retain a former employee's email after he or she leaves your organization. If another employee takes on the job responsibilities of the departed employee or if that employee returns to your organization, there are two ways that you can make the contents of the inactive mailbox available to a user:

- **Restore an inactive mailbox** If another employee takes on the job responsibilities of the departed employee, or if another user needs access to the contents of the inactive mailbox, you can restore (or merge) the contents of the inactive mailbox to an existing mailbox. You can also restore the archive from an inactive mailbox. After it's restored, the inactive mailbox is preserved and is retained as an inactive mailbox. This topic describes the procedures for restoring an inactive mailbox.
- **Recover an inactive mailbox** If the departed employee returns to your organization, or if a new employee is hired to take on the job responsibilities of the departed employee, you can recover the contents of the inactive mailbox. This method converts the inactive mailbox to a new mailbox that contains the contents of the inactive mailbox. After it's recovered, the inactive mailbox no longer exists. For the step-by-step procedures, see Recover an inactive mailbox in Office 365.

See the More information section in this article for more details about the differences between restoring and recovering an inactive mailbox.

> **NOTE**
>
> You can't recover or restore an inactive mailbox that's configured with an auto-expanding archive. If you need to recover data from an inactive mailbox with an auto-expanding archive, use content search to export the data from the mailbox and then import to another mailbox. For instructions, see following topics:
>
> - Content search
> - Export content search results

Requirements to restore an inactive mailbox

- You have to use Exchange Online PowerShell to restore an inactive mailbox. You can't use the Exchange admin center (EAC). For step-by-step instructions, see Connect to Exchange Online PowerShell.
- Run the following command in Exchange Online PowerShell to get identity information for the inactive mailboxes in your organization.

  ```
  Get-Mailbox -InactiveMailboxOnly | Format-List Name,DistinguishedName,ExchangeGuid,PrimarySmtpAddress
  ```

  Use the information returned by this command to restore a specific inactive mailbox.
- For more information about inactive mailboxes, see Inactive mailboxes in Office 365.

Restore inactive mailboxes

Use the **New-MailboxRestoreRequest** cmdlet with the *SourceMailbox* and *TargetMailbox* parameters to restore the contents of an inactive mailbox to an existing mailbox. For more information about using this cmdlet, see New-MailboxRestoreRequest.

1. Create a variable that contains the properties of the inactive mailbox.

```
$InactiveMailbox = Get-Mailbox -InactiveMailboxOnly -Identity <identity of inactive mailbox>
```

> **IMPORTANT**
> In the previous command, use the value of the **DistinguishedName** or **ExchangeGUID** property to identify the inactive mailbox. These properties are unique for each mailbox in your organization, whereas it's possible that an active and an inactive mailbox might have the same primary SMTP address.

2. Restore the contents of the inactive mailbox to an existing mailbox. The contents of the inactive mailbox (source mailbox) will be merged into the corresponding folders in the existing mailbox (target mailbox).

```
New-MailboxRestoreRequest -SourceMailbox $InactiveMailbox.DistinguishedName -TargetMailbox
newemployee@contoso.com -AllowLegacyDNMismatch
```

   Alternatively, you can specify a top-level folder in the target mailbox in which to restore the contents from the inactive mailbox. If the specified target folder or target folder structure doesn't already exist in the target mailbox, it is created during the restore process.

   This example copies mailbox items and subfolders from an inactive mailbox to a folder named "Inactive Mailbox" in the top-level folder structure of the target mailbox.

```
New-MailboxRestoreRequest -SourceMailbox $InactiveMailbox.DistinguishedName -TargetMailbox
newemployee@contoso.com -TargetRootFolder "Inactive Mailbox" -AllowLegacyDNMismatch
```

Restore the archive from an inactive mailbox

If an inactive mailbox has an archive mailbox, you can also restore it to the archive mailbox of an existing mailbox. To restore the archive from an inactive mailbox, you have to add the *SourceIsArchive* and *TargetIsArchive* switches to the command used to restore an inactive mailbox.

1. Create a variable that contains the properties of the inactive mailbox.

```
$InactiveMailbox = Get-Mailbox -InactiveMailboxOnly -Identity <identity of inactive mailbox>
```

> **NOTE**
> In the previous command, use the value of the **DistinguishedName** or **ExchangeGUID** property to identify the inactive mailbox. These properties are unique for each mailbox in your organization, whereas it's possible that an active and an inactive mailbox might have the same primary SMTP address.

2. Restore the contents of the archive from the inactive mailbox (source archive) to the archive of an existing mailbox (target archive). In this example, the contents from the source archive are copied to a folder named "Inactive Mailbox Archive" in the archive of the target mailbox.

```
New-MailboxRestoreRequest -SourceMailbox $InactiveMailbox.DistinguishedName -SourceIsArchive -
TargetMailbox newemployee@contoso.com -TargetIsArchive -TargetRootFolder "Inactive Mailbox Archive" -
AllowLegacyDNMismatch
```

More information

- **What's the main difference between recovering and restoring an inactive mailbox?** When you recover an inactive mailbox, the mailbox is basically converted to a new mailbox, the contents and folder structure of the inactive mailbox are retained, and the mailbox is linked to a new user account. After it's recovered, the inactive mailbox no longer exists, and any changes made to the content in the new mailbox will affect the content that was originally on hold in the inactive mailbox. Conversely, when you restore an inactive mailbox, the contents are merely copied to another mailbox. The inactive mailbox is preserved and remains an inactive mailbox. Any changes made to the content in the target mailbox won't affect the original content held in the inactive mailbox. The inactive mailbox can still be searched by using the Content Search tool, its contents can be restored to another mailbox, or it can be recovered or deleted at a later date.

- **How do you find inactive mailboxes?** To get a list of the inactive mailboxes in your organization and display information that is useful for restoring an inactive mailbox, you can run this command.

  ```
  Get-Mailbox -InactiveMailboxOnly | Format-List
  Name,PrimarySMTPAddress,DistinguishedName,ExchangeGUID,LegacyExchangeDN,ArchiveStatus
  ```

- **Use a Litigation Hold or Microsoft 365 retention policy to retain inactive mailbox content.** If you want to retain the state of an inactive mailbox after it's restored, you can place the target mailbox on Litigation Hold or apply an Microsoft 365 retention policy before you restore the inactive mailbox. This will prevent the permanent deletion of any items from the inactive mailbox after they're restored to the target mailbox.

- **Enable retention hold on the target mailbox before you restore an inactive mailbox.** Because mailbox items from an inactive mailbox could be old, you might consider enabling retention hold on the target mailbox before you restore an inactive mailbox. When you put a mailbox on retention hold, the retention policy that's assigned to it won't be processed until the retention hold is removed or until the retention hold period expires. This gives the owner of the target mailbox time to manage old messages from the inactive mailbox. Otherwise, the retention policy might delete old items (or move items to the archive mailbox, if it's enabled) that have expired based on the retention settings configured for the target mailbox. For more information, see Place a mailbox on retention hold in Exchange Online.

- **What does the AllowLegacyDNMismatch switch do?** In the previous examples to restore an inactive mailbox, the **AllowLegacyDNMismatch** switch is used to allow restoring the inactive mailbox to a different target mailbox. In a typical restore scenario, the goal is to restore content where the source and target mailboxes are the same mailbox. So by default, the **New-MailboxRestoreRequest** cmdlet checks to make sure that the value of the **LegacyExchangeDN** property on the source and target mailboxes is the same. This helps prevents you from accidentally restoring a source mailbox into the wrong target mailbox. If you try to restore an inactive mailbox without using the **AllowLegacyDNMismatch** switch, the command might fail if the source and target mailboxes have different values for the **LegacyExchangeDN** property.

- **You can use other parameters with the New-MailboxRestoreRequest cmdlet to implement different restore scenarios for inactive mailboxes.** For example, you can run this command to restore the archive from the inactive mailbox into the primary mailbox of the target mailbox.

  ```
  New-MailboxRestoreRequest -SourceMailbox <inactive mailbox> -SourceIsArchive -TargetMailbox <target
  mailbox> -TargetRootFolder "Inactive Mailbox Archive" -AllowLegacyDNMismatch
  ```

  You can also restore the inactive primary mailbox into the archive of the target mailbox by running this command.

```
New-MailboxRestoreRequest -SourceMailbox <inactive mailbox> -TargetMailbox <target mailbox> -
TargetIsArchive -TargetRootFolder "Inactive Mailbox" -AllowLegacyDNMismatch
```

- **What does the TargetRootFolder parameter do?** As previously explained, you can use the **TargetRootFolder** parameter to specify a folder in the top of the folder structure (also called the root) in the target mailbox in which to restore the contents of the inactive mailbox. If you don't use this parameter, mailbox items from the inactive mailbox are merged into the corresponding default folders of the target mailbox, and custom folders are re-created in the root of the target mailbox. The following illustrations highlight these differences between not using and using the **TargetRootFolder** parameter.

Folder hierarchy in the target mailbox when the TargetRootFolder parameter isn't used

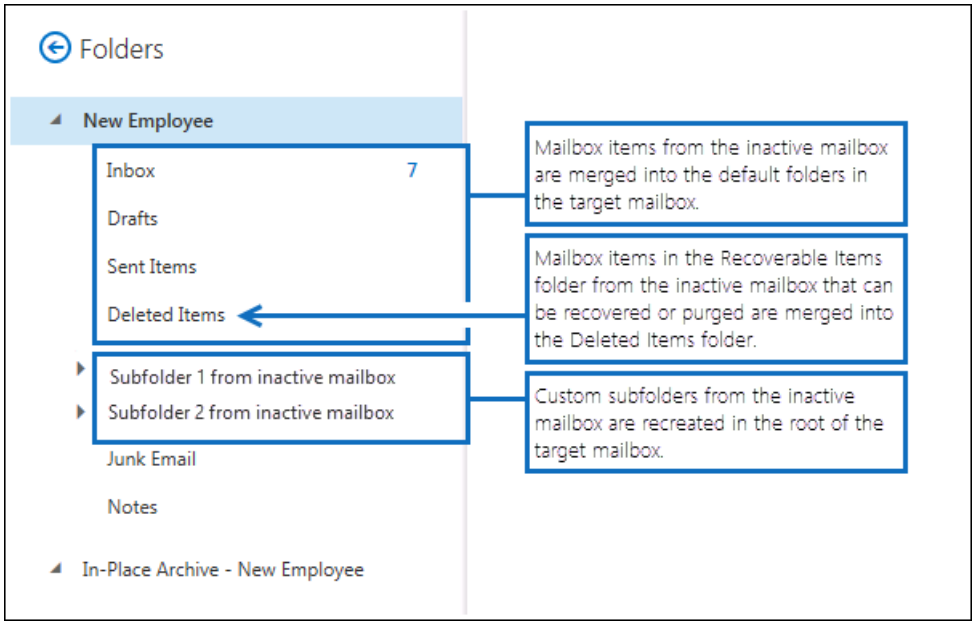

Folder hierarchy in the target mailbox when the TargetRootFolder parameter is used

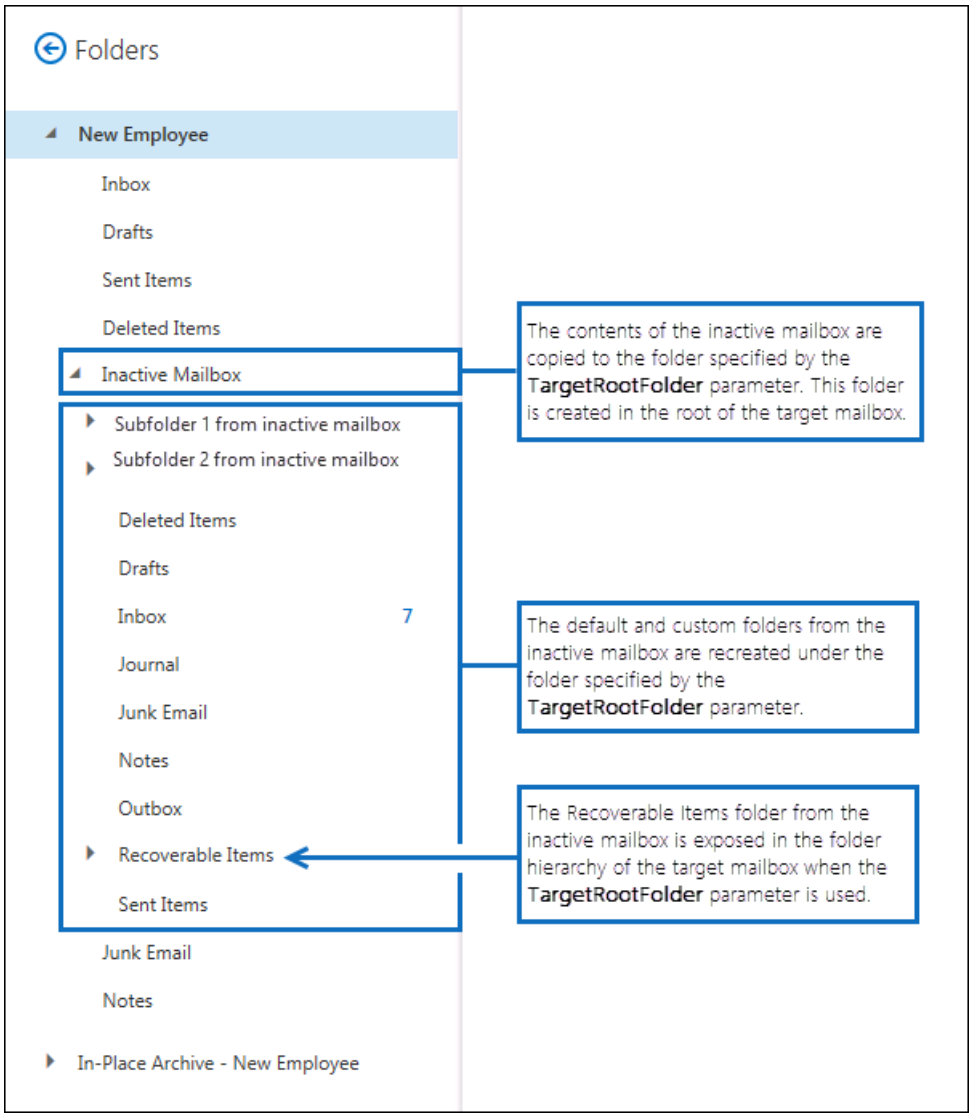
Folders
New Employee
Inbox
Drafts
Sent Items
Deleted Items
Inactive Mailbox
Subfolder 1 from inactive mailbox
Subfolder 2 from inactive mailbox
Deleted Items
Drafts
Inbox 7
Journal
Junk Email
Notes
Outbox
Recoverable Items
Sent Items
Junk Email
Notes
In-Place Archive - New Employee
The contents of the inactive mailbox are copied to the folder specified by the TargetRootFolder parameter. This folder is created in the root of the target mailbox.
The default and custom folders from the inactive mailbox are recreated under the folder specified by the TargetRootFolder parameter.
The Recoverable Items folder from the inactive mailbox is exposed in the folder hierarchy of the target mailbox when the TargetRootFolder parameter is used.

Delete an inactive mailbox

11/2/2020 • 10 minutes to read • Edit Online

An inactive mailbox is used to preserve a former employee's email after he or she leaves your organization. When you no longer need to preserve the contents of an inactive mailbox, you can permanently delete the inactive mailbox by removing the hold. Also, it's possible that multiple holds might be placed on an inactive mailbox. For example, an inactive mailbox might be placed on Litigation Hold and on one or more In-Place Holds. Additionally, a retention policy (created in the security and compliance center in Office 365 or Microsoft 365) might be applied to the inactive mailbox. You have to remove all holds and retention policies from an inactive mailbox to delete it. After you remove the holds and retention policies, the inactive mailbox is marked for deletion and is permanently deleted after it's processed.

IMPORTANT

As we continue to invest in different ways to preserve mailbox content, we're announcing the retirement of In-Place Holds in the Exchange admin center. That means you should use Litigation Holds and retention policies to create an inactive mailbox. Starting July 1, 2020 you won't be able to create new In-Place Holds in Exchange Online. But you'll still be able to change the hold duration of an In-Place Hold placed on an inactive mailbox. However, starting October 1, 2020, you won't be able to change the hold duration. You'll only be able to delete an inactive mailbox by removing the In-Place Hold. Existing inactive mailboxes that are on In-Place Hold will still be preserved until the hold is removed. For more information about the retirement of In-Place Holds, see Retirement of legacy eDiscovery tools.

See the More information section for a description of what happens after holds are removed from an inactive mailbox.

Before you delete an inactive mailbox

- You have to use Exchange Online PowerShell to remove a Litigation Hold from an inactive mailbox. You can't use the Exchange admin center (EAC). For step-by-step instructions, see Connect to Exchange Online PowerShell. You can use Exchange Online PowerShell or the EAC to remove an In-Place Hold from an inactive mailbox.
- You can copy the contents of an inactive mailbox to another mailbox before you remove the hold and delete an inactive mailbox. For details, see Restore an inactive mailbox in Office 365.
- If you remove the hold or retention policy from an inactive mailbox and the soft-deleted mailbox retention period for the mailbox has expired, the mailbox will be permanently deleted. After it's deleted, it can't be recovered. Before you remove the hold, be sure that you no longer need the contents in the mailbox. If you want to re-activate an inactive mailbox, you can recover it. For details, see Recover an inactive mailbox in Office 365.
- For more information about inactive mailboxes, see Inactive mailboxes in Office 365.

Step 1: Identify the holds on an inactive mailbox

As previously stated, a Litigation Hold, In-Place Hold, or retention policy might be placed on an inactive mailbox. The first step is to identify the holds on an inactive mailbox.

Run the following command to display the hold information for all inactive mailboxes in your organization.

```
Get-Mailbox -InactiveMailboxOnly | FL DisplayName,Name,IsInactiveMailbox,LitigationHoldEnabled,InPlaceHolds
```

The value of **True** for the **LitigationHoldEnabled** property indicates that the inactive mailbox is on Litigation Hold. If an In-Place Hold is placed on an inactive mailbox, the GUID for the hold is displayed as the value for the **InPlaceHolds** property. For example, the following results for two inactive mailboxes show that a Litigation Hold is placed on Ann Beebe and that two In-Place Holds are placed on Pilar Pinilla.

```
DisplayName           : Ann Beebe
Name                  : annb
IsInactiveMailbox     : True
LitigationHoldEnabled : True
InPlaceHolds          : {}
...
DisplayName           : Pilar Pinilla
Name                  : pilarp
IsInactiveMailbox     : True
LitigationHoldEnabled : False
InPlaceHolds          : {c0ba3ce811b6432a8751430937152491, ba6f4ba25b62490aaaa253eea27426ab}
```

TIP

If a lot of In-Place Holds are placed on an inactive mailbox, not all of the In-Place Hold GUIDs will be displayed. You can run the following command to display all the In-Place Hold GUIDs:

```
Get-Mailbox -InactiveMailboxOnly -Identity <identity of inactive mailbox> | Select-Object -ExpandProperty InPlaceHolds
```

Step 2: Remove a hold from an inactive mailbox

After you identify what type of hold is placed on the inactive mailbox (and whether there are multiple holds), the next step is to remove the holds on the mailbox. As previously stated, you have to remove all holds to permanently delete an inactive mailbox.

Remove a Litigation Hold

As previously stated, you have to use Windows PowerShell to remove a Litigation Hold from an inactive mailbox. You can't use the EAC. Run the following command to remove a Litigation Hold.

```
Set-Mailbox -InactiveMailbox -Identity <identity of inactive mailbox> -LitigationHoldEnabled $false
```

TIP

The best way to identify an inactive mailbox is by using its Distinguished Name or Exchange GUID value. Using one of these values helps prevent accidentally specifying the wrong mailbox.

Remove In-Place Holds

There are two ways to remove an In-Place Hold from an inactive mailbox:

- **Delete the In-Place Hold object** If the inactive mailbox that you want to permanently delete is the only source mailbox for an In-Place Hold, you can just delete the In-Place Hold object.

> **NOTE**
> You have to disable the hold before you can delete an In-Place Hold object. If you try to delete an In-Place Hold object that has the hold enabled, you'll receive an error message.

- **Remove the inactive mailbox as a source mailbox of an In-Place Hold** If you want to retain other source mailboxes for an In-Place Hold, you can remove the inactive mailbox from the list of source mailboxes and keep the In-Place Hold object.

Use the EAC to delete an In-Place Hold

1. If you know the name of the In-Place Hold that you want to delete, you can go to the next step. Otherwise, run the following command to get the name of the In-Place Hold that is placed on the inactive mailbox that you want to permanently delete. Use the In-Place Hold GUID that you obtained in Step 1: Identify the holds on an inactive mailbox.

   ```
   Get-MailboxSearch -InPlaceHoldIdentity <In-Place Hold GUID> | FL Name
   ```

2. In the EAC, go to **Compliance management** > **In-Place eDiscovery & Hold**.
3. Select the In-Place Hold you want to delete, and then click **Edit** .
4. On the **In-Place eDiscovery & Hold** properties page, click **In-Place Hold**, uncheck the **Place content matching the search query in selected mailboxes on hold** box, and then click **Save**.
5. On the **In-Place eDiscovery & Hold** page, select the In-Place Hold again, and then click **Delete** .
6. On the warning, click **Yes** to delete the In-Place Hold.

Use Exchange Online PowerShell to delete an In-Place Hold

1. Create a variable that contains the properties of the In-Place Hold that you want to delete. Use the In-Place Hold GUID that you obtained in Step 1: Identify the holds on an inactive mailbox.

   ```
   $InPlaceHold = Get-MailboxSearch -InPlaceHoldIdentity <In-Place Hold GUID>
   ```

2. Disable the hold on the In-Place Hold.

   ```
   Set-MailboxSearch $InPlaceHold.Name -InPlaceHoldEnabled $false
   ```

3. Delete the In-Place Hold.

   ```
   Remove-MailboxSearch $InPlaceHold.Name
   ```

Use the EAC to remove an inactive mailbox from an In-Place Hold

1. If you know the name of the In-Place Hold that's placed on the inactive mailbox, you can go to the next step. Otherwise, run the following command to get the name of the In-Place Hold placed on the mailbox. Use the In-Place Hold GUID that you obtained in Step 1: Identify the holds on an inactive mailbox.

   ```
   Get-MailboxSearch -InPlaceHoldIdentity <In-Place Hold GUID> | FL Name
   ```

2. In the EAC, go to **Compliance management** > **In-Place eDiscovery & Hold**.
3. Select the In-Place Hold that is placed on the inactive mailbox, and then click **Edit** .

4. On the **In-Place eDiscovery & Hold** properties page, click **Sources**.
5. In the list of source mailboxes, click the name of the inactive mailbox that you want to remove, and then click **Remove**—.
6. Click **Save** to save the change. A message is displayed saying the operation was successfully completed.
7. Repeat steps 1 through 6 to remove other In-Place Holds placed on the inactive mailbox.

Use Exchange Online PowerShell to remove an inactive mailbox from an In-Place Hold

If the In-Place Hold contains a large number of source mailboxes, it's possible the inactive mailbox won't be listed on the **Sources** page in the EAC. Up to 3,000 mailboxes are displayed on the **Sources** page when you edit an In-Place Hold. If an inactive mailbox isn't listed on the **Sources** page, you can use Exchange Online PowerShell to remove it from the In-Place Hold.

1. Create a variable that contains the properties of the In-Place Hold placed on the inactive mailbox. Use the In-Place Hold GUID that you obtained in Step 1: Identify the holds on an inactive mailbox.

   ```
   $InPlaceHold = Get-MailboxSearch -InPlaceHoldIdentity <In-Place Hold GUID>
   ```

2. Verify that the inactive mailbox is listed as a source mailbox for the In-Place Hold.

   ```
   $InPlaceHold.Sources
   ```

   Note: The *Sources* property of the In-Place Hold identifies the source mailboxes by their *LegacyExchangeDN* properties. Because this property uniquely identifies inactive mailboxes, using the *Sources* property from the In-Place Hold helps prevent removing the wrong mailbox. This also helps to avoid issues if two mailboxes have the same alias or SMTP address.

3. Remove the inactive mailbox from the list of source mailboxes in the variable. Be sure to use the **LegacyExchangeDN** of the inactive mailbox that's returned by the command in the previous step.

   ```
   $InPlaceHold.Sources.Remove("<LegacyExchangeDN of the inactive mailbox>")
   ```

   For example, the following command removes the inactive mailbox for Pilar Pinilla.

   ```
   $InPlaceHold.Sources.Remove("/o=contoso/ou=Exchange Administrative Group
   (FYDIBOHF23SPDLT)/cn=Recipients/ cn=9c8dfff651ec4908950f5df60cbbda06-pilarp")
   ```

4. Verify that the inactive mailbox is removed from the list of source mailboxes in the variable.

   ```
   $InPlaceHold.Sources
   ```

5. Modify the In-Place Hold with the updated list of source mailboxes, which doesn't include the inactive mailbox.

   ```
   Set-MailboxSearch $InPlaceHold.Name -SourceMailboxes $InPlaceHold.Sources
   ```

6. Verify that the inactive mailbox is removed from the list of source mailboxes for the In-Place Hold.

   ```
   Get-MailboxSearch $InPlaceHold.Name | FL Sources
   ```

More information

- **An inactive mailbox is a type of soft-deleted mailbox.** In Exchange Online, a soft-deleted mailbox is a mailbox that's been deleted but can be recovered within a specific retention period. The soft-deleted mailbox retention period in Exchange Online is 30 days. This means that the mailbox can be recovered within 30 days of being soft-deleted. After 30 days, a soft-deleted mailbox is marked for permanent deletion and can't be recovered.

- **What happens after you remove the hold on an inactive mailbox?** The mailbox is treated like other soft-deleted mailboxes and is marked for permanent deletion after the 30-day soft-deleted mailbox retention period expires. This retention period starts on the date when the mailbox was first made inactive. This date is known as the soft-deleted date, which is the date the corresponding user account was deleted or when the Exchange Online mailbox was deleted with the **Remove-Mailbox** cmdlet. The soft-deleted date isn't the date on which you remove the hold.

- **Is an inactive mailbox permanently deleted immediately after the hold is removed?** If the soft-deleted date for an inactive mailbox is older than 30 days, the mailbox won't be permanently deleted as soon as you remove the hold. The mailbox will be marked for permanent deletion and is deleted the next time it's processed.

- **How does the soft-deleted mailbox retention period affect inactive mailboxes?** If the soft-deleted date for an inactive mailbox is more than 30 days before the date the hold was removed, the mailbox is marked for permanent deletion. But if an inactive mailbox has a soft-deleted date within the last 30 days and you remove the hold, you can recover the mailbox up until the soft-deleted mailbox retention period expires. For details, see Delete or restore user mailboxes in Exchange Online. After the soft-deleted mailbox retention period expires, you have follow the procedures for recovering an inactive mailbox. For details, see Recover an inactive mailbox in Office 365.

- **How do you display information about an inactive mailbox after the hold is removed?** After a hold is removed and the inactive mailbox is reverted back to a soft-deleted mailbox, it won't be returned by using the *InactiveMailboxOnly* parameter with the **Get-Mailbox** cmdlet. But you can display information about the mailbox by using the **Get-Mailbox -SoftDeletedMailbox** command. For example:

  ```
  Get-Mailbox -SoftDeletedMailbox -Identity pilarp | FL Name,Identity,LitigationHoldEnabled,In
  Placeholds,WhenSoftDeleted,IsInactiveMailbox
  Name                  : pilarp
  Identity              : Soft Deleted Objects\pilarp
  LitigationHoldEnabled : False
  InPlaceHolds          : {}
  WhenSoftDeleted       : 10/30/2014 1:19:04 AM
  IsInactiveMailbox     : False
  ```

  In the above example, the *WhenSoftDeleted* property identifies the soft-deleted date, which in this example is October 30, 2014. If this soft-deleted mailbox was previously an inactive mailbox for which the hold was removed, it will be permanently deleted 30 days after the value of the *WhenSoftDeleted* property. In this case, the mailbox is permanently deleted after November 30, 2014.

Learn about records management in Microsoft 365

2/18/2021 • 4 minutes to read • Edit Online

> *Microsoft 365 licensing guidance for security & compliance*.

Organizations of all types require a records-management solution to manage regulatory, legal, and business-critical records across their corporate data. Records management in Microsoft 365 helps an organization manage their legal obligations, provides the ability to demonstrate compliance with regulations, and increases efficiency with regular disposition of items that are no longer required to be retained, no longer of value, or no longer required for business purposes.

Use the following capabilities to support your records management solution in Microsoft 365:

- **Label content as a record**. Create and configure retention labels to mark content as a record that can then be applied by users or automatically applied by identifying sensitive information, keywords, or content types.
- **Migrate and manage your retention requirements with file plan**. By using a file plan, you can bring in an existing retention plan to Microsoft 365, or build a new one for enhanced management capabilities.
- **Configure retention and deletion settings with retention labels**. Configure retention labels with the retention periods and actions based on various factors that include the date last modified or created.
- **Start different retention periods when an event occurs** with event-based retention.
- **Review and validate disposition** with disposition reviews and proof of records deletion.
- **Export information about all disposed items** with the export option.
- **Set specific permissions** for records manager functions in your organization to have the right access.

Using these capabilities, you can incorporate your organization's retention schedules and requirements into a records management solution that manages retention, records declaration, and disposition, to support the full lifecycle of your content.

In addition to the online documentation, you might find it useful to listen to the webinar recording for records management, and download the accompanying deck with FAQs.

Records

When content is declared a record:

- Restrictions are placed on the items in terms of what actions are allowed or blocked.
- Additional activities about the item are logged.
- You have proof of disposition when the items are deleted at the end of their retention period.

You use retention labels to mark content as a **record**, or a **regulatory record**. The difference between these two are explained in the next section. You can either publish those labels so that users and administrators can manually apply them to content, or auto-apply those labels to content that you want to mark as a record or a regulatory record.

By using retention labels to declare records, you can implement a single and consistent strategy for managing

records across your Microsoft 365 environment.

Compare restrictions for what actions are allowed or blocked

Use the following table to identify what restrictions are placed on content as a result of applying a standard retention label, and retention labels that mark content as a record or regulatory record.

A standard retention label has retention settings and actions but doesn't mark content as a record or a regulatory record.

> **NOTE**
>
> For completeness, the table includes columns for a locked and unlocked record, which is applicable to SharePoint and OneDrive, but not Exchange. The ability to lock and unlock a record uses record versioning that isn't supported for Exchange items. So for all Exchange items that are marked as a record, the behavior maps to the **Record - locked** column, and the **Record - unlocked column** is not relevant.

ACTION	RETENTION LABEL	RECORD - LOCKED	RECORD - UNLOCKED	REGULATORY RECORD
Edit contents	Allowed	**Blocked**	Allowed	**Blocked**
Edit properties, including rename	Allowed	Allowed	Allowed	**Blocked**
Delete	Allowed [1]	**Blocked**	**Blocked**	**Blocked**
Copy	Allowed	Allowed	Allowed	Allowed
Move within container [2]	Allowed	Allowed	Allowed	Allowed
Move across containers [2]	Allowed	Allowed if never unlocked	**Blocked**	**Blocked**
Open/Read	Allowed	Allowed	Allowed	Allowed
Change label	Allowed	Allowed - container admin only	Allowed - container admin only	**Blocked**
Remove label	Allowed	Allowed - container admin only	Allowed - container admin only	**Blocked**

Footnotes:

[1] Supported by OneDrive and Exchange by retaining a copy in a secured location, but blocked by SharePoint.

Message a user sees if they try to delete a labeled document in SharePoint:

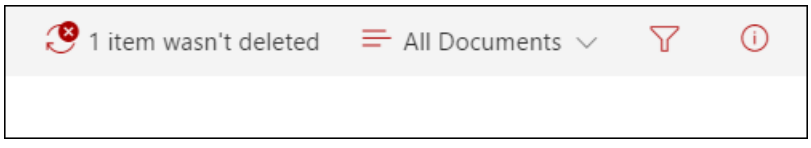

[2] Containers include SharePoint document libraries and Exchange mailboxes.

IMPORTANT

The most important difference for a regulatory record is that after it is applied to content, nobody, not even a global administrator, can remove the label.

Retention labels configured for regulatory records also have the following admin restrictions:

- The retention period can't be made shorter after the label is saved, only extended.
- These labels aren't supported by auto-labeling policies, and must be applied by using retention label policies.

In addition, a regulatory label can't be applied to a document that's checked out in SharePoint.

Because of the restrictions and irreversible actions, make sure you really do need to use regulatory records before you select this option for your retention labels. To help prevent accidental configuration, this option is not available by default but must first be enabled by using PowerShell. Instructions are included in Declare records by using retention labels.

Configuration guidance

See Get started with records management.

To mark content as a record, see Declare records by using retention labels.

Get started with records management

2/18/2021 • 4 minutes to read • Edit Online

> *Microsoft 365 licensing guidance for security & compliance*.

Ready to start managing your organization's high-value content for legal, business, or regulatory obligations by using a records management solution in Microsoft 365? Use the following high-level guidance to get started:

1. **Understand the records management solution** and what actions are allowed or blocked when documents and emails are declared records: Learn about records management.
2. **Understand retention labels and how retention works** for SharePoint and Exchange, because retention labels are used to declare records: Learn about retention policies and retention labels
3. **Create your file plan for retention settings and actions** by importing an existing plan if you have one, or create new retention labels that declare records.
4. **Publish and apply your retention labels**. Retention labels are reusable building blocks that can be used in multiple policies and can be incorporated into user workflows:
   - Create retention labels and apply them in apps
   - Apply a retention label to content automatically

Subscription and licensing requirements for records management

A number of different subscriptions support records management and the licensing requirements for users depend on the features you use.

To see the options for licensing your users to benefit from Microsoft 365 compliance features, see the Microsoft 365 licensing guidance for security & compliance. For records management, see the Records Management section and related PDF or Excel download for feature-level licensing requirements.

Permissions required for records management

Members of your compliance team who are responsible for records management need permissions to the Microsoft 365 compliance center. By default, the tenant admin (global administrator) has access to this location and can give compliance officers and other people access without giving them all the permissions of a tenant admin. To grant permissions for this limited administration, we recommend that you add users to the **Records Management** admin role group, which grants permissions for all features related to records management, including disposition review and verification.

For a read-only role, you can create a new role group and add the **View-Only Record Management** role to this group.

For more information about role groups and roles, see Permissions in the Security & Compliance Center.

For instructions to add users to role groups and assign roles, see Give users access to the Security & Compliance Center.

These permissions are required only to create, configure, and apply retention labels that declare records, and manage disposition. The person configuring these labels doesn't require access to the content.

Common scenarios for records management

Use the following table to help you map your business requirements to the scenarios that are supported by records management.

> **NOTE**
>
> Because records management uses retention labels to mark an item as a record, many scenarios in this table are also listed as common scenarios for retention policies and retention labels.

I WANT TO ...	DOCUMENTATION
Declare a record	Declare records by using retention labels
Update a record	Use record versioning to update records stored in SharePoint or OneDrive
Let admins and users manually apply retain and delete actions for documents and emails: - SharePoint - OneDrive - Outlook and Outlook on the web	Create retention labels and apply them in apps
Let site admins set default retain and delete actions for all content in a SharePoint library, folder, or document set	Create retention labels and apply them in apps
Let users automatically apply retain and delete actions to emails by using Outlook rules	Create retention labels and apply them in apps
Let admins apply retain and delete actions to a document understanding model, so that these are automatically applied to identified documents in a SharePoint library	Create retention labels and apply them in apps
Automatically apply retain and delete actions to documents and emails	Apply a retention label to content automatically
Start the retention period when an event occurs, such as: - Employees leave the organization - Contracts expire - End of product lifetime	Start retention when an event occurs
Restrict changes to policies to help meet regulatory requirements or safeguard against rogue administrators	Use Preservation Lock to restrict changes to retention policies and retention label policies
Manage the lifecycle of different document types in SharePoint	Use retention labels to manage the lifecycle of documents stored in SharePoint
Make sure somebody reviews and approves before content is deleted at the end of its retention period	Disposition reviews
Have proof of disposition for content that is permanently deleted at the end of its retention period	Disposition of records
Monitor how and where retain and delete settings are applied to items	Monitoring retention labels

End-user documentation for records

Retention labels that are used for records management have a UI presence in Microsoft 365 apps. Make sure you provide guidance for end users and your help desk before you deploy retention labels to your production network.

The most effective end-user documentation will be customized guidance and instructions you provide for the retention label names and configurations you choose. See the following post for a download package that you can use to train users and drive adoption: End User Training for Retention Labels in M365 – How to Accelerate Your Adoption.

You will also find basic user instructions in the following section: Manually apply retention labels.

Declare records by using retention labels

11/2/2020 • 2 minutes to read • Edit Online

> *Microsoft 365 licensing guidance for security & compliance.*

To declare documents and emails as records, you use retention labels that mark the content as a **record** or a **regulatory record**.

If you're not sure whether to use a record or a regulatory record, see Compare restrictions for what actions are allowed or blocked. If you need to use regulatory records, you must first run a PowerShell command, as described in the next section.

You can then either publish those labels in a retention label policy so that users and administrators can apply them to content, or for labels that mark items as records (but not regulatory records), auto-apply those labels to content that you want to declare a record.

How to display the option to mark content as a regulatory record

> **NOTE**
> The following procedure is an auditable action, logging **Enabled regulatory record option for retention labels** in the Retention policy and retention label activities section of the audit log.

By default, the retention label option to mark content as a regulatory record isn't displayed in the retention label wizard. To display this option, you must first run a PowerShell command:

1. Connect to the Office 365 Security & Compliance Center Powershell.
2. Run the following cmdlet:

   ```
   Set-RegulatoryComplianceUI -Enabled $true
   ```

   There is no prompt to confirm and the setting takes effect immediately.

If you change your mind about seeing this option in the retention label wizard, you can hide it again by running the same cmdlet with the **false** value: `Set-RegulatoryComplianceUI -Enabled $false`

Configuring retention labels to declare records

When you create or edit a retention label from the **Records Management** solution in the Microsoft 365 compliance center, you have the option to mark items as a record. If you ran the PowerShell command from the previous section, you can alternatively mark items as a regulatory record.

For example:

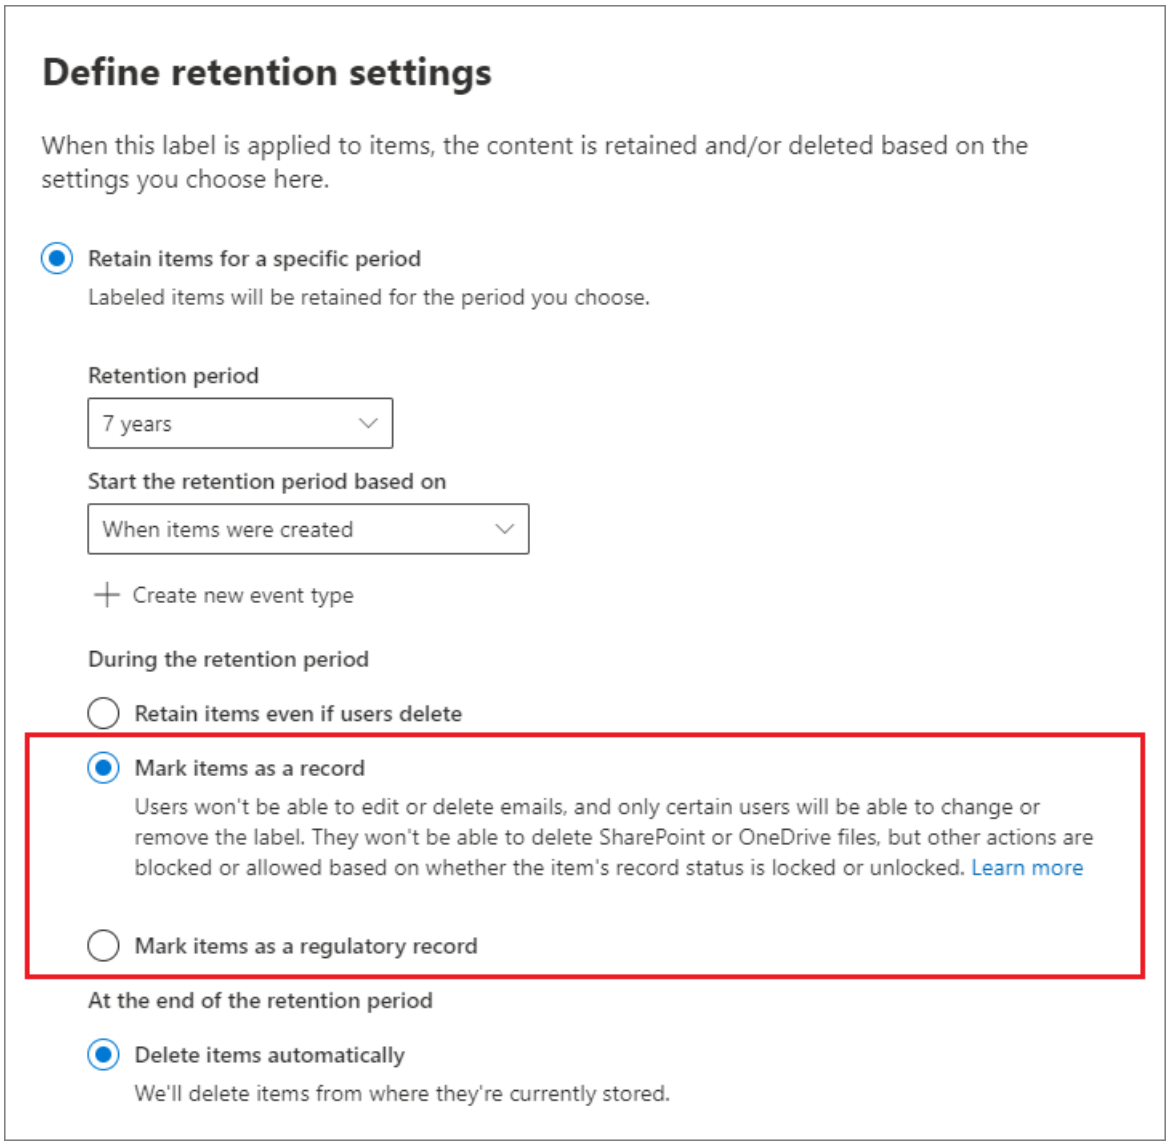

Using this retention label, you can now apply it to SharePoint or OneDrive documents and Exchange emails, as needed.

For full instructions:

- Create retention labels and apply them in apps
- Apply a retention label to content automatically (not supported for regulatory records)

Applying the configured retention label to content

When retention labels that mark items as a record or regulatory record are made available for users to apply them in apps:

- For Exchange, any user with write-access to the mailbox can apply these labels.
- For SharePoint and OneDrive, any user in the default Members group (the Contribute permission level) can apply these labels.

Example of a document marked as record by using a retention label:

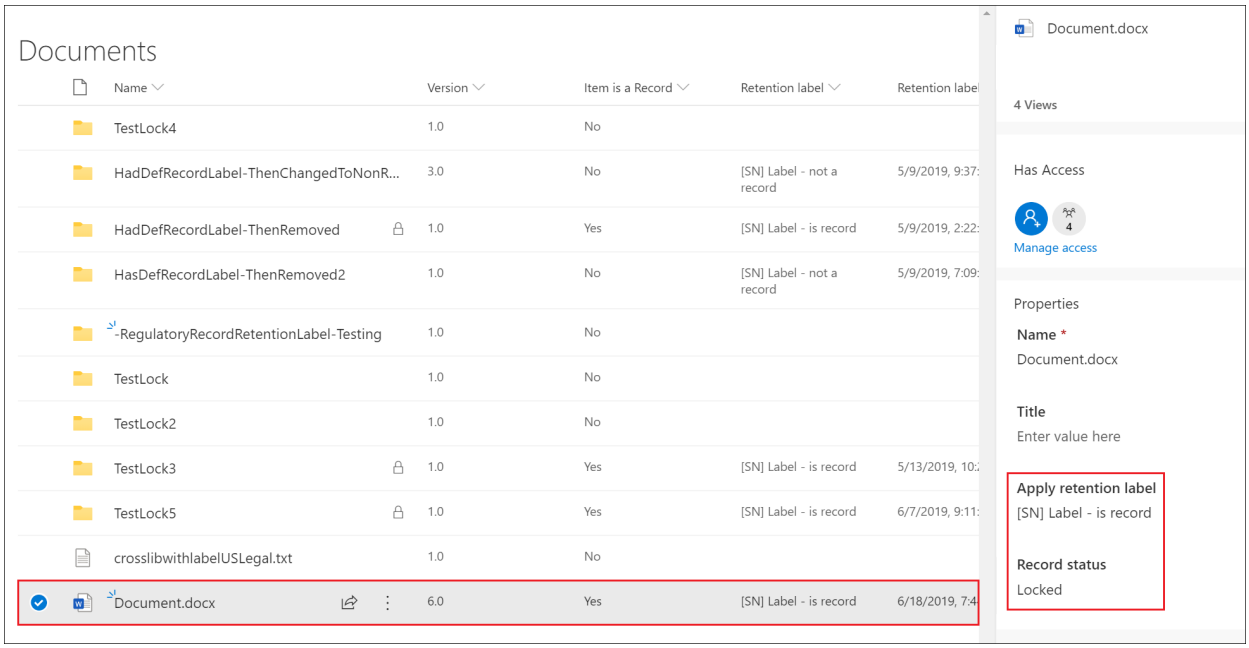

Next steps

For a list of scenarios supported by records management, see Common scenarios for records management.

Use record versioning to update records stored in SharePoint or OneDrive

11/2/2020 • 4 minutes to read • Edit Online

Microsoft 365 licensing guidance for security & compliance.

NOTE

Because regulatory records block editing, record versioning is not available for regulatory records.

The ability to mark a document as a record and restrict actions that can be performed on the record is an essential goal for any records management solution. However, collaboration might also be needed for people to create subsequent versions.

For example, you might mark a sales contract as a record, but then need to update the contract with new terms and mark the latest version as a new record while still retaining the previous record version. For these types of scenarios, SharePoint and OneDrive support *record versioning*. OneNote notebook folders don't support record versioning.

To use record versioning, you first label the document and mark it as a record. At this point, a document property, called *Record status* is displayed next to the retention label, and the initial record status is **Locked**.

You can now do the following things:

- **Continually edit and retain individual versions of the document as records, by unlocking and locking the Record status property.** Only when the **Record status** property is set to **Locked** is a new version of the record retained. This toggle of locked and unlocked reduces the risk of retaining unnecessary versions and copies of the document.
- **Have the records automatically stored in an in-place records repository located within the site collection.** Each site collection in SharePoint and OneDrive preserves content in its Preservation Hold library. Record versions are stored in the Records folder in this library.
- **Maintain an evergreen document that contains all versions.** By default, each SharePoint and OneDrive document has a version history available on the item menu. In this version history, you can easily see which versions are records and view those documents.

Record versioning is automatically available for any document that has a retention label that marks the item as a record. When a user views the document properties by using the details pane, they can toggle the **Record status** from **Locked** to **Unlocked**. This action creates a record in the Records folder in the Preservation Hold library, where it resides for the remainder of its retention period.

While the document is unlocked, any user with standard edit permissions can edit the file. However, users can't delete the file, because it's still a record. When editing is complete, a user can then toggle the **Record status** from **Unlocked** to **Locked**, which prevents further edits while in this status.

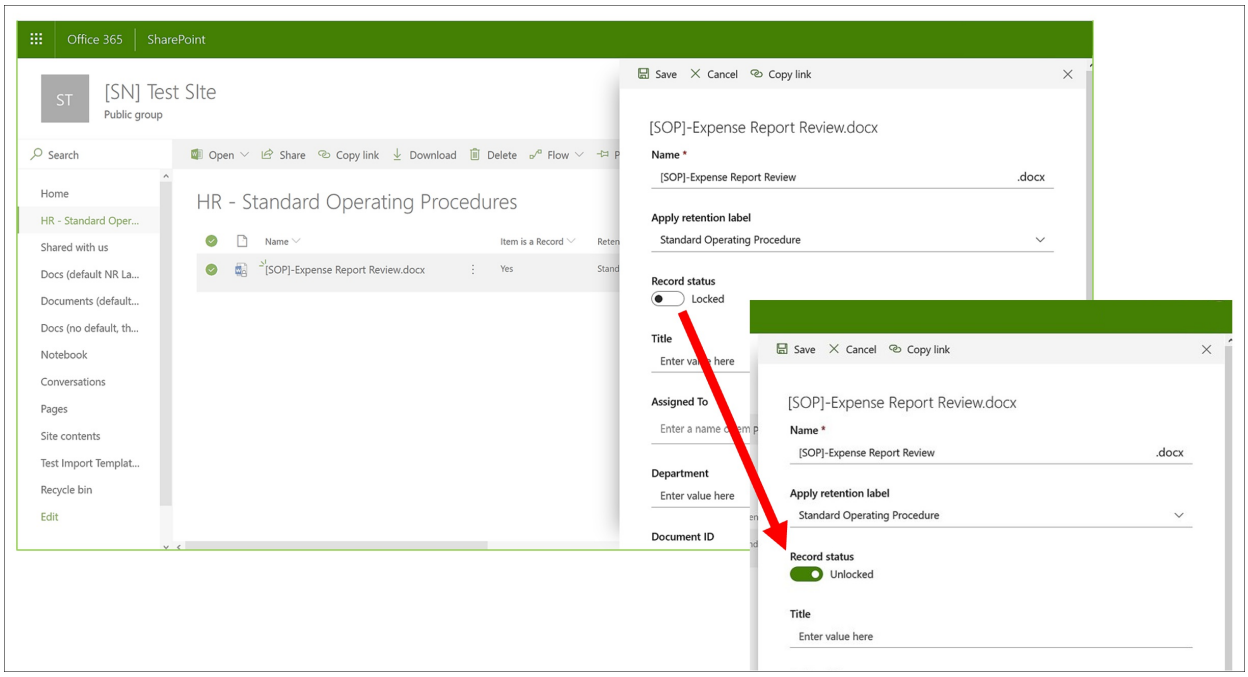

Locking and unlocking a record

After a retention label that marks content as a record is applied to a document, any user with Contribute permissions or a narrower permission level can unlock a record or lock an unlocked record.

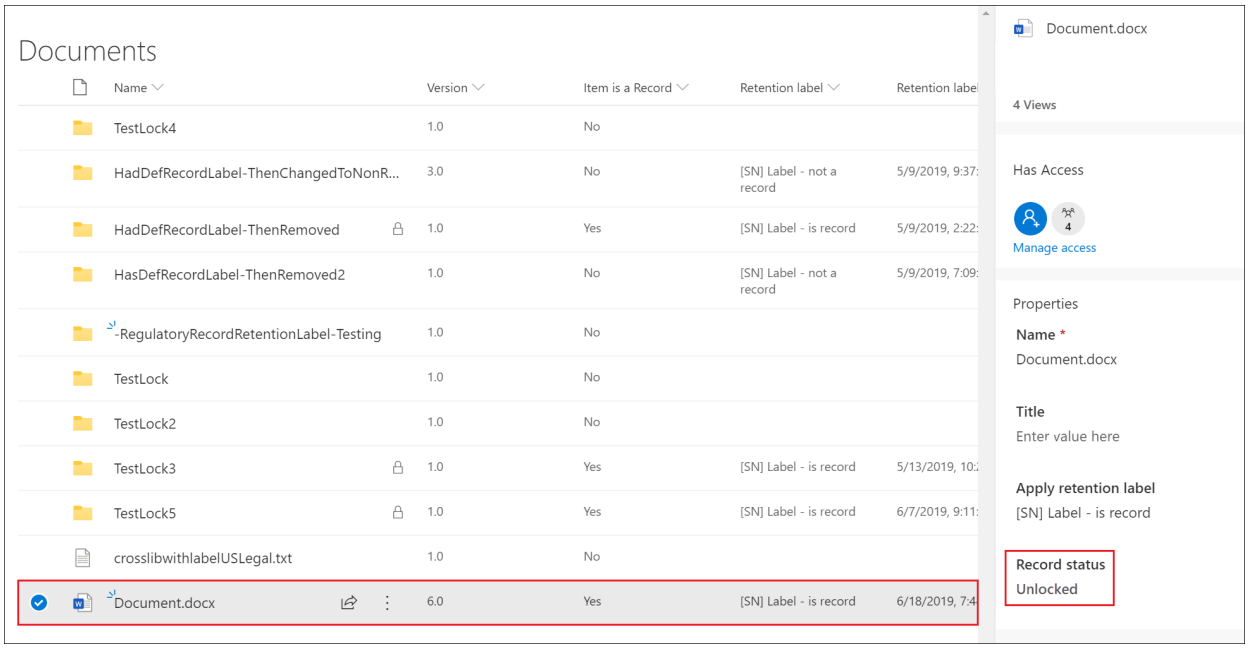

When a user unlocks a record, the following actions occur:

1. If the current site collection doesn't have a Preservation Hold library, one is created.
2. If the Preservation Hold library doesn't have a Records folder, one is created.
3. A **Copy to** action copies the latest version of the document to the Records folder. The **Copy to** action includes only the latest version and no prior versions. This copied document is now considered a record version of the document, and its file name has the format: [Title GUID Version#]
4. The copy created in the Records folder is added to the version history of the original document, and this version shows the word **Record** in the comments field.
5. The original document is a new version that can be edited, but not deleted. The document library column

Item is a Record still shows the **Yes** value because the document is still a record, even if it can now be edited.

When a user locks a record, the original document again can't be edited. But it is the action of unlocking a record that copies a version to the Records folder in the Preservation Hold library.

Record versions

Each time a user unlocks a record, the latest version is copied to the Records folder in the Preservation Hold library, and that version contains the value of **Record** in the **Comments** field of the version history.

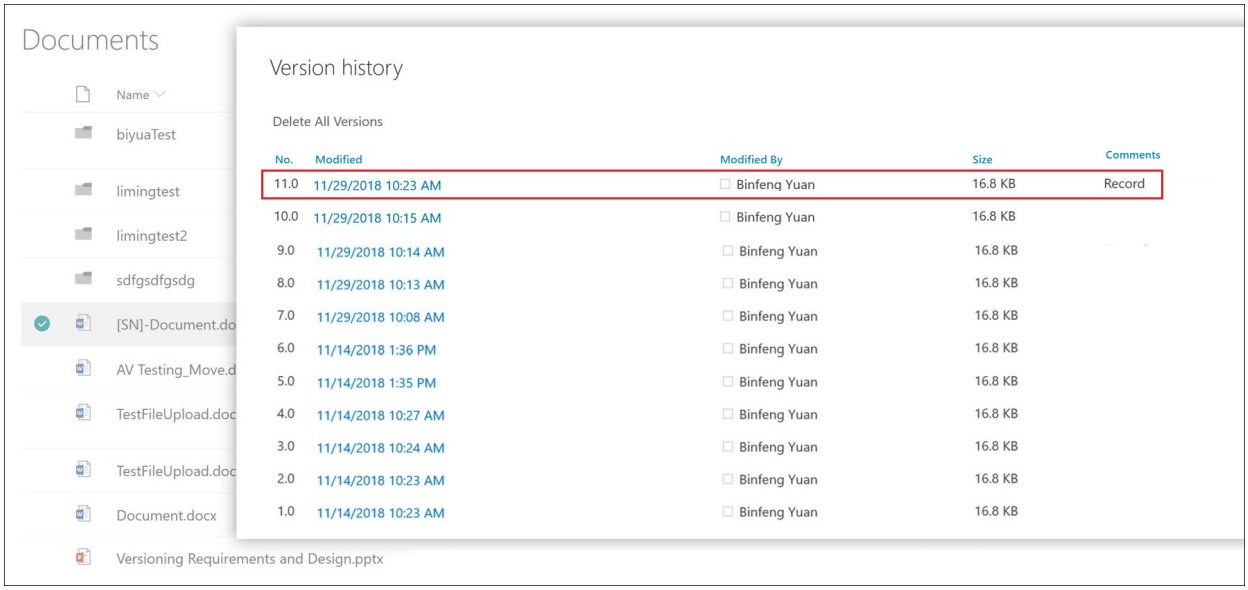

To view the version history, select a document in the document library and then click **Version history** in the item menu.

Where records are stored

Records are stored in the Records folder in the Preservation Hold library in the top-level site in the site collection. In the left navigation on the top-level site, choose **Site contents** > **Preservation Hold Library**.

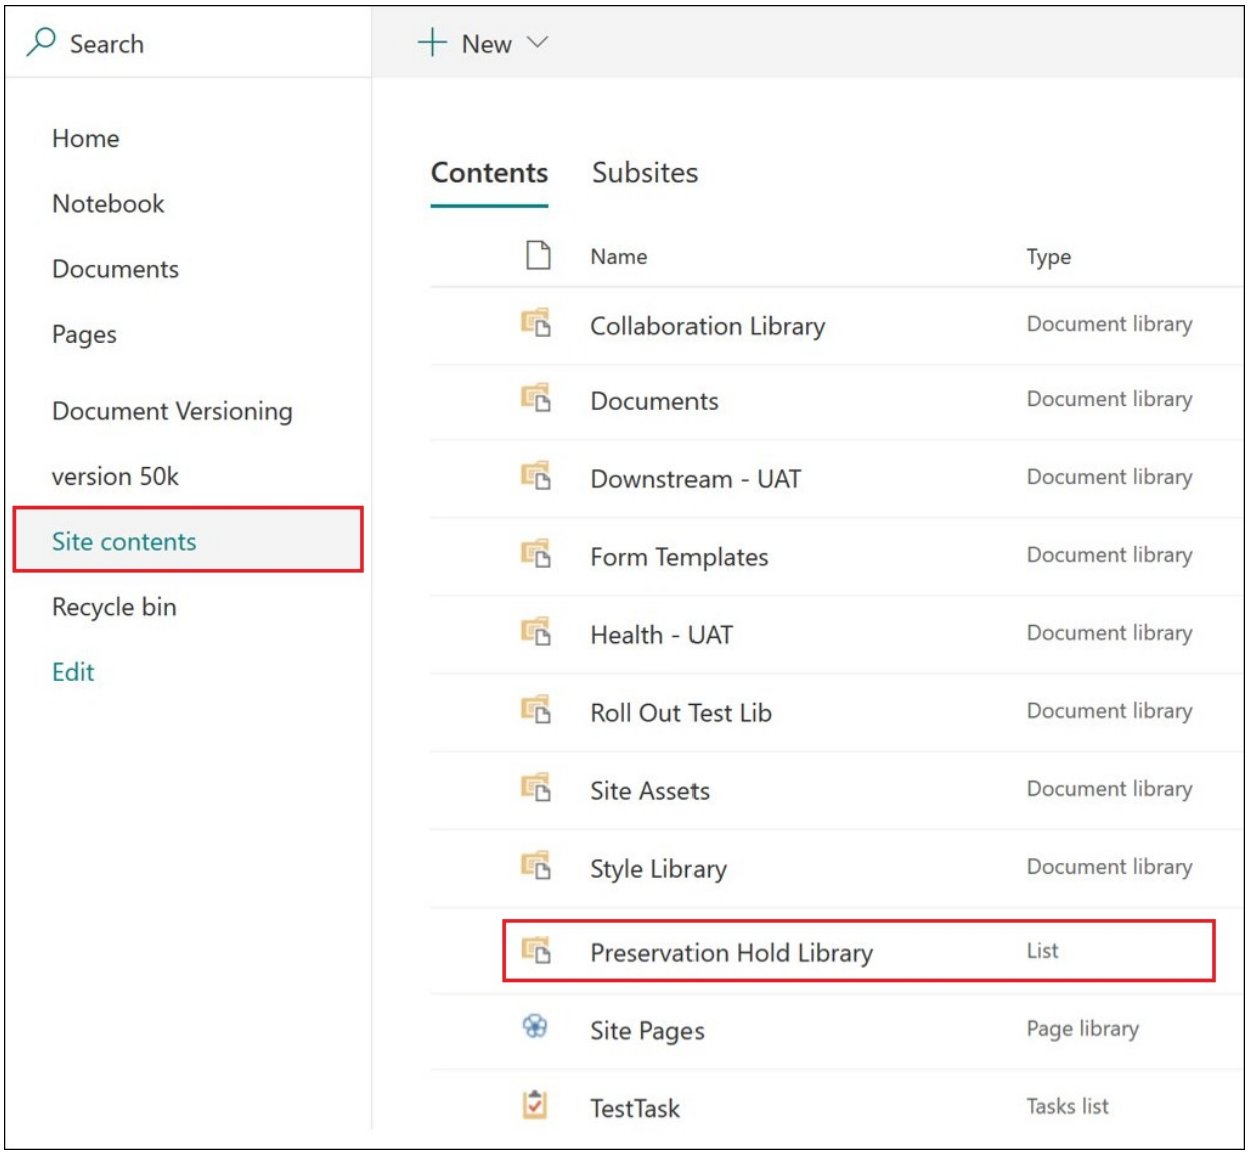

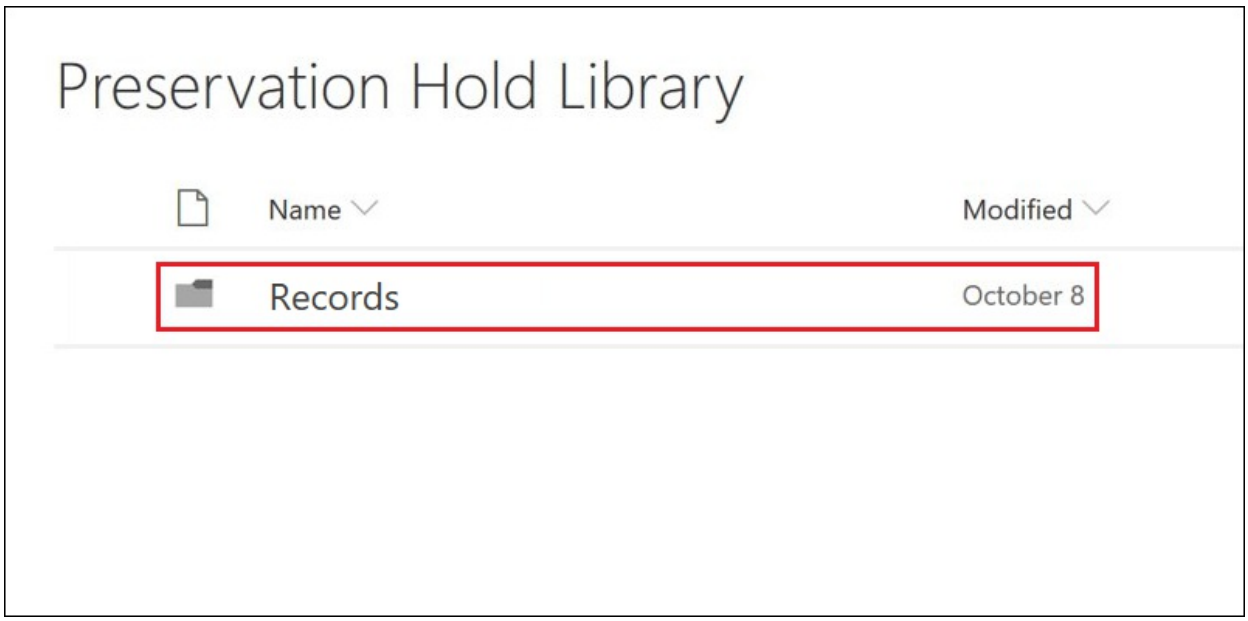

The Preservation Hold library is visible only to site collection admins. Also, the Preservation Hold library doesn't exist by default. It's created only when content subject to a retention label or retention policy is deleted for the first time in the site collection.

Searching the audit log for record versioning events

The actions of locking and unlocking records are logged in the audit log. You can search for the specific activities **Changed record status to locked** and **Changed record status to unlocked**, which are located in the **File and page activities** section in the **Activities** dropdown list on the **Audit log search** page in the security and compliance center.

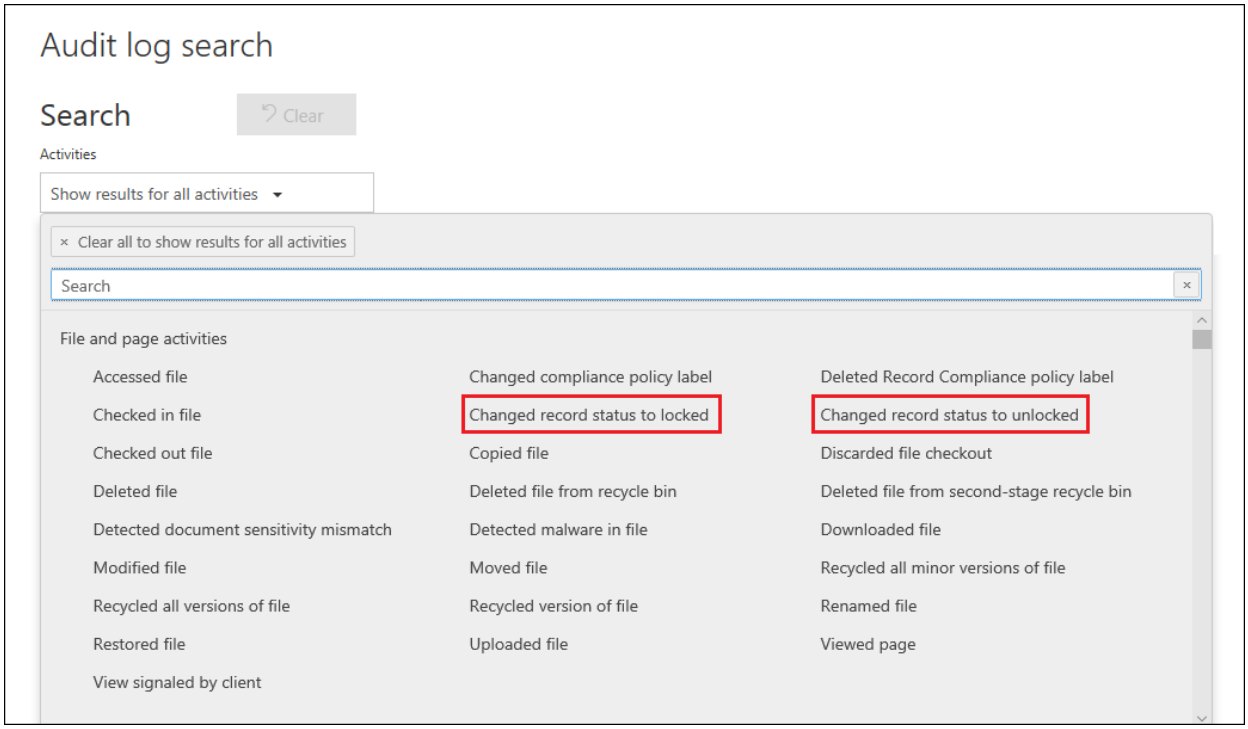

For more information about searching for these events, see the "File and page activities" section in Search the audit log in the Security & Compliance Center.

Next steps

For other scenarios supported by records management, see Common scenarios for records management.

Regulatory requirements for information governance and records management

2/18/2021 • 2 minutes to read • Edit Online

> *Microsoft 365 licensing guidance for security & compliance*.

Use the resources on this page to help you meet specific regulatory requirements for information governance and records management in Microsoft 365. Each section of this document focuses on one or more related regulations and includes any existing guidance or third-party assessment of how to configure Microsoft 365 to help with the requirements outlined.

These resources are available to download from the Data Protection Resources, FAQ and White Papers page of the Service Trust Portal.

New Zealand Public Records Act

Supporting New Zealand's Public Records Act compliance obligations with Microsoft 365 - Download assessment

Applicable workloads: SharePoint, OneDrive, Teams, and Exchange

Released January 2021, this report has been produced in partnership with Microsoft New Zealand to assess the capabilities of Microsoft 365 services for recording, storing, and managing requirements for electronic records, as specified by:

- New Zealand Public Records Act 2005, which sets guidelines for preservation of public archives and local authority archives in New Zealand.

This report helps you understand how the system aspects of the New Zealand Public Records Act 2005 (PRA) are achievable when using Microsoft 365.

SEC 17a-4(f), FINRA 4511(c), and CFTC 1.31(c)-(d)

Cohasset Assessment - Microsoft 365 - SEC Rule 17a-4(f) - Immutable Storage for SharePoint, OneDrive, Teams, Exchange, and Skype - Download assessment

Applicable workloads: SharePoint, OneDrive, Teams, Exchange, and Skype for Business

Released November 2020, this report has been produced in partnership with Cohasset Associates, Inc. (Cohasset) to assess the capabilities of Microsoft 365 services for recording, storing, and managing requirements for electronic records, as specified by:

- Securities and Exchange Commission (SEC) in 17CFR §240.17a-4(f), which regulates exchange members, brokers or dealers.
- Financial Industry Regulatory Authority (FINRA) Rule 4511(c), which defers to the format and media requirements of SEC Rule 17a-4(f).
- The principles-based electronic records requirements of the Commodity Futures Trading Commission (CFTC) in 17CFR §1.31(c)-(d).

The opinion from Cohasset is that when compliance features are properly configured and carefully applied and

managed as described in their report, the assessed Microsoft 365 services meet the five requirements related to the recording and non-rewriteable, non-erasable storage of electronic records.

eDiscovery solutions in Microsoft 365

2/18/2021 • 6 minutes to read • Edit Online

Electronic discovery, or eDiscovery, is the process of identifying and delivering electronic information that can be used as evidence in legal cases. You can use eDiscovery tools in Microsoft 365 to search for content in Exchange Online mailboxes, Microsoft 365 Groups, Microsoft Teams, SharePoint Online and OneDrive for Business sites, and Skype for Business conversations, and Yammer teams. You can search mailboxes and sites in the same eDiscovery search by using the Content Search tool. And you can use Core eDiscovery cases to identify, hold, and export content found in mailboxes and sites. If your organization has an Office 365 E5 or Microsoft 365 E5 subscription (or related E5 add-on subscriptions), you can further manage custodians and analyze content by using the Advanced eDiscovery solution in Microsoft 365.

Microsoft 365 provides the following eDiscovery tools:

- Content search
- Core eDiscovery
- Advanced eDiscovery

Content search

The following table contains links to articles that will help you use the Content search tool.

ARTICLE	DESCRIPTION
Run a search	Learn how to use the Content Search tool to search mailboxes, public folders, Microsoft 365 Groups, Microsoft Teams, SharePoint Online sites, One Drive for Business locations, and Skype for Business conversations in your organization in a single search.
Keyword queries and search conditions	Learn about the email and file properties and search conditions you can use to search for content in mailboxes and sites in your organization.
View keyword statistics for search results	Learn how to use search statistics to display and compare the statistics for one or more content searches, and to configure new and existing searches to return statistics for each keyword in the search query.
Export search results	Learn how to export the results of a Content search.
Configure permissions filtering for Content search	Learn how to use permissions filtering to let an eDiscovery manager search only a subset of mailboxes and sites in your organization.
Export a search report	Learn how to download the export report without having to export the actual search results.
Content search limits	Learn about the limits of the Content Search tool, such as the maximum number of searches that you can run at one time.

ARTICLE	DESCRIPTION
Unindexed items in Content search	Learn about unindexed items in Exchange and SharePoint that you can include in the estimated search result statistics when you run a search. You can also include unindexed items when you export search results.
Search for and delete email messages	Learn how to use Content search to search for and delete an email message from *all* mailboxes in your organization. This can help you find and remove potentially harmful or high-risk email.
Search the mailbox and OneDrive accounts for a list of users	Learn how to use a script to search the mailbox and One Drive for Business site for a group of users. See Create a list of all OneDrive locations for steps on how to quickly generate a list of email addresses that you can use for the source content locations when you create and run content searches.
Use Content search for targeted collections	Learn how to use the Windows PowerShell script in this article to perform targeted collections using Content search. A targeted collection means you want to search a specific folder because you're confident that items responsive to a case (or privileged items) are located in that folder. Use the script in this article to obtain the folder ID or path for the specific mailbox or site folders that you want to search.

Core eDiscovery

The following table contains links to topics that will help you use Core eDiscovery cases. You can use Core eDiscovery cases to add eDiscovery managers who can access the case, place an eDiscovery hold on content locations relevant to the case, search for content, and export the search results from the case.

ARTICLE	DESCRIPTION
Get started with Core eDiscovery	Learn how to assign eDiscovery permissions and create Core eDiscovery cases. This topic also provides an overview of the Core eDiscovery workflow.
Assign eDiscovery permissions	Learn how to assign permissions to users so they can search for content, place content locations on hold, and perform other eDiscovery-related tasks in a Core eDiscovery case.
Set up compliance boundaries for Core eDiscovery	Learn how to use compliance boundaries to create logical boundaries within an organization that control the content locations that an eDiscovery manager can search.
Create an eDiscovery hold	Learn how to create eDiscovery holds that associated with a Core eDiscovery case to preserve content relevant to the case you're investigating.
Search for content in a case	Learn how to search for content that's relevant to a case. You can quickly create searches that search the content locations on hold.

ARTICLE	DESCRIPTION
Export content from a case	Learn how to export and download content from a Core eDiscovery case.
Close, reopen, and delete a case	Learn how to manage the lifecycle of a Core eDiscovery case.

Advanced eDiscovery

The Advanced eDiscovery solution in Microsoft 365 (also called *Advanced eDiscovery v2.0*) builds on the existing eDiscovery and analytics capabilities in Microsoft 365. This eDiscovery solution provides an end-to-end workflow to preserve, collect, review, analyze, and export content that's responsive to your organization's internal and external investigations. It also lets legal teams manage custodians and the entire legal hold notification workflow to communicate with custodians involved in a case.

ARTICLE	DESCRIPTION
Overview of Advanced eDiscovery	This article introduces Advanced eDiscovery, outlines the business justification for using this tool, presents Advanced eDiscovery architecture, and provides a high-level overview of the built-in workflow of Advanced eDiscovery.
Set up Advanced eDiscovery	Learn how to get started using Advanced eDiscovery, including the required licensing and necessary eDiscovery permission.
Create and manage a case	This article shows you how to create an Advanced eDiscovery case and provides a walk-through of the Advanced eDiscovery workflow.
Manage custodians	Learn about working with custodians in an Advanced eDiscovery. This topic links to step-by-step instructions to add custodians to a case, managing custodians in a case, and viewing custodian activity in Microsoft 365 by searching the audit log.
Manage custodian communications	Learn about managing the legal hold notification process in Advanced eDiscovery. This includes creating and automating the notification workflow and how a user acknowledged a hold notification.
Manage processing errors	Learn about Advanced indexing and how to remediate indexing errors in content from custodial and non-custodial content locations, such as Exchange mailboxes, SharePoint sites, and OneDrive accounts. You can bulk-remediate errors and then upload remediated files to a review set or remediate individual processing errors within a review set.
Collect data for a case	Learn about searching for content in custodial content locations, and then adding relevant case data to a review set. When you copy content to a review set, the data is copied from the original content locations to a Microsoft-provided Azure Storage location. This provides a static set of documents for the review process.

ARTICLE	DESCRIPTION
Manage review sets	Learn about reviewing case data in a review set. This includes viewing, querying, filtering, and tagging documents in a review set.
Analyze data in a review set	Learn about running analysis on the documents in a review set. The results of running analysis include near-duplication detection, email threading, and themes identification.
Export case data	Learn about exporting data from a case for external review.

eDiscovery roadmap

To see what eDiscovery features have been launched, are rolling out, or in development, see the Microsoft 365 Roadmap.

Search for content using the Content Search tool

2/18/2021 • 3 minutes to read • Edit Online

Use the Content Search tool in the Security & Compliance Center to quickly find email in Exchange mailboxes, documents in SharePoint sites and OneDrive locations, and instant messaging conversations in Skype for Business. You can use the content search tool to search for email, documents, and instant messaging conversations in collaboration tools such as Microsoft Teams and Microsoft 365 Groups.

Search for content

The first step is to starting using the Content Search tool to choose content locations to search and configure a keyword query to search for specific items. Or, you can just leave the query blank and return all items in the target locations.

- Create and run a content search
- Build search queries and use conditions to narrow your search
- Configure search permissions filtering so that an eDiscovery manager can only search subset of mailboxes or sites in your organization
- Run an ID list search to search for specific email messages
- Search cloud-based mailboxes for on-premises users in Microsoft 365
- View keyword statistics for the results of a search and then refine the query if necessary
- Search for third-party data that your organization has imported to Microsoft 365
- Bulk edit the query and content locations for multiple searches
- Retry a Content Search to resolve a content location error
- Preserve Bcc recipients so you can search for them

Perform actions on content you find

After you run a search and refine it as necessary, the next step is to do something with the results returned by the search. You can export and download the results to your local computer or in the case of a email attack on your organization, you can delete the results of a search from user mailboxes.

- Export the results of a content search and download them to your local computer
- Search for and delete email messages , such as messages that content a virus, dangerous attachments, or phishing messages
- Export a report about the results of a content search, without exporting the actual results

Learn more about content search

Content Search is easy to use, but it's also a powerful tool. Behind-the-scenes, there's a lot going on. The more you know about it and understand its behavior and its limitations, the more successful you'll be using it for your organization's search and investigation needs. Learn about:

- Partially indexed items in Exchange and SharePoint and how to include or exclude them when you export

and download search results

- Investigate partially indexed items and determine your organization's exposure to them
- Limits of the Content Search tool, such as the maximum number of searches that you can run at one time and the maximum number of content locations you can include in a single search
- Estimated and actual search results and the reasons why there might be differences between them when you export and download search results
- De-duplication in search results that you can enable when you export email messages that are the results of a search

Use scripts for advanced scenarios

Sometimes you have to perform more advanced, complex, and repetitive content search tasks. In these cases, it's easier and fast to use PowerShell commands in the Security & Compliance Center. To help make this easier, we've created a number of Security & Compliance Center PowerShell scripts to help you complete complex content search-related tasks.

- Search specific mailbox and site folders (called a *targeted collection) when you're confident that items responsive to a case are located in that folder
- Search the mailbox and OneDrive location for a list of users
- Create, report on, and delete multiple searches to quickly and efficiently identify and cull search data
- Clone a content search and quickly compare the results of different keyword search queries run on the same content locations; or use the script to save time by not having to re-enter a large number of content locations when you create a new search

Content Search

2/18/2021 • 25 minutes to read • Edit Online

You can use the Content search eDiscovery tool in the compliance center in Office 365 or Microsoft 365 to search for in-place items such as email, documents, and instant messaging conversations in your organization. Use this tool to search for items in these services:

- Exchange Online mailboxes
- SharePoint Online sites and OneDrive for Business accounts
- Microsoft Teams
- Microsoft 365 Groups
- Yammer Groups
- Skype for Business conversations

After you run a Content search, the number of content locations and an estimated number of search results are displayed in the search statistics. You can also quickly view statistics, such as the content locations that have the most items that match the search query. After you run a search, you can preview the results or export them to a local computer.

Create a search

To have access to the **Content search** page to run searches and preview and export search results, an administrator, compliance officer, or eDiscovery manager must be a member of the eDiscovery Manager role group in the Security & Compliance Center. For more information, see Assign eDiscovery permissions.

1. Go to https://compliance.microsoft.com and sign in using your Microsoft email address and password.
2. In the left navigation pane of the Microsoft 365 compliance center, click **Show all**, and then click **Content search**.
3. On the **Content search** page, click **New search**.

   You can also choose one of the other search options:

   - **Guided search:** This option starts a wizard that guides you through the creating the search. The user interface to select content locations and build the search query are the same as the **New search** option.
   - **Search by ID list:** This option lets you search for specific email messages and other mailbox items using a list of Exchange IDs. To create an ID list search, you submit a comma-separated value (CSV) file that identifies the specific mailbox items to search for. For instructions, see Prepare a CSV file for an ID list search.
4. Under **Search query**, specify the following things:

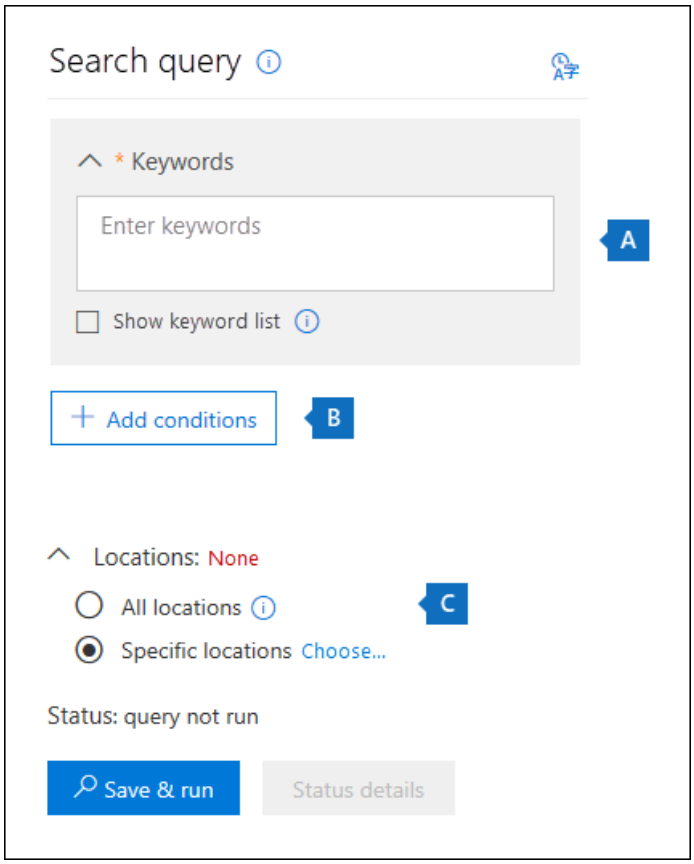

- **Keywords to search for:** Type a search query in **Keywords** box. You can specify keywords, message properties such as sent and received dates, or document properties such as file names or the date that a document was last changed. You can use more complex queries that use a Boolean operator, such as **AND**, **OR**, **NOT**, and **NEAR**. You can also search for sensitive information (such as social security numbers) in documents, or search for documents that have been shared externally. If you leave the keyword box empty, all content located in the specified content locations is included in the search results.

  Alternatively, you can click the **Show keyword list** checkbox and the type a keyword in each row. If you do this, the keywords on each row are connected by a logical operator (**c:s**) that is similar in functionality to the **OR** operator in the search query that's created.

  Why use the keyword list? You can get statistics that show how many items match each keyword. This can help you quickly identify which keywords are the most (and least) effective. You can also use a keyword phrase (surrounded by parentheses) in a row. For more information about search statistics, see View keyword statistics for Content Search results.

  > **NOTE**
  > To help reduce issues caused by large keyword lists, you're now limited to a maximum of 20 rows in the keyword list.

- **Conditions:** You can add search conditions to narrow a search and return a more refined set of results. Each condition adds a clause to the search query that is created and run when you start the search. A condition is logically connected to the keyword query (specified in the keyword box) by a logical operator (**c:c**) that is similar in functionality to the **AND** operator. That means that items have to satisfy both the keyword query and one or more conditions to be included in the results. This is how conditions help to narrow your results. For a list and description of conditions that you can use in a search query, see the "Search conditions" section in Keyword queries and search conditions for Content Search.

- **Locations:** Choose the content locations to search.
- **All locations:** Use this option to search all content locations in your organization. This includes email in all Exchange mailboxes (including all inactive mailboxes, and mailboxes for all Microsoft Teams, Yammer Groups, and Microsoft 365 Groups), all Skype for Business conversations, all SharePoint and OneDrive for Business sites (including the sites for all Microsoft Teams, Yammer Groups, and Microsoft 365 Groups), and items in all Exchange public folders.
- **Specific locations:** Use this option to search specific content locations. You can search all content locations for a specific Office 365 service (such as searching all Exchange mailboxes or search all SharePoint sites) or you can search for content in specific locations of any of the Office 365 services that are displayed.

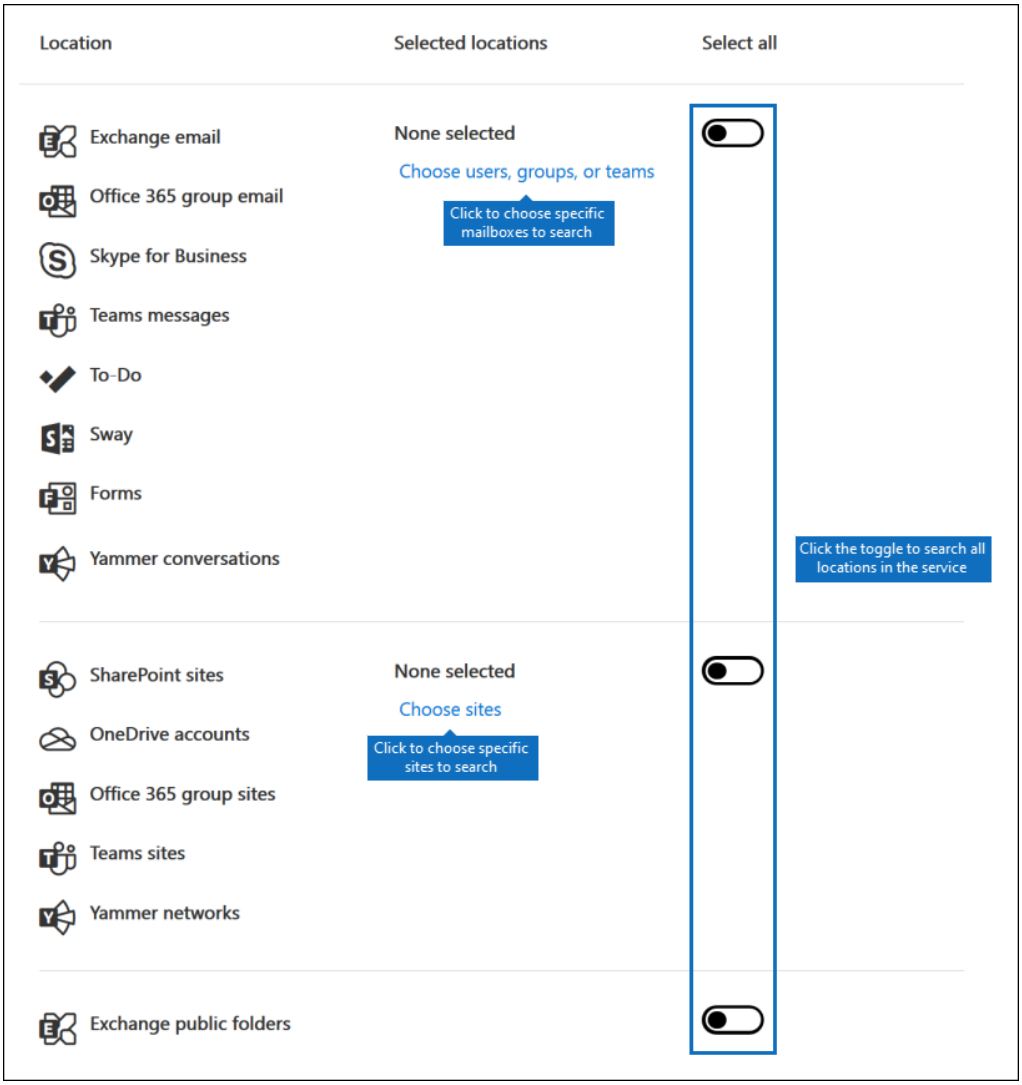

You can also add distribution groups to the list of Exchange mailboxes to search. For distribution groups, the mailboxes of group members are searched. Dynamic distribution groups aren't supported.

> **NOTE**
>
> When you search all mailbox locations or just specific mailboxes, data from other Office 365 applications that's saved to user mailboxes is included when you export the results of a Content Search. This data won't be included in the estimated search results and isn't available for preview. It is included when you export and download the search results. For more information, see Content stored in Exchange Online mailboxes.

5. After you've set up your search query, click **Save & run**.

6. On the **Save search** page, type a name for the search, and an optional description that helps identify the search. The name of the search has to be unique in your organization.

7. Click **Save** to start the search.

   After you save and run the search, any results returned by the search are displayed in the results pane. Depending on how you have the preview setting configured, the search results are display or you have to click **Preview results** to view them. See the next section for details.

To access this content search again or access other content searches listed on the **Content search** page, select the search and then click **Open**.

To clear the results or create another search, click **+ New search**.

Preview search results

There are two configuration settings for previewing search results. After you run a new search or open an existing search, click **Individual results** to view the following preview settings:

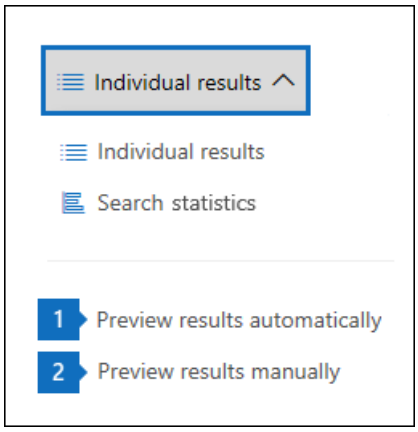

1. **Preview results automatically:** This setting displays the search results after you run a search.

2. **Preview results manually:** This setting displays placeholders in the search results pane, and displays the **Preview results** button that you have to click to display the search results. This is the default setting. It helps enhance search performance by not automatically displaying the search results when you open an existing search.

There are limits related to how many items are available to be previewed. For more information, see Limits for Content Search.

For a list of supported file types that can be previewed, see Previewing search results in the "More information about content search" section. If a file type isn't supported for preview or to download a copy of a document, you can click **Download original file** to download it to your local computer. For .aspx Web pages, the URL for the page is included though you might not have permissions to access the page.

Also note that unindexed items aren't available for previewing.

View information and statistics about a search

After you create and run a content search, you can view statistics about the estimated search results. This includes a summary of the search results, the query statistics such as the number of content locations with items that match the search query, and the name of content locations that have the most matching items. You can display statistics for one or more content searches. This lets you quickly compare the results for multiple searches and make decisions about the effectiveness of your search queries.

You can also download the search statistics and keyword statistics to a CSV file. This lets you use the filtering

and sorting features in Excel to compare results, and prepare reports for your search results.

To view search statistics:

1. On the **Content search** page, click **Open** and then click the search that you want to view the statistic for.
2. On the flyout page, click **Open query**.
3. In the **Individual results** drop down list, click **Search statistics**.
4. In the **Type** drop down list, click one of the following options depending on the search statistics you want to view:

   - **Summary:** Displays statistics for each type of content locations searched. This contents the number of content locations that contained items that matched the search query, and the total number and size of search result items. This is the default setting.
   - **Queries:** Displays statistics about the search query. This includes the type of content location the query statistics are applicable to, part of the search query the statistics are applicable to (note that **Primary** indicates the entire search query), the number of the content locations that contain items that match the search query, and the total number and size and items that were found (in the specified content location) that match the search query. Statistics for unindexed items (also called *partially indexed items*) are also displayed. However, only partially indexed items from mailboxes are included in the statistics. Partially indexed items from SharePoint and OneDrive are not included in the statistics.
   - **Top locations:** Displays statistics about the number of items that match the search query in each content location. The top 1,000 locations are displayed.

For more detailed information about search statistics, see View keyword statistics for Content Search results.

Export search results

After a search is successfully run, you can export the search results to a local computer. When you export email results, they can be downloaded to your computer as PST files or as individual messages (.msg files). When you export content from SharePoint and OneDrive sites, copies of native Office documents are exported. There are also other documents and reports that are included with the exported search results. You can also export the search results report and not the actual items.

To export search results:

1. On the **Content search** page, click the search that you want to export the search results for.
2. On the flyout page, click **Export results**. You can also export a search results report.
3. Complete the sections on the **Export results** fly out page. Be sure to use the scroll bar to view all export options.

For more detailed instructions and troubleshooting tips, see:

- Export Content search results
- Export a Content search report

More information about content search

See the following sections for more information about Content searches.

Content search limits

Building a search query

Searching OneDrive accounts

Searching Microsoft Teams and Microsoft 365 Groups

Searching Yammer Groups

Searching inactive mailboxes

Searching disconnected or de-licensed mailboxes

Previewing search results

Partially indexed items

Searching for content in a SharePoint Multi-Geo environment

Content search limits

- For a description of the limits that are applied to Content search, see Limits for Content search.
- Microsoft collects performance information for Content searches run by all organizations. While the complexity of the search query can impact search times, the biggest factor that affects how long searches take is the number of mailboxes searched. Although Microsoft doesn't provide a Service Level Agreement for search times, the following table lists average search times for a Content Search based on the number of mailboxes included in the search.

NUMBER OF MAILBOXES	AVERAGE SEARCH TIME
100	30 seconds
1,000	45 seconds
10,000	4 minutes
25,000	10 minutes
50,000	20 minutes
100,000	25 minutes

Building a search query

For detailed information about creating a search query, using Boolean search operators and search conditions, and searching for sensitive information types and content shared with users outside your organization, see Keyword queries and search conditions for Content Search .

Keep the following things in mind when using the keyword list to create a search query.

- You have to select the **Show keyword list** checkbox and then type each keyword in a separate row to create a search query where the keywords (or keyword phrases) in each row are connected by the **OR** operator. If you paste a list of keywords in the keyword box or press the **Enter** key after typing a keyword, they won't be connected by the **OR** operator. Here are incorrect and correct examples of how to add a list of keywords.

  Incorrect

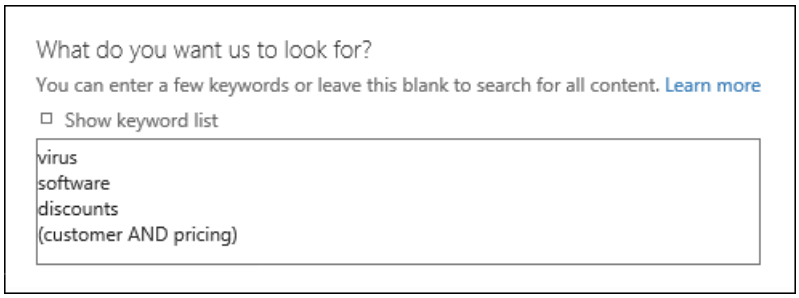

Correct

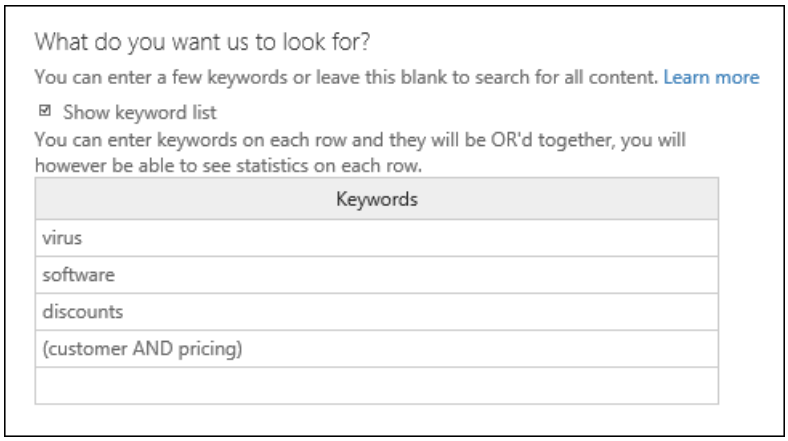

- You can also prepare a list of keywords or keyword phrases in an Excel file or a plain text file, and then copy and paste your list into the keyword list. To do this, you have to select the **Show keyword list** check box. Then, click the first row in the keyword list and paste your list. Each line from the Excel or text file is pasted into separate row in the keyword list.

- After you create a query using the keyword list, it's a good idea to verify the search query syntax to make the search query is what you intended. In the search query that's displayed under **Query** in the details pane, the keywords are separated by the text **(c:s)**. This indicates that the keywords are connected by a logical operator similar in functionality to the **OR** operator. Similarly, if your search query includes conditions, the keywords and the conditions are separated by the text **(c:c)**. This indicates that the keywords are connected to the conditions with a logical operator similar in functionality to the **AND** operator. Here's an example of the search query (displayed in the Details pane) that results when using the keyword list and a condition.

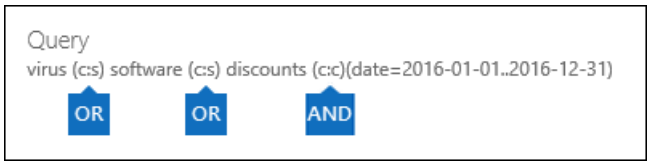

- When you run a content search, Microsoft 365 automatically checks your search query for unsupported characters and for Boolean operators that may not be capitalized. Unsupported characters are often hidden and typically cause a search error or return unintended results. For more information about the unsupported characters that are checked, see Check your Content Search query for errors.

- If you have a search query that contains keywords for non-English characters (such as Chinese characters), you can click **Query language-country/region** and select a language-country culture code value for the search. The default language/region is neutral. How can you tell if you need to change the language setting for a content search? If you're certain content locations contain the non-English characters you're searching for, but the search returns no results, the language setting may be the cause.

Searching OneDrive accounts

- To collect a list of the URLs for the OneDrive sites in your organization, see Create a list of all OneDrive locations in your organization. This script in this article creates a text file that contains a list of all

OneDrive sites. To run this script, you have to install and use the SharePoint Online Management Shell. Be sure to append the URL for your organization's MySite domain to each OneDrive site that you want to search. This is the domain that contains all your OneDrive; for example, `https://contoso-my.sharepoint.com`. Here's an example of a URL for a user's OneDrive site: `https://contoso-my.sharepoint.com/personal/sarad_contoso_onmicrosoft.com`.

In the rare case of a person's user principal name (UPN) being changed, the URL for their OneDrive location is changed to incorporate the new UPN. If this happens, you have to modify a content search by adding the user's new OneDrive URL and removing the old one. For more information, see How UPN changes affect the OneDrive URL.

Searching Microsoft Teams and Microsoft 365 Groups

You can search the mailbox that's associated with a Microsoft Team or Microsoft 365 Group. Because Microsoft Teams is built on Microsoft 365 Groups, searching them is similar. In both cases, only the group or team mailbox is searched. The mailboxes of the group or team members aren't searched. To search them, you have to specifically add them to the search.

Keep the following things in mind when searching for content in Microsoft Teams and Microsoft 365 Groups.

- To search for content located in Teams and Microsoft 365 Groups, you have to specify the mailbox and SharePoint site that are associated with a team or group.
- Content from private channels is stored in each user's mailbox, not the team mailbox. To search for content in private channels, see eDiscovery of private channels.
- Run the **Get-UnifiedGroup** cmdlet in Exchange Online to view properties for a team or a Microsoft 365 Group. This is a good way to get the URL for the site that's associated with a team or a group. For example, the following command displays selected properties for a Microsoft 365 Group named Senior Leadership Team:

```
Get-UnifiedGroup "Senior Leadership Team" | FL DisplayName,Alias,PrimarySmtpAddress,SharePointSiteUrl
DisplayName            : Senior Leadership Team
Alias                  : seniorleadershipteam
PrimarySmtpAddress     : seniorleadershipteam@contoso.onmicrosoft.com
SharePointSiteUrl      : https://contoso.sharepoint.com/sites/seniorleadershipteam
```

> **NOTE**
> To run the **Get-UnifiedGroup** cmdlet, you have to be assigned the View-Only Recipients role in Exchange Online or be a member of a role group that's assigned the View-Only Recipients role.

- When a user's mailbox is searched, any team or Microsoft 365 Group that the user is a member of won't be searched. Similarly, when you search a team or a Microsoft 365 Group, only the group mailbox and group site that you specify is searched. The mailboxes and OneDrive for Business accounts of group members aren't searched unless you explicitly add them to the search.
- To get a list of the members of a team or a Microsoft 365 Group, you can view the properties on the **Home > Groups** page in the Microsoft 365 admin center. Alternatively, you can run the following command in Exchange Online PowerShell:

```
Get-UnifiedGroupLinks <group or team name> -LinkType Members | FL DisplayName,PrimarySmtpAddress
```

> **NOTE**
>
> To run the **Get-UnifiedGroupLinks** cmdlet, you have to be assigned the View-Only Recipients role in Exchange Online or be a member of a role group that's assigned the View-Only Recipients role.

- Conversations that are part of a Teams channel are stored in the mailbox that's associated with the team. Similarly, files that team members share in a channel are stored on the team's SharePoint site. Therefore, you have to add the team mailbox and SharePoint site as a content location to search conversations and files in a channel.
- Alternatively, conversations that are part of the Chat list in Teams are stored in the Exchange Online mailbox of the users who participate in the chat. And files that a user shares in Chat conversations are stored in the OneDrive for Business account of the user who shares the file. Therefore, you have to add the individual user mailboxes and OneDrive for Business accounts as content locations to search conversations and files in the Chat list.

  > **NOTE**
  >
  > In an Exchange hybrid deployment, users with an on-premises mailbox might participate in conversations that are part of the Chat list in Teams. In this case, content from these conversations is also searchable because it's saved to a cloud-based storage area (called a *cloud-based mailbox for on-premises users*) for users who have an on-premises mailbox. For more information, see Search for Teams chat data for on-premises users.

- Every team or team channel contains a Wiki for note-taking and collaboration. The Wiki content is automatically saved to a file with a .mht format. This file is stored in the Teams Wiki Data document library on the team's SharePoint site. You can use the Content Search tool to search the Wiki by specifying the team's SharePoint site as the content location to search.

  > **NOTE**
  >
  > The capability to search the Wiki for a team or channel (when you search the team's SharePoint site) was released on June 22, 2017. Wiki pages that were saved or updated on that date or after are available to be searched. Wiki pages last saved or updated before that date aren't available for search.

- Summary information for meetings and calls in a Teams channel are also stored in the mailboxes of users who dialed into the meeting or call. This means you can use Content Search to search these summary records. Summary information includes:
  - Date, start time, end time, and duration of a meeting or call
  - The date and time when each participant joined or left the meeting or call
  - Calls sent to voice mail
  - Missed or unanswered calls
  - Call transfers, which are represented as two separate calls

  It can take up to 8 hours for meeting and call summary records to be available to be searched.

  In the search results, meeting summaries are identified as **Meeting** in the **Type field**, and call summaries are identified as **Call**. Also, conversations that are part of a Teams channel and 1xN chats are identified as **IM** in the **Type** field.

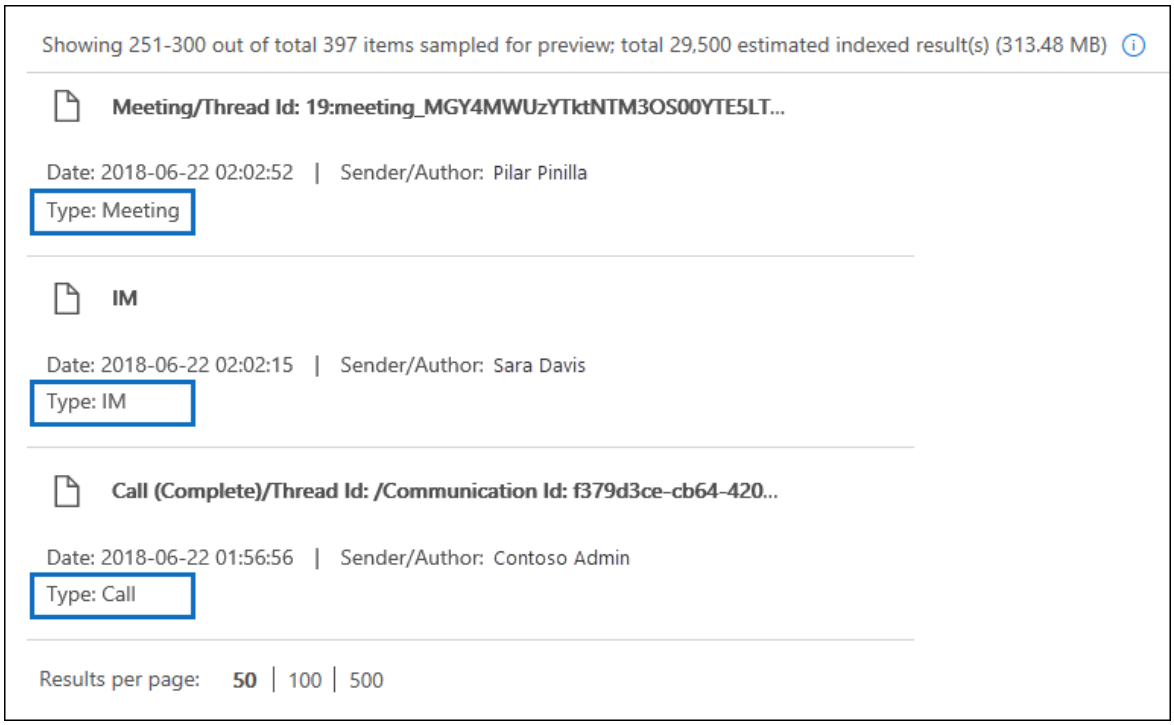

For more information, see Microsoft Teams launches eDiscovery for calls and meetings.

- Card content generated by apps in Teams channels, 1:1 chats, and 1xN chats is stored in mailboxes and can be searched. A *card* is a UI container for short pieces of content. Cards can have multiple properties and attachments, and can include buttons that can trigger card actions. For more information, see Cards

  Like other Teams content, where card content is stored is based on where the card was used. Content for cards used in a Teams channel is stored in the Teams group mailbox. Card content for 1:1 and 1xN chats are stored in the mailboxes of the chat participants.

  To search for card content, you can use the `kind:microsoftteams` or `itemclass:IPM.SkypeTeams.Message` search conditions. When reviewing search results, card content generated by bots in a Teams channel have the **Sender/Author** email property as `<appname>@teams.microsoft.com`, where `appname` is the name of the app that generated the card content. If card content was generated by a user, the value of **Sender/Author** identifies the user.

  When viewing card content in Content search results, the content appears as an attachment to the message. The attachment is named `appname.html`, where `appname` is the name of the app that generated the card content. The following screenshots show how card content (for an app named Asana) appears in Teams and in the results of a search.

  #### Card content in Teams

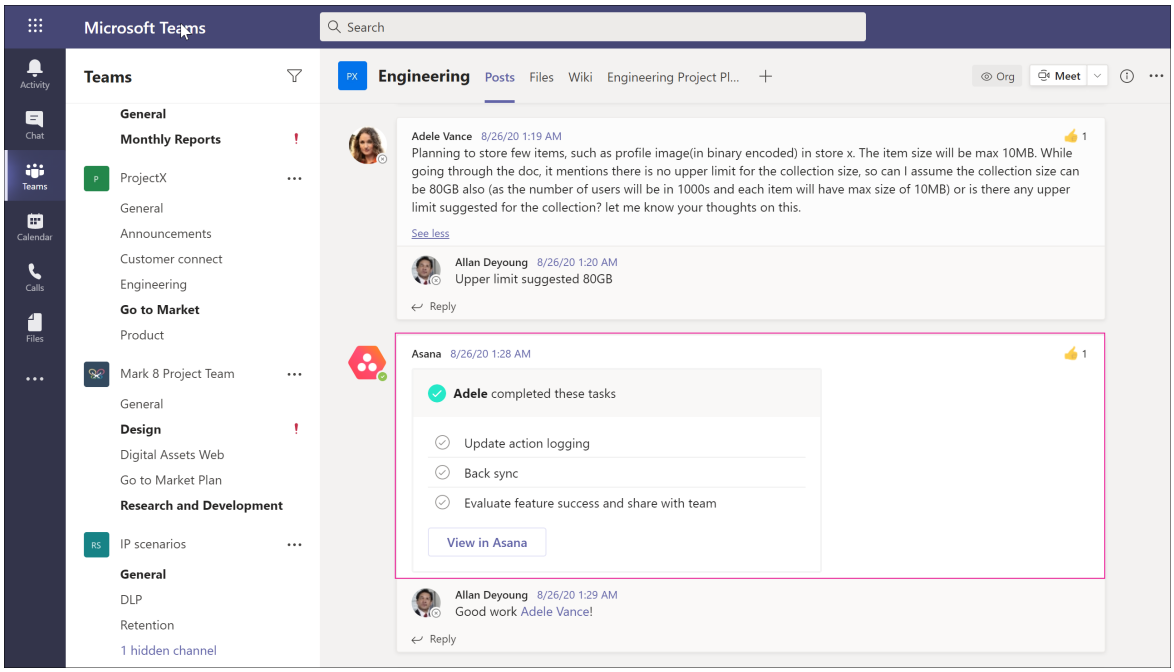

Card content in search results

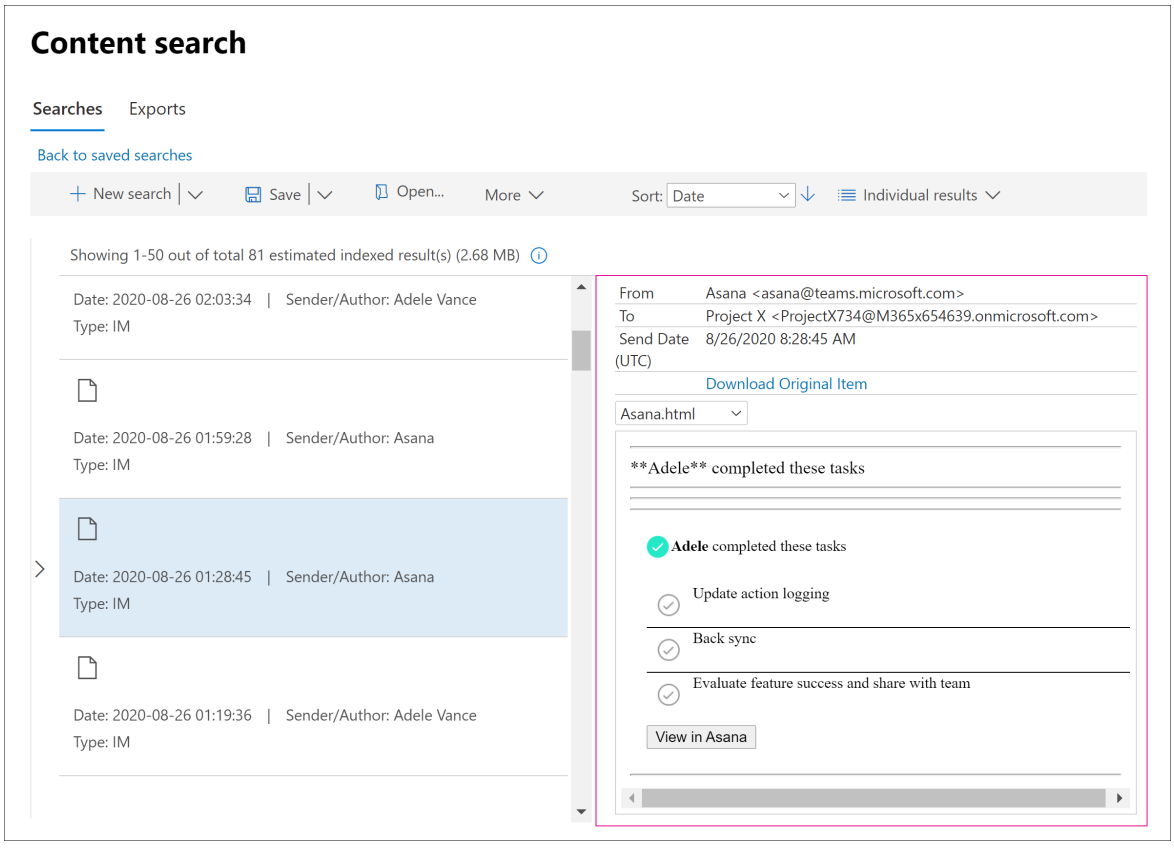

> **NOTE**
>
> To display images from card content in search results at this time (such as the checkmarks in the previous screenshot), you have to be signed into Teams (at https://teams.microsoft.com) in a different tab in the same browser session that you use to view the search results. Otherwise, image placeholders are displayed.

- You can use the **Kind** email property or the **Message kind** search condition to search specifically for content in Teams.
  - To use the **Kind** property as part of the keyword search query, in the **Keywords** box of a search query, type `kind:microsoftteams`.

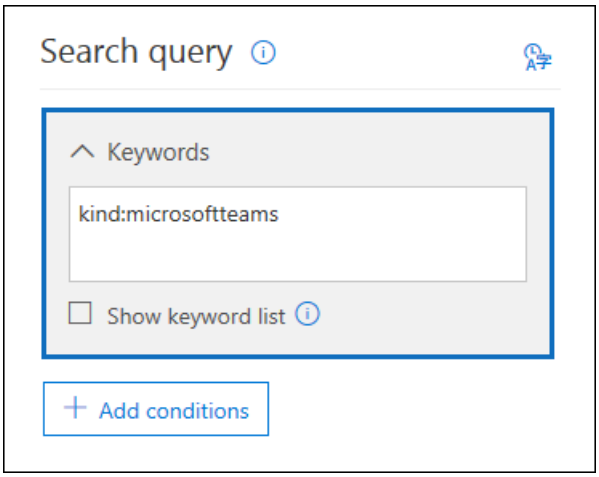

- To use a search condition, add the **Message kind** condition and use the value `microsoftteams`.

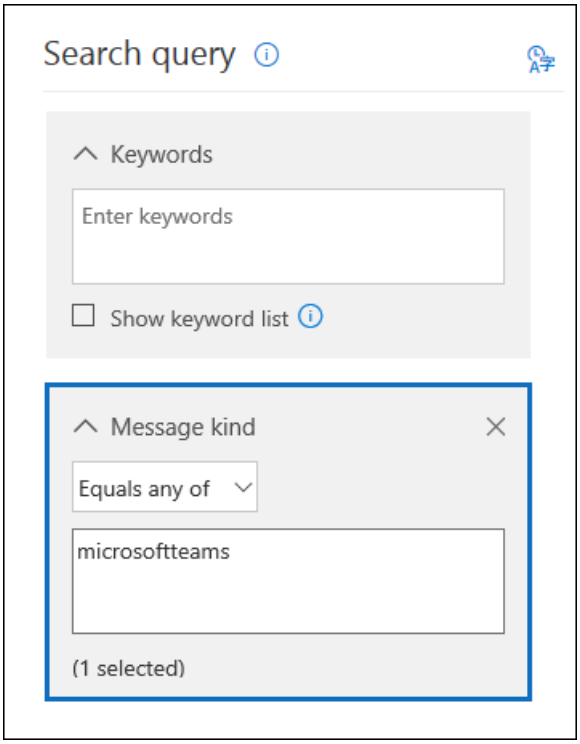

Conditions are logically connected to the keyword query by the **AND** operator. That means an item must match both the keyword query and the search condition to be returned in the search results. For more information, see the "Guidelines for using conditions" section in Keyword queries and search conditions for Content Search.

Searching Yammer Groups

You can use the **ItemClass** email property or the **Type** search condition to search specifically for conversation items in Yammer Groups.

- To use the **ItemClass** property as part of the keyword search query, in the **Keywords** box of a search query, you can type one (or all) of the following property:value pairs:
  - ItemClass:IPM.Yammer.message
  - ItemClass:IPM.Yammer.poll
  - ItemClass:IPM.Yammer.praise
  - ItemClass:IPM.Yammer.question

  For example, you can use the following search query to return Yammer messages and Yammer praise items:

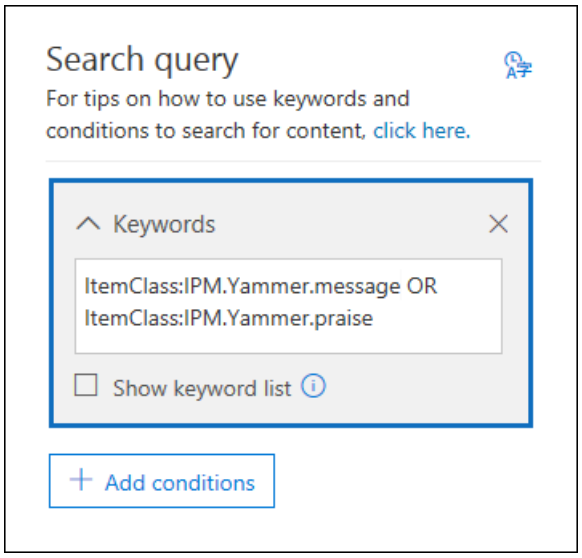

- Alternatively, you can use the **Type** email condition and select **Yammer messages** to return Yammer items. For example, the following search query will return all Yammer conversation items that contain the keyword "confidential".

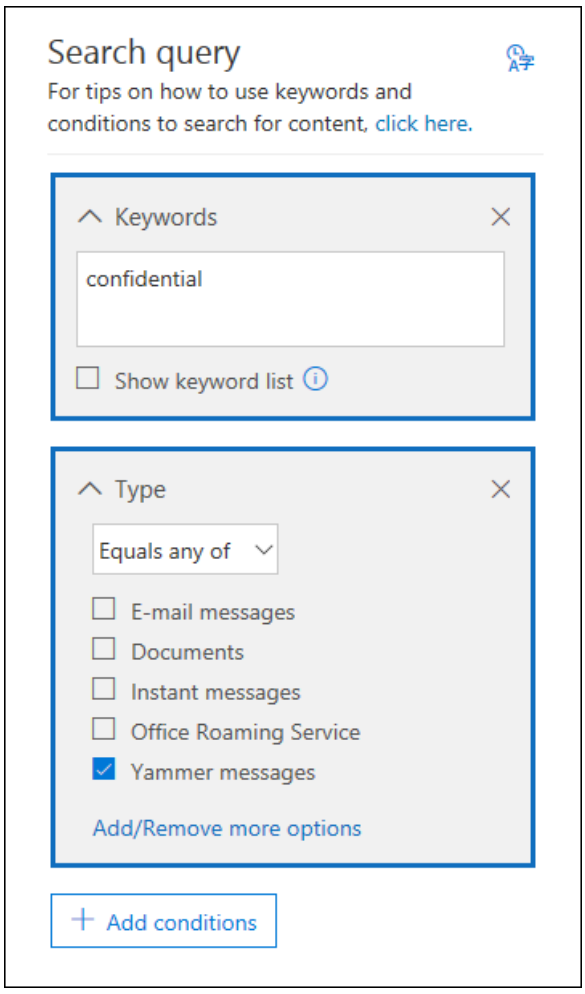

Searching inactive mailboxes

You can search inactive mailboxes in a content search. To get a list of the inactive mailboxes in your organization, run the command `Get-Mailbox -InactiveMailboxOnly` in Exchange Online PowerShell. Alternatively, you can go to **Information governance** > **Retention** in the Security & Compliance Center, and then click **More** > **Inactive mailboxes**.

Here are a few things to keep in mind when searching inactive mailboxes.

- If an existing content search includes a user mailbox and that mailbox is made inactive, the content search

will continue to search the inactive mailbox when you rerun the search after it becomes inactive.

- Sometimes a user may have an active mailbox and an inactive mailbox that have the same SMTP address. In this case, only the specific mailbox that you select as a location for a content search is searched. In other words, if you add a user's mailbox to a search, you can't assume that both their active and inactive mailboxes are searched. Only the mailbox that you explicitly add to the search is searched.

- You can use Security & Compliance Center PowerShell to create a content search to search an inactive mailbox. To do this, you have to pre-append a period (.) to the email address of the inactive mailbox. For example, the following command creates a content search that searches an inactive mailbox with the email address pavelb@contoso.onmicrosoft.com:

  ```
  New-ComplianceSearch -Name InactiveMailboxSearch -ExchangeLocation .pavelb@contoso.onmicrosoft.com -
  AllowNotFoundExchangeLocationsEnabled $true
  ```

- We strongly recommend that you avoid having an active mailbox and inactive mailbox with the same SMTP address. If you need to reuse the SMTP address that is assigned to an inactive mailbox, we recommend that you recover the inactive mailbox or restore the contents of an inactive mailbox to an active mailbox (or the archive of an active mailbox), and then delete the inactive mailbox. For more information, see one of the following topics:

  - Recover an inactive mailbox in Office 365
  - Restore an inactive mailbox in Office 365
  - Delete an inactive mailbox in Office 365

Searching disconnected or de-licensed mailboxes

If the Exchange Online license (or the entire Microsoft 365 license) is removed from a user account or in Azure Active Directory, the user's mailbox becomes a *disconnected* mailbox. This means that the mailbox is no longer associated with the user account. Here's what happens when searching disconnected mailboxes:

- If the license is removed from a mailbox, the mailbox is no longer searchable.
- If an existing content search includes a mailbox in which the license is removed, no search results from the disconnected mailbox will be returned if you rerun the content search.
- If you use the **New-ComplianceSearch** cmdlet to create a content search and specify a disconnected mailbox as the Exchange content location to search, the content search won't return any search results from the disconnected mailbox.

If you need to preserve the data in a disconnected mailbox so that it's searchable, you must place a hold on the mailbox before removing the license. This preserves the data and keeps the disconnected mailbox searchable until the hold is removed. For more information about holds, see How to identify the type of hold placed on an Exchange Online mailbox.

Previewing search results

You can preview supported file types in the preview pane. If a file type isn't supported, you have to download a copy of the file to your local computer to view it. The following file types are supported and can be previewed in the search results pane.

- .txt, .html, .mhtml
- .eml
- .doc, .docx, .docm
- .pptm, .pptx

- .pdf

Also, the following file container types are supported. You can view the list of files in the container in the preview pane.

- .zip
- .gzip

Partially indexed items

- As previously explained, partially indexed items in mailboxes are included in the estimated search results. Partially indexed items from SharePoint and OneDrive aren't included in the estimated search results.
- If a partially indexed item matches the search query (because other message or document properties meet the search criteria), it isn't included in the estimated number of unindexed items. If a partially indexed item is excluded by the search criteria, it isn't included in the estimated number of unindexed items. For more information, see Partially indexed items in Content Search in Office 365.

Searching for content in a SharePoint Multi-Geo environment

If it's necessary for an eDiscovery manager to search for content in SharePoint and OneDrive in different regions in a SharePoint multi-geo environment, then you need to do the following things to make that happen:

1. Create a separate user account for each satellite geo location that the eDiscovery manager needs to search. To search for content in sites in that geo location, the eDiscovery manager must sign in to the account you created for that location and then run a content search.
2. Create a search permissions filter for each satellite geo location (and corresponding user account) the eDiscovery manager needs to search. Each of these search permissions filters limits the scope of the content search to a specific geo location when the eDiscovery manager is signed in to the user account associated with that location.

> **TIP**
> You don't have to use this strategy when using the search tool in Advanced eDiscovery. That's because all datacenters are searched when you search SharePoint sites and OneDrive accounts in Advanced eDiscovery. You have to use this strategy of region-specific user accounts and search permissions filters only when using the Content Search tool and running searches associated with eDiscovery cases.

For example, let's say that an eDiscovery manager needs to search for SharePoint and OneDrive content in satellite locations in North American, Europe, and Asia Pacific. The first step is to create three users accounts, one for each location. The next step is to create three search permissions filters, one for each location *and* corresponding user account. Here are examples of the three search permissions filters for this scenario. In each of these examples, the **Region** specifies the SharePoint datacenter location for that geo and the **Users** parameter specifies the corresponding user account.

North America

```
New-ComplianceSecurityFilter -FilterName "SPMultiGeo-NAM" -Users ediscovery-nam@contoso.com -Region NAM -
Action ALL
```

Europe

```
New-ComplianceSecurityFilter -FilterName "SPMultiGeo-EUR" -Users ediscovery-eur@contoso.com -Region EUR -
Action ALL
```

Asia Pacific

```
New-ComplianceSecurityFilter -FilterName "SPMultiGeo-APC" -Users ediscovery-apc@contoso.com -Region APC -
Action ALL
```

Keep the following things in mind when using search permissions filters to search for content in multi-geo environments:

- The **Region** parameter directs searches to the specified satellite location. If an eDiscovery manager only searches SharePoint and OneDrive sites outside of the region specified in the search permissions filter, no search results are returned.
- The **Region** parameter doesn't control searches of Exchange mailboxes. All datacenters are searched when you search mailboxes.

For more information about using search permissions filters in a multi-geo environment, see the "Searching and exporting content in Multi-Geo environments" section in Set up compliance boundaries for eDiscovery investigations.

Keyword queries and search conditions for Content Search and eDiscovery

2/18/2021 • 34 minutes to read • Edit Online

This topic describes the email and document properties that you can search for in email items in Exchange Online and documents stored on SharePoint and OneDrive for Business sites by using the Content Search feature in the Microsoft 365 compliance center. You can also use the ***-ComplianceSearch** cmdlets in Security & Compliance Center PowerShell to search for these properties. The topic also describes:

- Using Boolean search operators, search conditions, and other search query techniques to refine your search results.
- Searching for sensitive data types and custom sensitive data types in SharePoint and OneDrive for Business.
- Searching for site content that's shared with users outside of your organization

For step-by-step instructions on how to create a Content Search, see Content Search.

> **NOTE**
> Content Search in the Microsoft 365 compliance center and the corresponding ***-ComplianceSearch** cmdlets in Security & Compliance Center PowerShell use the Keyword Query Language (KQL). For more detailed information, see Keyword Query Language syntax reference.

Searchable email properties

The following table lists email message properties that can be searched by using the Content Search feature in the Microsoft 365 compliance center or by using the **New-ComplianceSearch** or the **Set-ComplianceSearch** cmdlet. The table includes an example of the *property:value* syntax for each property and a description of the search results returned by the examples. You can type these `property:value` pairs in the keywords box for a Content Search.

> **NOTE**
> When searching email properties, it's not possible to search for items in which the specified property is empty or blank. For example, using the *property:value* pair of **subject:""** to search for email messages with an empty subject line will return zero results. This also applies when searching site and contact properties.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLES	SEARCH RESULTS RETURNED BY THE EXAMPLES

PROPERTY	PROPERTY DESCRIPTION	EXAMPLES	SEARCH RESULTS RETURNED BY THE EXAMPLES

AttachmentNames	The names of files attached to an email message.	`attachmentnames:annualreport.ppt`
		`attachmentnames:annual*`
		`attachmentnames:.pptx`
			Messages that have an attached file named annualreport.ppt. In the second example, using the wildcard returns messages with the word "annual" in the file name of an attachment. The third example returns all attachments with the pptx file extension.

Bcc	The Bcc field of an email message.[1]	`bcc:pilarp@contoso.com`
		`bcc:pilarp`
		`bcc:"Pilar Pinilla"`
			All examples return messages with Pilar Pinilla included in the Bcc field.

Category	The categories to search. Categories can be defined by users by using Outlook or Outlook on the web (formerly known as Outlook Web App). The possible values are: blue green orange purple red yellow	`category:"Red Category"`	Messages that have been assigned the red category in the source mailboxes.
Cc	The Cc field of an email message.[1]	`cc:pilarp@contoso.com` `cc:"Pilar Pinilla"`	In both examples, messages with Pilar Pinilla specified in the Cc field.
Folderid	The folder ID (GUID) of a specific mailbox folder. If you use this property, be sure to search the mailbox that the specified folder is located in. Only the specified folder will be searched. Any subfolders in the folder won't be searched. To search subfolders, you need to use the Folderid property for the subfolder you want to search. For more information about searching for the Folderid property and using a script to obtain the folder IDs for a specific mailbox, see [Use Content Search for targeted collections]().	`folderid:4D6DD7F943C29041A65787E30F02AD1F00000000013A0000` `folderid:2370FB455F82FC44BE31397F4B6D52A70000000001160000 AND participants:garthf@contoso.com`	The first example returns all items in the specified mailbox folder. The second example returns all items in the specified mailbox folder that were sent or received by garthf@contoso.com.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLES	SEARCH RESULTS RETURNED BY THE EXAMPLES
From	The sender of an email message.[1]	`from:pilarp@contoso.com` `from:contoso.com`	Messages sent by the specified user or sent from a specified domain.
HasAttachment	Indicates whether a message has an attachment. Use the values **true** or **false**.	`from:pilar@contoso.com AND hasattachment:true`	Messages sent by the specified user that have attachments.

Importance	The importance of an email message, which a sender can specify when sending a message. By default, messages are sent with normal importance, unless the sender sets the importance as **high** or **low**.	`importance:high`
		`importance:medium`
		`importance:low`
			Messages that are marked as high importance, medium importance, or low importance.

IsRead	Indicates whether messages have been read. Use the values **true** or **false**.	`isread:true` `isread:false`	The first example returns messages with the IsRead property set to **True**. The second example returns messages with the IsRead property set to **False**.
ItemClass	Use this property to search specific third-party data types that your organization imported to Office 365. Use the following syntax for this property: `itemclass:ipm.externaldata.<third-party data type>*`	`itemclass:ipm.externaldata.Facebook* AND subject:contoso` `itemclass:ipm.externaldata.Twitter* AND from:"Ann Beebe" AND "Northwind Traders"`	The first example returns Facebook items that contain the word "contoso" in the Subject property. The second example returns Twitter items that were posted by Ann Beebe and that contain the keyword phrase "Northwind Traders". For a complete list of values to use for third-party data types for the ItemClass property, see Use Content Search to search third-party data that was imported to Office 365.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLES	SEARCH RESULTS RETURNED BY THE EXAMPLES

Kind	The type of email message to search for. Possible values:
	contacts
	docs
	email
	externaldata
	faxes
	im
	journals
	meetings
	microsoftteams (returns items from chats, meetings, and calls in Microsoft Teams)
	notes
	posts
	rssfeeds
	tasks
	voicemail
		`kind:email`
		`kind:email OR kind:im OR kind:voicemail`
		`kind:externaldata`
			The first example returns email messages that meet the search criteria. The second example returns email messages, instant messaging conversations (including Skype for Business conversations and chats in Microsoft Teams), and voice messages that meet the search criteria. The third example returns items that were imported to mailboxes in Microsoft 365 from third-party data sources, such as Twitter, Facebook, and Cisco Jabber, that meet the search criteria. For more information, see Archiving third-party data in Office 365.

Participants	All the people fields in an email message. These fields are From, To, Cc, and Bcc.[1]	`participants:garthf@contoso.com` `participants:contoso.com`	Messages sent by or sent to garthf@contoso.com. The second example returns all messages sent by or sent to a user in the contoso.com domain.
Received	The date that an email message was received by a recipient.	`received:04/15/2016` `received>=01/01/2016 AND received<=03/31/2016`	Messages that were received on April 15, 2016. The second example returns all messages received between January 1, 2016 and March 31, 2016.
Recipients	All recipient fields in an email message. These fields are To, Cc, and Bcc.[1]	`recipients:garthf@contoso.com` `recipients:contoso.com`	Messages sent to garthf@contoso.com. The second example returns messages sent to any recipient in the contoso.com domain.
Sent	The date that an email message was sent by the sender.	`sent:07/01/2016` `sent>=06/01/2016 AND sent<=07/01/2016`	Messages that were sent on the specified date or sent within the specified date range.
Size	The size of an item, in bytes.	`size>26214400` `size:1..1048567`	Messages larger than 25?? MB. The second example returns messages from 1 through 1,048,567 bytes (1 MB) in size.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLES	SEARCH RESULTS RETURNED BY THE EXAMPLES
Subject	The text in the subject line of an email message. **Note:** When you use the Subject property in a query, the search returns all messages in which the subject line contains the text you're searching for. In other words, the query doesn't return only those messages that have an exact match. For example, if you search for `subject:"Quarterly Financials"`, your results will include messages with the subject "Quarterly Financials 2018".	`subject:"Quarterly Financials"` `subject:northwind`	Messages that contain the phrase "Quarterly Financials" anywhere in the text of the subject line. The second example returns all messages that contain the word northwind in the subject line.

To	The To field of an email message.[1]	`to:annb@contoso.com`
		`to:annb`
		`to:"Ann Beebe"`
			All examples return messages where Ann Beebe is specified in the To: line.

NOTE

[1] For the value of a recipient property, you can use email address (also called *user principal name* or UPN), display name, or alias to specify a user. For example, you can use annb@contoso.com, annb, or "Ann Beebe" to specify the user Ann Beebe.

Recipient expansion

When searching any of the recipient properties (From, To, Cc, Bcc, Participants, and Recipients), Microsoft 365 attempts to expand the identity of each user by looking them up in Azure Active Directory (Azure AD). If the user is found in Azure AD, the query is expanded to include the user's email address (or UPN), alias, display name, and LegacyExchangeDN. For example, a query such as `participants:ronnie@contoso.com` expands to `participants:ronnie@contoso.com OR participants:ronnie OR participants:"Ronald Nelson" OR participants:"<LegacyExchangeDN>"`.

To prevent recipient expansion, add a wild card character (asterisk) to the end of the email address and use a reduced domain name; for example, `participants:"ronnie@contoso*"` Be sure to surround the email address with double quotation marks.

However, be aware that preventing recipient expansion in the search query may result in relevant items not being returned in the search results. Email messages in Exchange can be saved with different text formats in the recipient fields. Recipient expansion is intended to help mitigate this fact by returning messages that may contain different text formats. So preventing recipient expansion may result in the search query not returning all items that may be relevant to your investigation.

NOTE

If you need to review or reduce the items returned by a search query due to recipient expansion, consider using Advanced eDiscovery. You can search for messages (taking advantage of recipient expansion), add them to a review set, and then use review set queries or filters to review or narrow the results. For more information, see Collect data for a case and Query the data in a review set.

Searchable site properties

The following table lists some of the SharePoint and OneDrive for Business properties that can be searched by using the Content Search feature in the Security & Compliance Center or by using the **New-ComplianceSearch** or the **Set-ComplianceSearch** cmdlet. The table includes an example of the *property:value* syntax for each property and a description of the search results returned by the examples.

For a complete list of SharePoint properties that can be searched, see Overview of crawled and managed properties in SharePoint. Properties marked with a **Yes** in the **Queryable** column can be searched.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLE	SEARCH RESULTS RETURNED BY THE EXAMPLES
Author	The author field from Office documents, which persists if a document is copied. For example, if a user creates a document and the emails it to someone else who then uploads it to SharePoint, the document will still retain the original author. Be sure to use the user's display name for this property.	`author:"Garth Fort"`	All documents that are authored by Garth Fort.
ContentType	The SharePoint content type of an item, such as Item, Document, or Video.	`contenttype:document`	All documents would be returned.
Created	The date that an item is created.	`created>=06/01/2016`	All items created on or after June 1, 2016.
CreatedBy	The person that created or uploaded an item. Be sure to use the user's display name for this property.	`createdby:"Garth Fort"`	All items created or uploaded by Garth Fort.
DetectedLanguage	The language of an item.	`detectedlanguage:english`	All items in English.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLE	SEARCH RESULTS RETURNED BY THE EXAMPLES
DocumentLink	The path (URL) of a specific folder on a SharePoint or OneDrive for Business site. If you use this property, be sure to search the site that the specified folder is located in. To return items located in subfolders of the folder that you specify for the documentlink property, you have to add /* to the URL of the specified folder; for example, `documentlink: "https://contoso.sharepoint.com/Shared Documents/*"` For more information about searching for the documentlink property and using a script to obtain the documentlink URLs for folders on a specific site, see Use Content Search for targeted collections.	`documentlink:"https://contoso-my.sharepoint.com/personal/garthf_contoso_com/Documents/Private` `documentlink:"https://contoso-my.sharepoint.com/personal/garthf_contoso_com/Documents/Shared with Everyone/*" AND filename:confidential`	The first example returns all items in the specified OneDrive for Business folder. The second example returns documents in the specified site folder (and all subfolders) that contain the word "confidential" in the file name.
FileExtension	The extension of a file; for example, docx, one, pptx, or xlsx.	`fileextension:xlsx`	All Excel files (Excel 2007 and later)
FileName	The name of a file.	`filename:"marketing plan"` `filename:estimate`	The first example returns files with the exact phrase "marketing plan" in the title. The second example returns files with the word "estimate" in the file name.
LastModifiedTime	The date that an item was last changed.	`lastmodifiedtime>=05/01/2016` `lastmodifiedtime>=05/10/2016 AND lastmodifiedtime<=06/1/2016`	The first example returns items that were changed on or after May 1, 2016. The second example returns items changed between May 1, 2016 and June 1, 2016.
ModifiedBy	The person who last changed an item. Be sure to use the user's display name for this property.	`modifiedby:"Garth Fort"`	All items that were last changed by Garth Fort.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLE	SEARCH RESULTS RETURNED BY THE EXAMPLES
Path	The path (URL) of a specific site in a SharePoint or OneDrive for Business site. To return items located in folders in the site that you specify for the path property, you have to add /* to the URL of the specified site; for example, `path: "https://contoso.sharepoint.com/Shared Documents/*"` **Note:** Using the `Path` property to search OneDrive locations won't return media files, such as .png, .tiff, or .wav files, in the search results. Use a different site property in your search query to search for media files in OneDrive folders.	`path:"https://contoso-my.sharepoint.com/personal/garthf_contoso_com/"` `path:"https://contoso-my.sharepoint.com/personal/garthf_contoso_com/*" AND filename:confidential`	The first example returns all items in the specified OneDrive for Business site. The second example returns documents in the specified site (and folders in the site) that contain the word "confidential" in the file name.
SharedWithUsersOWSUser	Documents that have been shared with the specified user and displayed on the **Shared with me** page in the user's OneDrive for Business site. These are documents that have been explicitly shared with the specified user by other people in your organization. When you export documents that match a search query that uses the SharedWithUsersOWSUser property, the documents are exported from the original content location of the person who shared the document with the specified user. For more information, see Searching for site content shared within your organization.	`sharedwithusersowsuser:garthf` `sharedwithusersowsuser:"garthf@contoso.com"`	Both examples return all internal documents that have been explicitly shared with Garth Fort and that appear on the **Shared with me** page in Garth Fort's OneDrive for Business account.
Site	The URL of a site or group of sites in your organization.	`site:"https://contoso-my.sharepoint.com"` `site:"https://contoso.sharepoint.com/sites/teamA"`	The first example returns items from the OneDrive for Business sites for all users in the organization. The second example returns items from all team sites.
Size	The size of an item, in bytes.	`size>=1` `size:1..10000`	The first example returns items larger than 1 byte. The second example returns items from 1 through 10,000 bytes in size.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLE	SEARCH RESULTS RETURNED BY THE EXAMPLES
Title	The title of the document. The Title property is metadata that's specified in Microsoft Office documents. It's different from the file name of the document.	`title:"communication plan"`	Any document that contains the phrase "communication plan" in the Title metadata property of an Office document.

Searchable contact properties

The following table lists the contact properties that are indexed and that you can search for using Content Search. These are the properties that are available for users to configure for the contacts (also called personal contacts) that are located in the personal address book of a user's mailbox. To search for contacts, you can select the mailboxes to search and then use one or more contact properties in the keyword query.

> **TIP**
> To search for values that contain spaces or special characters, use double quotation marks (" ") to contain the phrase; for example, `businessaddress:"123 Main Street"`.

PROPERTY	PROPERTY DESCRIPTION		
BusinessAddress	The address in the **Business Address** property. The property is also called the **Work** address on the contact properties page.		
BusinessPhone	The phone number in any of the **Business Phone** number properties.		
CompanyName	The name in the **Company** property.		
Department	The name in the **Department** property.		
DisplayName	The display name of the contact. This is the name in the **Full Name** property of the contact.		
EmailAddress	The address for any email address property for the contact. Users can add multiple email addresses for a contact. Using this property would return contacts that match any of the contact's email addresses.		

PROPERTY	PROPERTY DESCRIPTION		
FileAs	The **File as** property. This property is used to specify how the contact is listed in the user's contact list. For example, a contact could be listed as *FirstName,LastName* or *LastName,FirstName*.		
GivenName	The name in the **First Name** property.		
HomeAddress	The address in any of the **Home** address properties.		
HomePhone	The phone number in any of the **Home** phone number properties.		
IMAddress	The IM address property, which is typically an email address used for instant messaging.		
MiddleName	The name in the **Middle** name property.		
MobilePhone	The phone number in the **Mobile** phone number property.		
Nickname	The name in the **Nickname** property.		
OfficeLocation	The value in **Office** or **Office location** property.		
OtherAddress	The value for the **Other** address property.		
Surname	The name in the **Last** name property.		
Title	The title in the **Job title** property.		

Searchable sensitive data types

You can use eDiscovery search tools in the Microsoft 365 compliance center to search for sensitive data, such as credit card numbers or social security numbers, that is stored in documents on SharePoint and OneDrive for Business sites. You can do this by using the `SensitiveType` property and the name (or ID) of a sensitive information type in a keyword query. For example, the query `SensitiveType:"Credit Card Number"` returns documents that contain a credit card number. The query `SensitiveType:"U.S. Social Security Number (SSN)"` returns documents that contain a U.S. social security number.

To see a list of the sensitive information types that you can search for, go to **Data classifications** > **Sensitive info types** in the Microsoft 365 compliance center. Or you can use the **Get-DlpSensitiveInformationType** cmdlet in Security & Compliance Center PowerShell to display a list of sensitive information types.

For more information about creating queries using the `SensitiveType` property, see Form a query to find sensitive data stored on sites.

Limitations for searching sensitive data types

- To search for custom sensitive information types, you have to specify the ID of the sensitive information type in the `SensitiveType` property. Using the name of a custom sensitive information type (as shown in the example for built-in sensitive information types in the previous section) will return no results. Use the **Publisher** column on the **Sensitive info types** page in the compliance center (or the **Publisher** property in PowerShell) to differentiate between built-in and custom sensitive information types. Built-in sensitive data types have a value of `Microsoft Corporation` for the **Publisher** property.

  To display the name and ID for the custom sensitive data types in your organization, run the following command in Security & Compliance Center PowerShell:

  ```
  Get-DlpSensitiveInformationType | Where-Object {$_.Publisher -ne "Microsoft Corporation"} | FT
  Name,Id
  ```

  Then you can use the ID in the `SensitiveType` search property to return documents that contain the custom sensitive data type; for example, `SensitiveType:7e13277e-6b04-3b68-94ed-1aeb9d47de37`

- You can't use sensitive information types and the `SensitiveType` search property to search for sensitive data at-rest in Exchange Online mailboxes. However, you can use data loss prevention (DLP) policies to protect sensitive email data in transit. For more information, see Overview of data loss prevention policies and Search for and find personal data.

Search operators

Boolean search operators, such as **AND**, **OR**, and **NOT**, help you define more-precise searches by including or excluding specific words in the search query. Other techniques, such as using property operators (such as `>=` or `..`), quotation marks, parentheses, and wildcards, help you refine a search query. The following table lists the operators that you can use to narrow or broaden search results.

OPERATOR	USAGE	DESCRIPTION	
AND	keyword1 AND keyword2	Returns items that include all of the specified keywords or `property:value` expressions. For example, `from:"Ann Beebe" AND subject:northwind` would return all messages sent by Ann Beebe that contained the word northwind in the subject line. [2]	

OPERATOR	USAGE	DESCRIPTION	
+	keyword1 + keyword2 + keyword3	Returns items that contain *either* `keyword2` or `keyword3` *and* that also contain `keyword1`. Therefore, this example is equivalent to the query `(keyword2 OR keyword3) AND keyword1`. The query `keyword1 + keyword2` (with a space after the **+** symbol) isn't the same as using the **AND** operator. This query would be equivalent to `"keyword1 + keyword2"` and return items with the exact phase `"keyword1 + keyword2"`.	
OR	keyword1 OR keyword2	Returns items that include one or more of the specified keywords or `property:value` expressions. [2]	
NOT	keyword1 NOT keyword2 NOT from:"Ann Beebe" NOT kind:im	Excludes items specified by a keyword or a `property:value` expression. In the second example excludes messages sent by Ann Beebe. The third example excludes any instant messaging conversations, such as Skype for Business conversations that are saved to the Conversation History mailbox folder. [2]	
-	keyword1 -keyword2	The same as the **NOT** operator. So this query returns items that contain `keyword1` and would exclude items that contain `keyword2`.	
NEAR	keyword1 NEAR(n) keyword2	Returns items with words that are near each other, where n equals the number of words apart. For example, `best NEAR(5) worst` returns any item where the word "worst" is within five words of "best". If no number is specified, the default distance is eight words. [2]	

OPERATOR	USAGE	DESCRIPTION	
:	property:value	The colon (:) in the `property:value` syntax specifies that the value of the property being searched for contains the specified value. For example, `recipients:garthf@contoso.com` returns any message sent to garthf@contoso.com.	
=	property=value	The same as the : operator.	
<	property<value	Denotes that the property being searched is less than the specified value. [1]	
>	property>value	Denotes that the property being searched is greater than the specified value.[1]	
<=	property<=value	Denotes that the property being searched is less than or equal to a specific value.[1]	
>=	property>=value	Denotes that the property being searched is greater than or equal to a specific value.[1]	
..	property:value1..value2	Denotes that the property being searched is greater than or equal to value1 and less than or equal to value2.[1]	
" "	"fair value" subject:"Quarterly Financials"	Use double quotation marks (" ") to search for an exact phrase or term in keyword and `property:value` search queries.	

OPERATOR	USAGE	DESCRIPTION	
*	cat* subject:set*	Prefix wildcard searches (where the asterisk is placed at the end of a word) match for zero or more characters in keywords or `property:value` queries. For example, `title:set*` returns documents that contain the word set, setup, and setting (and other words that start with "set") in the document title. **Note:** You can use only prefix wildcard searches; for example, **cat*** or **set***. Suffix searches (***cat**), infix searches (**c*t**), and substring searches (***cat***) are not supported.	
( )	(fair OR free) AND (from:contoso.com) (IPO OR initial) AND (stock OR shares) (quarterly financials)	Parentheses group together Boolean phrases, `property:value` items, and keywords. For example, `(quarterly financials)` returns items that contain the words quarterly and financials.	

NOTE

[1] Use this operator for properties that have date or numeric values.

[2] Boolean search operators must be uppercase; for example, **AND**. If you use a lowercase operator, such as **and**, it will be treated as a keyword in the search query.

Search conditions

You can add conditions to a search query to narrow a search and return a more refined set of results. Each condition adds a clause to the KQL search query that is created and run when you start the search.

Conditions for common properties

Conditions for mail properties

Conditions for document properties

Operators used with conditions

Guidelines for using conditions

Examples of using conditions in search queries

Conditions for common properties

Create a condition using common properties when searching mailboxes and sites in the same search. The following table lists the available properties to use when adding a condition.

CONDITION	DESCRIPTION
Date	For email, the date a message was received by a recipient or sent by the sender. For documents, the date a document was last modified.
Sender/Author	For email, the person who sent a message. For documents, the person cited in the author field from Office documents. You can type more than one name, separated by commas. Two or more values are logically connected by the **OR** operator.
Size (in bytes)	For both email and documents, the size of the item (in bytes).
Subject/Title	For email, the text in the subject line of a message. For documents, the title of the document. As previously explained, the Title property is metadata specified in Microsoft Office documents. You can type the name of more than one subject/title, separated by commas. Two or more values are logically connected by the **OR** operator.
Compliance label	For both email and documents, retention labels that have been assigned to messages and documents automatically by autolabel policies or retention labels that have been manually assigned by users. Retention labels are used to classify email and documents for information governance and enforce retention rules based on the settings defined by the label. You can type part of the retention label name and use a wildcard or type the complete label name. For more information about retention labels, see Learn about retention policies and retention labels.

Conditions for mail properties

Create a condition using mail properties when searching mailboxes or public folders. The following table lists the email properties that you can use for a condition. These properties are a subset of the email properties that were previously described. These descriptions are repeated for your convenience.

CONDITION	DESCRIPTION
Message kind	The message type to search. This is the same property as the Kind email property. Possible values: contacts docs email externaldata faxes im journals meetings microsoftteams notes posts rssfeeds tasks voicemail

CONDITION	DESCRIPTION
Participants	All the people fields in an email message. These fields are From, To, Cc, and Bcc.
Type	The message class property for an email item. This is the same property as the ItemClass email property. It's also a multi-value condition. So to select multiple message classes, hold the **CTRL** key and then click two or more message classes in the drop-down list that you want to add to the condition. Each message class that you select in the list will be logically connected by the **OR** operator in the corresponding search query. For a list of the message classes (and their corresponding message class ID) that are used by Exchange and that you can select in the **Message class** list, see Item Types and Message Classes.
Received	The date that an email message was received by a recipient. This is the same property as the Received email property.
Recipients	All recipient fields in an email message. These fields are To, Cc, and Bcc.
Sender	The sender of an email message.
Sent	The date that an email message was sent by the sender. This is the same property as the Sent email property.
Subject	The text in the subject line of an email message.
To	The recipient of an email message in the To field.

Conditions for document properties

Create a condition using document properties when searching for documents on SharePoint and OneDrive for Business sites. The following table lists the document properties that you can use for a condition. These properties are a subset of the site properties that were previously described. These descriptions are repeated for your convenience.

CONDITION	DESCRIPTION
Author	The author field from Office documents, which persists if a document is copied. For example, if a user creates a document and the emails it to someone else who then uploads it to SharePoint, the document will still retain the original author.
Title	The title of the document. The Title property is metadata that's specified in Office documents. It's different than the file name of the document.
Created	The date that a document is created.
Last modified	The date that a document was last changed.

CONDITION	DESCRIPTION
File type	The extension of a file; for example, docx, one, pptx, or xlsx. This is the same property as the FileExtension site property.

Operators used with conditions

When you add a condition, you can select an operator that is relevant to type of property for the condition. The following table describes the operators that are used with conditions and lists the equivalent that is used in the search query.

OPERATOR	QUERY EQUIVALENT	DESCRIPTION
After	`property>date`	Used with date conditions. Returns items that were sent, received, or modified after the specified date.
Before	`property<date`	Used with date conditions. Returns items that were sent, received, or modified before the specified date.
Between	`date..date`	Use with date and size conditions. When used with a date condition, returns items there were sent, received, or modified within the specified date range. When used with a size condition, returns items whose size is within the specified range.
Contains any of	`(property:value) OR (property:value)`	Used with conditions for properties that specify a string value. Returns items that contain any part of one or more specified string values.
Doesn't contain any of	`-property:value` `NOT property:value`	Used with conditions for properties that specify a string value. Returns items that don't contain any part of the specified string value.
Doesn't equal any of	`-property=value` `NOT property=value`	Used with conditions for properties that specify a string value. Returns items that don't contain the specific string.
Equals	`size=value`	Returns items that are equal to the specified size.[1]
Equals any of	`(property=value) OR (property=value)`	Used with conditions for properties that specify a string value. Returns items that are an exact match of one or more specified string values.
Greater	`size>value`	Returns items where the specified property is greater than the specified value.[1]

OPERATOR	QUERY EQUIVALENT	DESCRIPTION
Greater or equal	`size>=value`	Returns items where the specified property is greater than or equal to the specified value.[1]
Less	`size<value`	Returns items that are greater than or equal to the specific value.[1]
Less or equal	`size<=value`	Returns items that are greater than or equal to the specific value.[1]
Not equal	`size<>value`	Returns items that don't equal the specified size.[1]

NOTE

[1] This operator is available only for conditions that use the Size property.

Guidelines for using conditions

Keep the following in mind when using search conditions.

- A condition is logically connected to the keyword query (specified in the keyword box) by the **AND** operator. That means that items have to satisfy both the keyword query and the condition to be included in the results. This is how conditions help to narrow your results.
- If you add two or more unique conditions to a search query (conditions that specify different properties), those conditions are logically connected by the **AND** operator. That means only items that satisfy all the conditions (in addition to any keyword query) are returned.
- If you add more than one condition for the same property, those conditions are logically connected by the **OR** operator. That means items that satisfy the keyword query and any one of the conditions are returned. So, groups of the same conditions are connected to each other by the **OR** operator and then sets of unique conditions are connected by the **AND** operator.
- If you add multiple values (separated by commas or semi-colons) to a single condition, those values are connected by the **OR** operator. That means items are returned if they contain any of the specified values for the property in the condition.
- The search query that is created by using the keywords box and conditions is displayed on the **Search** page, in the details pane for the selected search. In a query, everything to the right of the notation `(c:c)` indicates conditions that are added to the query.
- Conditions only add properties to the search query; the don't add operators. This is why the query displayed in the detail pane doesn't show operators to the right of the `(c:c)` notation. KQL adds the logical operators (according to the previously explained rules) when the executing the query.
- You can use the drag and drop control to resequence the order of conditions. Click on the control for a condition and move it up or down.
- As previously explained, some condition properties allow you to type multiple values. Each value is logically connected by the **OR** operator. This results in the same logic as having multiple instances of the same condition, where each has a single value. The following illustrations show an example of a single condition with multiple values and an example of multiple conditions (for the same property) with a single value. Both examples result in the same query:
  `(filetype:docx) OR (filetype:pptx) OR (filetype:xlsx)`

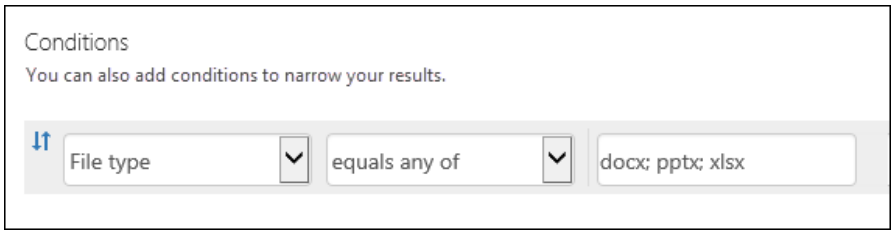

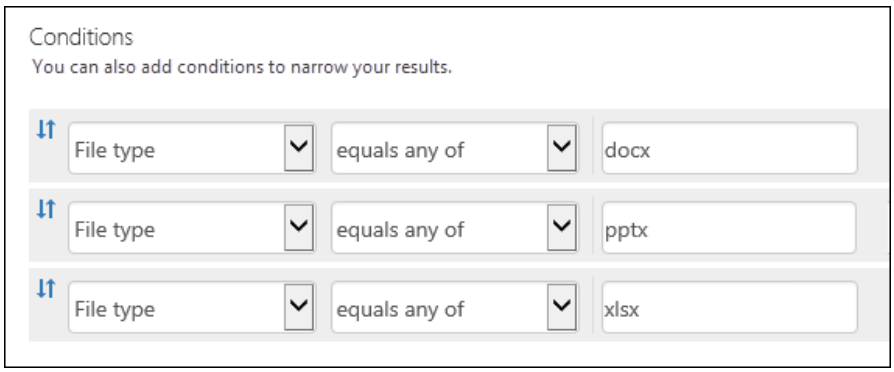

TIP

If a condition accepts multiple values, we recommend that you use a single condition and specify multiple values (separated by commas or semi-colons). This helps ensure the query logic that's applied is what you intend.

Examples of using conditions in search queries

The following examples show the GUI-based version of a search query with conditions, the search query syntax that is displayed in the details pane of the selected search (which is also returned by the **Get-ComplianceSearch** cmdlet), and the logic of the corresponding KQL query.

Example 1

This example returns documents on SharePoint and OneDrive for Business sites that contain a credit card number and were last modified before January 1, 2016.

GUI

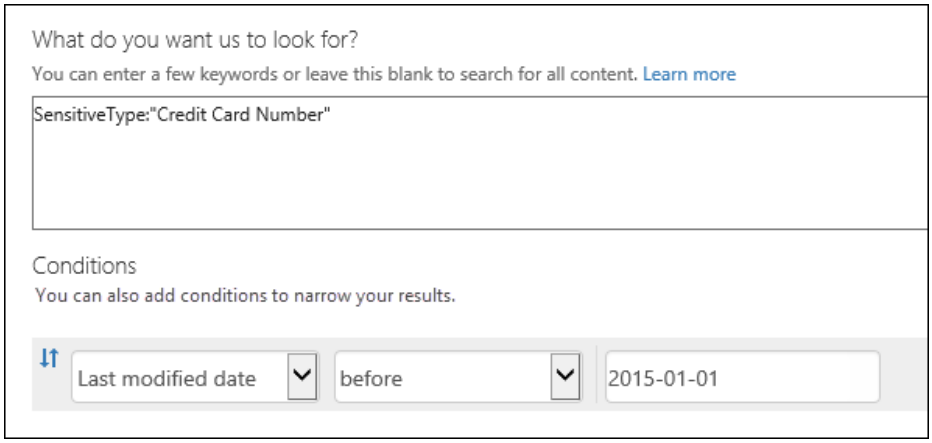

Search query syntax

```
SensitiveType:"Credit Card Number"(c:c)(lastmodifiedtime<2016-01-01)
```

Search query logic

```
SensitiveType:"Credit Card Number" AND (lastmodifiedtime<2016-01-01)
```

Example 2

This example returns email items or documents that contain the keyword "report", that were sent or created before April 1, 2105, and that contain the word "northwind" in the subject field of email messages or in the title

property of documents. The query excludes Web pages that meet the other search criteria.

GUI

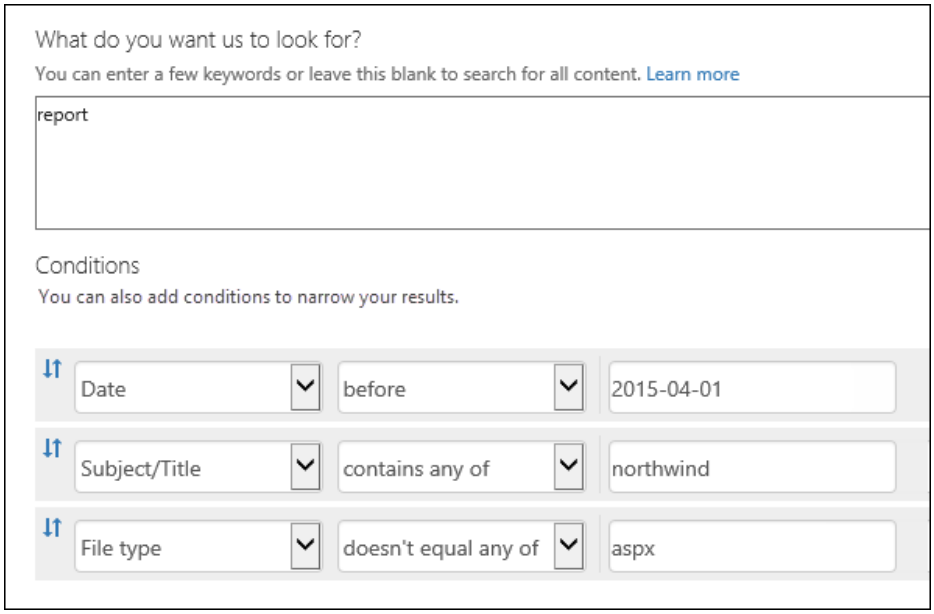

Search query syntax

```
report(c:c)(date<2016-04-01)(subjecttitle:"northwind")(-filetype:aspx)
```

Search query logic

```
report AND (date<2016-04-01) AND (subjecttitle:"northwind") NOT (filetype:aspx)
```

Example 3

This example returns email messages or calendar meetings that were sent between 12/1/2016 and 11/30/2016 and that contain words that start with "phone" or "smartphone".

GUI

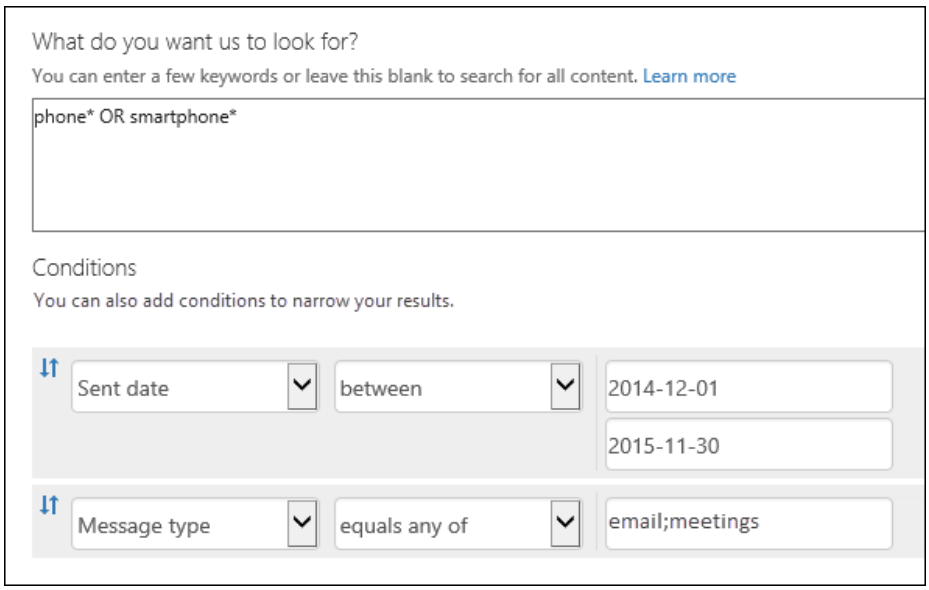

Search query syntax

```
phone* OR smartphone*(c:c)(sent=2016-12-01..2016-11-30)(kind="email")(kind="meetings")
```

Search query logic

```
phone* OR smartphone* AND (sent=2016-12-01..2016-11-30) AND ((kind="email") OR (kind="meetings"))
```

Special characters

Some special characters are not included in the search index and therefore are not searchable. This also includes the special characters that represent search operators in the search query. Here's a list of special characters that are either replaced by a blank space in the actual search query or cause a search error.

```
+ - = : ! @ # % ^ & ; _ / ? ( ) [ ] { }
```

Searching for site content shared with external users

You can also use the Content Search feature in the Security & Compliance Center to search for documents stored on SharePoint and OneDrive for Business sites that have been shared with people outside of your organization. This can help you identify sensitive or proprietary information that's being shared outside your organization. You can do this by using the `ViewableByExternalUsers` property in a keyword query. This property returns documents or sites that have been shared with external users by using one of the following sharing methods:

- A sharing invitation that requires users to sign in to your organization as an authenticated user.
- An anonymous guest link, which allows anyone with this link to access the resource without having to be authenticated.

Here are some examples:

- The query `ViewableByExternalUsers:true AND SensitiveType:"Credit Card Number"` returns all items that have been shared with people outside your organization and contain a credit card number.
- The query
  ```
  ViewableByExternalUsers:true AND ContentType:document AND
  site:"https://contoso.sharepoint.com/Sites/Teams"
  ```
  returns a list of documents on all team sites in the organization that have been shared with external users.

> **TIP**
> A search query such as `ViewableByExternalUsers:true AND ContentType:document` might return a lot of .aspx files in the search results. To eliminate these (or other types of files), you can use the `FileExtension` property to exclude specific file types; for example `ViewableByExternalUsers:true AND ContentType:document NOT FileExtension:aspx`.

What is considered content that is shared with people outside your organization? Documents in your organization's SharePoint and OneDrive for Business sites that are shared by sending a sharing invitation or that are shared in public locations. For example, the following user activities result in content that is viewable by external users:

- A user shares a file or folder with a person outside your organization.
- A user creates and sends a link to a shared file to a person outside your organization. This link allows the external user to view (or edit) the file.
- A user sends a sharing invitation or a guest link to a person outside your organization to view (or edit) a shared file.

Issues using the ViewableByExternalUsers property

While the `ViewableByExternalUsers` property represents the status of whether a document or site is shared with external users, there are some caveats to what this property does and doesn't reflect. In the following scenarios, the value of the `ViewableByExternalUsers` property won't be updated, and the results of a Content Search query that uses this property may be inaccurate.

- Changes to sharing policy, such as turning off external sharing for a site or for the organization. The property will still show previously shared documents as being externally accessible even though external access might have been revoked.
- Changes to group membership, such as adding or removing external users to Microsoft 365 Groups or Microsoft 365 security groups. The property won't automatically be updated for items the group has access to.
- Sending sharing invitations to external users where the recipient hasn't accepted the invitation, and therefore doesn't yet have access to the content.

In these scenarios, the `ViewableByExternalUsers` property won't reflect the current sharing status until the site or document library is recrawled and reindexed.

Searching for site content shared within your organization

As previously explained, you can use the `SharedWithUsersOWSUser` property so search for documents that have been shared between people in your organization. When a person shares a file (or folder) with another user inside your organization, a link to the shared file appears on the **Shared with me** page in the OneDrive for Business account of the person who the file was shared with. For example, to search for the documents that have been shared with Sara Davis, you can use the query `SharedWithUsersOWSUser:"sarad@contoso.com"`. If you export the results of this search, the original documents (located in the content location of the person who shared the documents with Sara) will be downloaded.

Documents must be explicitly shared with a specific user to be returned in search results when using the `SharedWithUsersOWSUser` property. For example, when a person shares a document in their OneDrive account, they have the option to share it with anyone (inside or outside the organization), share it only with people inside the organization, or share it with a specific person. Here's a screenshot of the **Share** window in OneDrive, that shows the three sharing options.

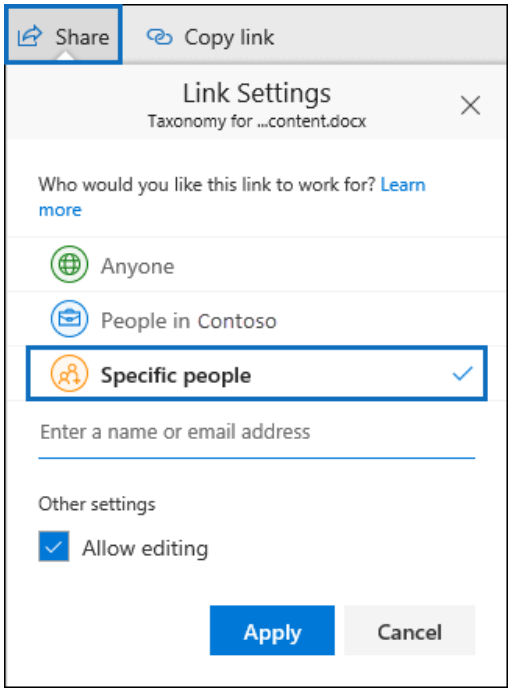

Only documents that are shared by using the third option (shared with **Specific people**) will be returned by a search query that uses the `SharedWithUsersOWSUser` property.

Searching for Skype for Business conversations

You can use the following keyword query to specifically search for content in Skype for Business conversations:

```
kind:im
```

The previous search query also returns chats from Microsoft Teams. To prevent this, you can narrow the search results to include only Skype for Business conversations by using the following keyword query:

```
kind:im AND subject:conversation
```

The previous keyword query excludes chats in Microsoft Teams because Skype for Business conversations are saved as email messages with a Subject line that starts with the word "Conversation".

To search for Skype for Business conversations that occurred within a specific date range, use the following keyword query:

```
kind:im AND subject:conversation AND (received=startdate..enddate)
```

Search tips and tricks

- Keyword searches are not case-sensitive. For example, **cat** and **CAT** return the same results.
- The Boolean operators **AND**, **OR**, **NOT**, and **NEAR** must be uppercase.
- A space between two keywords or two `property:value` expressions is the same as using **AND**. For example, `from:"Sara Davis" subject:reorganization` returns all messages sent by Sara Davis that contain the word reorganization in the subject line.
- Use syntax that matches the `property:value` format. Values are not case-sensitive, and they can't have a space after the operator. If there is a space, your intended value will be a full-text search. For example `to: pilarp` searches for "pilarp" as a keyword, rather than for messages that were sent to pilarp.
- When searching a recipient property, such as To, From, Cc, or Recipients, you can use an SMTP address, alias, or display name to denote a recipient. For example, you can use pilarp@contoso.com, pilarp, or "Pilar Pinilla".
- You can use only prefix wildcard searches; for example, **cat*** or **set***. Suffix searches (***cat**), infix searches (**c*t**), and substring searches (***cat***) are not supported.
- When searching a property, use double quotation marks (" ") if the search value consists of multiple words. For example `subject:budget Q1` returns messages that contain **budget** in the subject line and that contain **Q1** anywhere in the message or in any of the message properties. Using `subject:"budget Q1"` returns all messages that contain **budget Q1** anywhere in the subject line.
- To exclude content marked with a certain property value from your search results, place a minus sign (-) before the name of the property. For example, `-from:"Sara Davis"` excludes any messages sent by Sara Davis.
- You can export items based on message type. For example, to export Skype conversations and chats in Microsoft Teams, use the syntax `kind:im`. To return only email messages, you would use `kind:email`. To return chats, meetings, and calls in Microsoft Teams, use `kind:microsoftteams`.

View keyword statistics for Content Search results

2/18/2021 • 6 minutes to read • Edit Online

After you create and run a Content Search, you can view statistics about the estimated search results. This includes a summary of the search results (similar to the summary of the estimated search results displayed in the details pane), the query statistics such as the number of content locations with items that match the search query, and the name of content locations that have the most matching items. You can display statistics for one or more content searches. This lets you to quickly compare the results for multiple searches and make decisions about the effectiveness of your search queries.

Additionally, you can configure new and existing searches to return statistics for each keyword in a search query. This lets you compare the number of results for each keyword in a query and to compare the keyword statistics from multiple searches.

You can also download the search statistics and keyword statistics to a CSV file. This lets you use the filtering and sorting features in Excel to compare results, and prepare reports for your search results.

Get statistics for Content Searches

To display statistics for Content searches:

1. In the Microsoft 365 compliance center, go to **Show all** > **Content search**.
2. In the list of searches, select two or more searches, and then click **Search statistics** on the **Bulk actions** flyout page.

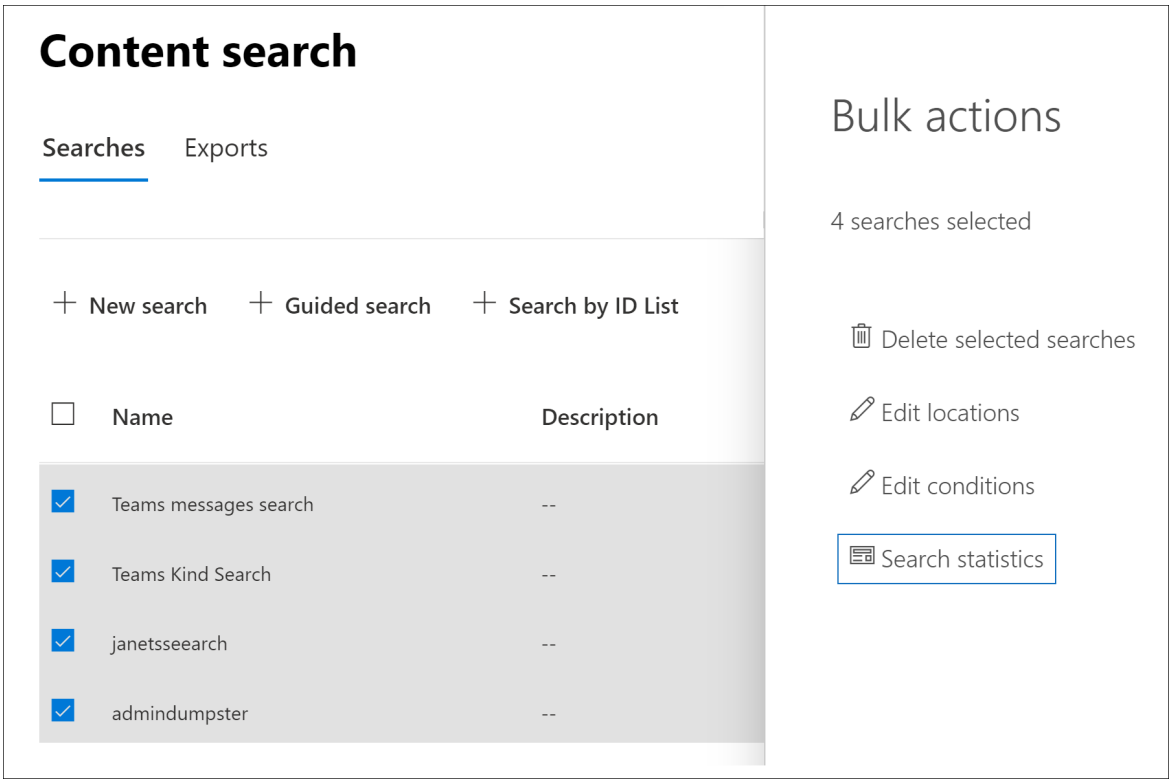

3. On the **Search statistics** page, click one of the following links to display statistics about the selected searches.

   Summary

   This page displays statistics similar to the ones displayed in the details pane on the **Content search**

page. Statistics for all selected searches are displayed. Note that you can also re-run the selected searches from this page to update the statistics.

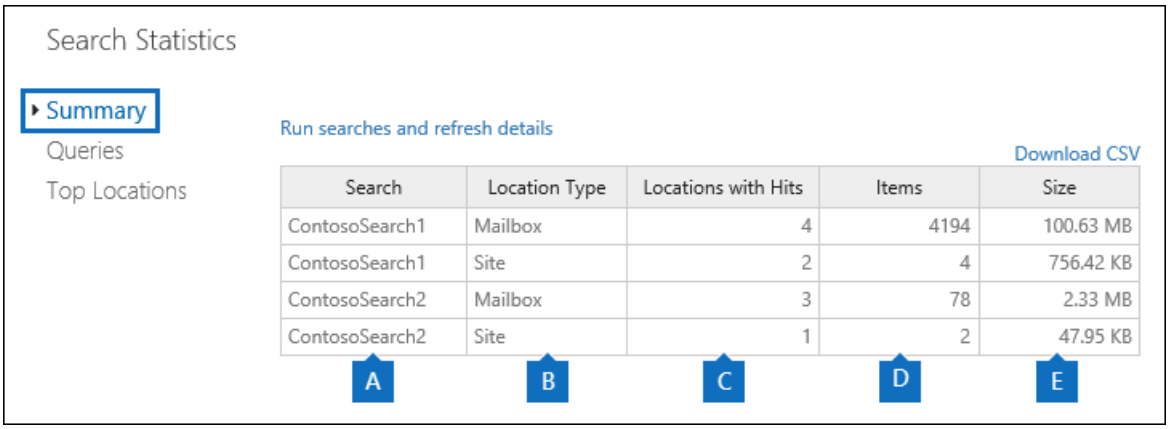

a. The name of the Content Search. As previously stated, you can display and compare statistics for multiple searches.

b. The type of content location that was searched. Each row displays statistics for mailboxes, sites, and public folders from the specified search.

c. The number of content locations containing items that match the search query. For mailboxes, this statistic also includes the number of archive mailboxes that contain items that match the search query.

d. The total number of items of all specified content locations that match the search query. Examples of item types include email messages, calendar items, and documents. If an item contains multiple instances of a keyword that is being searched for, it's only counted once in the total number of items. For example, if you're searching for words "stock" or "fraud" and an email message contains three instances of the word "stock", it's only counted once in the **Items** column.

e. The total size of all items that were found in the specified content location that match the search query.

Queries

This page displays statistics about the search query.

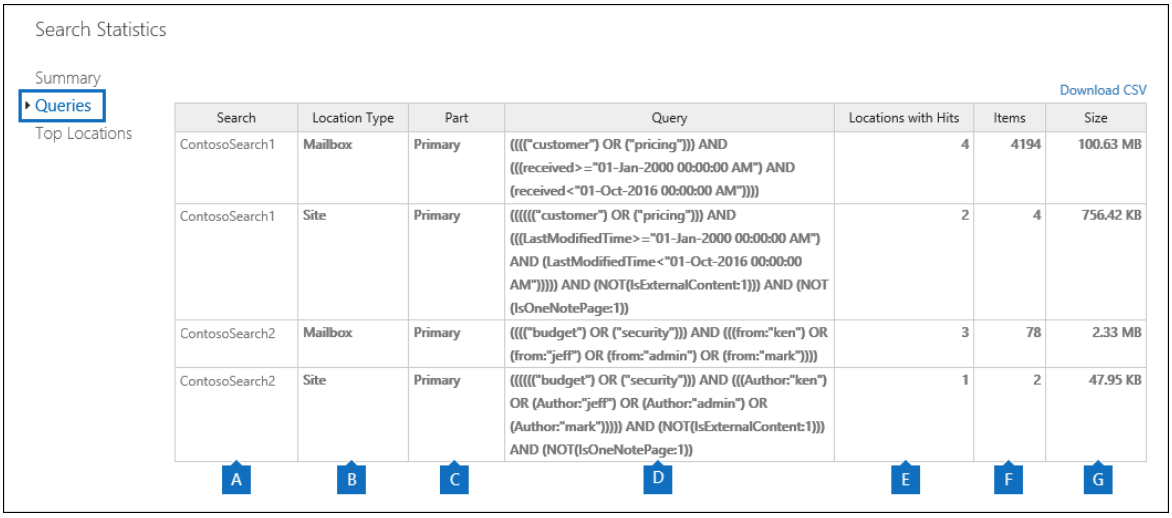

a. The name of the Content Search that the row contains query statistics for.

b. The type of content location that the query statistics are applicable to.

c. This column indicates which part of the search query the statistics are applicable to. **Primary** indicates the entire search query. If you use a keyword list when you create or edit a search query, statistics for each component of the query are included in this table. See the Get keyword statistics for Content

Searches section in this article for more information.

d. This column contains the actual search query that run by the Content Search tool. Note that the tool automatically adds a few additional components to the query that you create.

- When you search for all content in mailboxes (by not specifying any keywords), the actual key word query is `size>=0` so that all items are returned.
- When you search SharePoint Online and OneDrive for Business sites, the two following components are added:

  NOT IsExternalContent:1 - Excludes any content from an on-premises SharePoint organization.

  NOT IsOneNotePage:1 - Excludes all OneNote files because these would be duplicates of any document that matches the search query.

e. The number of the content locations (specified by the ** Location type ** column) that contain items that match the search query listed in the **Query** column.

f. The number of items (from the specified content location) that match the search query listed in the **Query** column. As previously explained, if an item contains multiple instances of a keyword that is being searched for, it's only counted once in the this column.

g. The total size of all items that were found (in the specified content location) that match the search query in the **Query** column.

Top locations

This page displays statistics about the number of items that match the search query in each content location that was searched. The top 1,000 locations are displayed. If you view statistics for multiple searches, the top 1,000 locations for each search are displayed. Note that a content location isn't included on this page if it doesn't contain any items that match the search query.

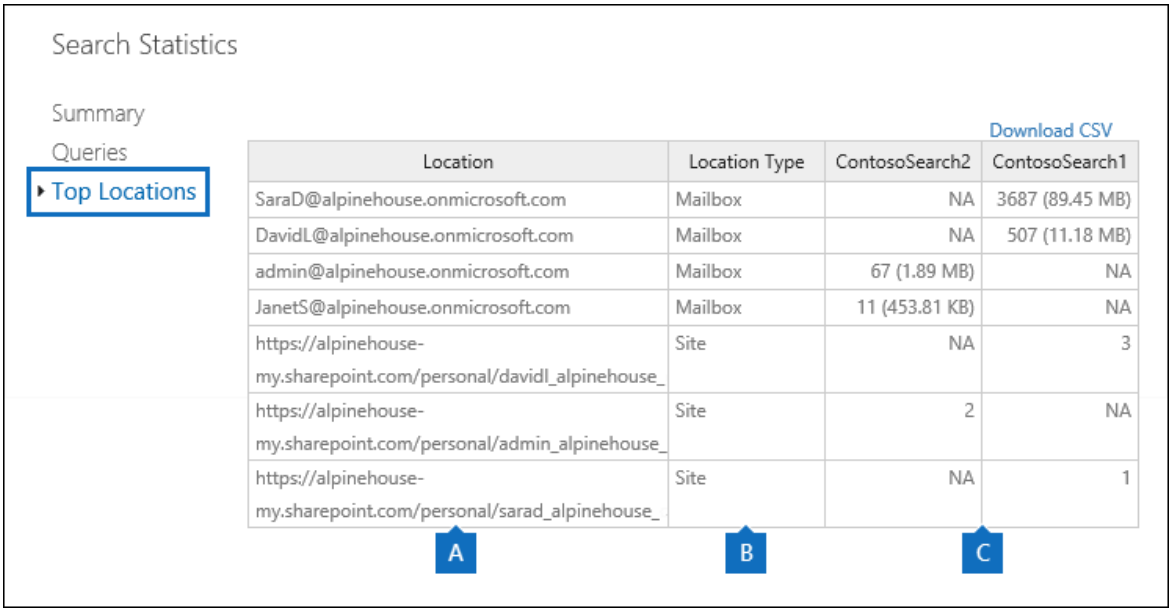

a. The name of the content location.

b. The type of content location that the location statistics are applicable to.

c. There are columns for each search that you're displaying statistics for. This column shows the number (and total size) of items that match the search query in each content location. Note that when you're displaying statistics for multiple searches, the "NA" in this column indicates that the content location wasn't included in that search.

Get keyword statistics for Content Searches

As previous explained, the **Queries** page shows the search query and the number (and size) of items that match the query. If you use a keyword list when you create or edit a search query, you can get enhanced statistics that show how many items match each keyword or keyword phrase. This can help you quickly identify which parts of the query are the most (and least) effective. For example, if a keyword returns a large number of items, you might choose to refine the keyword query to narrow the search results. You can set up a keyword list when you create or edit a Content Search.

To create a keyword list and view keyword statistics for a Content Search:

1. In the Microsoft 365 compliance center, go to **Show all** > **Content search**.
2. In the list of content searches, click and a search, and then click **Edit** ✎.
3. Click **Query** and then do the following things:

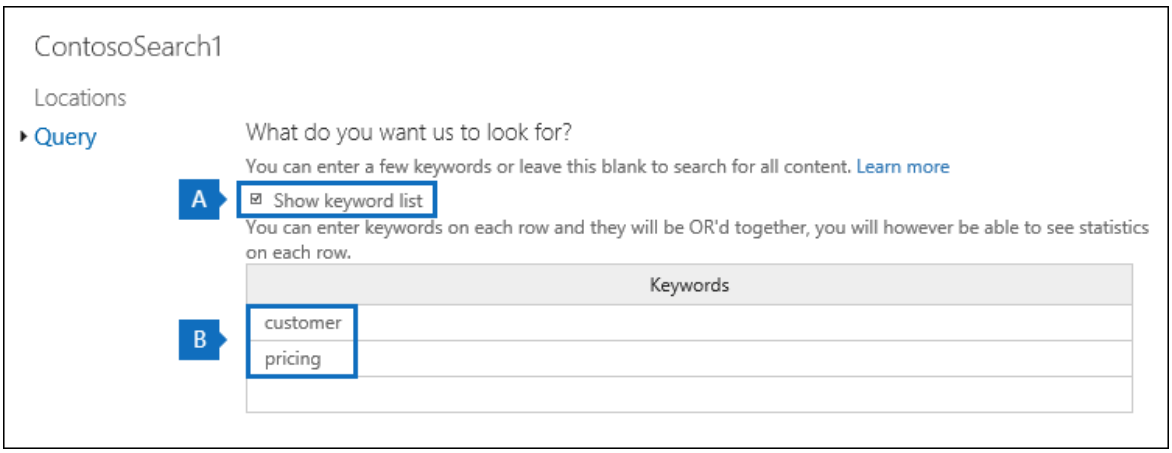

   a. Click the **Show keyword list** check box.

   b. Type a keyword or keyword phase in a row in the keywords table. For example, type **budget** in the first row and then type **security** in the second row.

4. After adding the keywords that you want to search and get statistics for, click **Search** to run the revised search.
5. When the search is completed, select it in the list of searches, and then click **Search statistics** 🖶. You can also display and compare keyword statistics for multiple searches.
6. On the **Search statistics** page, click **Query** to display the keyword statistics for the selected searches.

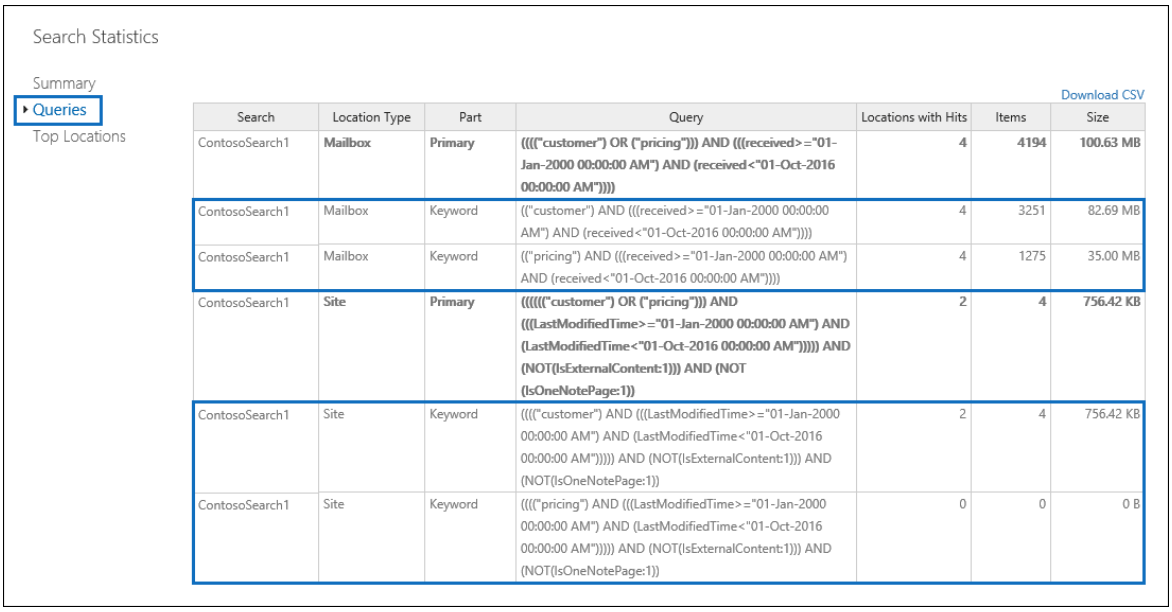

Search	Location Type	Part	Query	Locations with Hits	Items	Size
ContosoSearch1	**Mailbox**	**Primary**	**((("customer") OR ("pricing"))) AND (((received>="01-Jan-2000 00:00:00 AM") AND (received<"01-Oct-2016 00:00:00 AM")))**	**4**	**4194**	**100.63 MB**
ContosoSearch1	Mailbox	Keyword	(("customer") AND (((received>="01-Jan-2000 00:00:00 AM") AND (received<"01-Oct-2016 00:00:00 AM"))))	4	3251	82.69 MB
ContosoSearch1	Mailbox	Keyword	(("pricing") AND (((received>="01-Jan-2000 00:00:00 AM") AND (received<"01-Oct-2016 00:00:00 AM"))))	4	1275	35.00 MB
ContosoSearch1	**Site**	**Primary**	**((((("customer") OR ("pricing"))) AND (((LastModifiedTime>="01-Jan-2000 00:00:00 AM") AND (LastModifiedTime<"01-Oct-2016 00:00:00 AM"))))) AND (NOT(IsExternalContent:1))) AND (NOT (IsOneNotePage:1))**	**2**	**4**	**756.42 KB**
ContosoSearch1	Site	Keyword	((((("customer") AND (((LastModifiedTime>="01-Jan-2000 00:00:00 AM") AND (LastModifiedTime<"01-Oct-2016 00:00:00 AM"))))) AND (NOT(IsExternalContent:1))) AND (NOT(IsOneNotePage:1))	2	4	756.42 KB
ContosoSearch1	Site	Keyword	((((("pricing") AND (((LastModifiedTime>="01-Jan-2000 00:00:00 AM") AND (LastModifiedTime<"01-Oct-2016 00:00:00 AM"))))) AND (NOT(IsExternalContent:1))) AND (NOT(IsOneNotePage:1))	0	0	0 B

As shown in the previous screenshot, the statistics for each keyword are displayed; this includes:

- The keyword statistics for each type of content location included in the search.
- The actual search query for each keyword, which includes any conditions from the search query.
- The complete search query (identified as **Primary** in the **Part** column) and the statistics for the complete query. Note these are the same statistics displayed on the **Summary** page.

> **NOTE**
>
> To help reduce issues caused by large keyword lists, you're now limited to a maximum of 20 rows in the keyword list of a search query.

Export Content Search results

2/18/2021 • 22 minutes to read • Edit Online

After a Content Search is successfully run, you can export the search results to a local computer. When you export email results, they're downloaded to your computer as PST files. When you export content from SharePoint and OneDrive for Business sites, copies of native Office documents are exported. There are other documents and reports included with the exported search results.

Exporting the results of a Content Search involves preparing the results, and then downloading them to a local computer.

Before you export content search results

- To export search results, you have to be assigned the Export management role in the Security & Compliance Center. This role is assigned to the built-in eDiscovery Manager role group. It isn't assigned by default to the Organization Management role group. For more information, see Assign eDiscovery permissions.
- The computer you use to export the search results has to meet the following system requirements:
  - 32-bit or 64-bit versions of Windows 7 and later versions
  - Microsoft .NET Framework 4.7
- You have to use one of the following supported browsers to run the eDiscovery Export Tool[1]:
  - Microsoft Edge [2]

    OR
  - Microsoft Internet Explorer 10 and later versions

> **NOTE**
>
> [1] Microsoft doesn't manufacture third-party extensions or add-ons for ClickOnce applications. Exporting search results using an unsupported browser with third-party extensions or add-ons isn't supported.
>
> [2] As a result of recent changes to Microsoft Edge, ClickOnce support is no longer enabled by default. For instructions on enabling ClickOnce support in Edge, see Use the eDiscovery Export Tool in Microsoft Edge.

- We recommend downloading search results to a local computer. However, to eliminate your company's firewall or proxy infrastructure from causing issues when downloading search results, you might consider downloading search results to a virtual desktop outside of your network. This may decrease timeouts that occur in Azure data connections when exporting a large number of files. For more information about virtual desktops, see Windows Virtual Desktop.
- To improve performance when downloading search results, consider dividing searches that return a large set of results into smaller searches. For example, you can use date ranges in search queries to return a smaller set of results that can be downloaded faster.
- When you export search results, the data is temporarily stored in a Microsoft-provided Azure Storage location in the Microsoft cloud before it's downloaded to your local computer. Be sure that your organization can connect to the endpoint in Azure, which is ***.blob.core.windows.net** (the wildcard represents a unique identifier for your export). The search results data is deleted from the Azure Storage

location two weeks after it's created.

- If your organization uses a proxy server to communicate with the Internet, you need to define the proxy server settings on the computer that you use to export the search results (so the export tool can be authenticated by your proxy server). To do this, open the *machine.config* file in the location that matches your version of Windows.

  - 32-bit: `%windir%\Microsoft.NET\Framework\[version]\Config\machine.config`
  - 64-bit: `%windir%\Microsoft.NET\Framework64\[version]\Config\machine.config`

  Add the following lines to the *machine.config* file somewhere between the `<configuration>` and `</configuration>` tags. Be sure to replace `ProxyServer` and `Port` with the correct values for your organization; for example, `proxy01.contoso.com:80` .

```
<system.net>
   <defaultProxy enabled="true" useDefaultCredentials="true">
     <proxy proxyaddress="https://ProxyServer :Port "
            usesystemdefault="False"
            bypassonlocal="True"
            autoDetect="False" />
   </defaultProxy>
</system.net>
```

Step 1: Prepare search results for export

The first step is to prepare the search results for exporting. When you prepare results, they are uploaded to a Microsoft-provided Azure Storage location in the Microsoft cloud. Content from mailboxes and sites is uploaded at a maximum rate of 2 GB per hour.

1. Go to https://protection.office.com.
2. Sign in using your work or school account.
3. In the left pane of the Security & Compliance Center, click **Search** > **Content search**.
4. On the **Content search** page, select a search.
5. In the details pane, under **Export results to a computer**, click **Start export**.

> **NOTE**
> If the results for a search are older than 7 days, you are prompted to update the search results. If this happens, cancel the export, click **Update search results** in the details pane for the selected search, and then start the export again after the results are updated.

6. On the **Export the search results** page, under **Output options**, choose one of the following 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

   See the More information section for a description about how partially indexed items are exported. For more information about partially indexed items, see Partially indexed items in Content Search.

7. Under **Export Exchange content as**, choose one of the following options:

   - **One PST file for each mailbox:** Exports one PST file for each user mailbox that contains search results. Any results from the user's archive mailbox are included in the same PST file. This option reproduces the mailbox folder structure from the source mailbox.
   - **One PST file containing all messages:** Exports a single PST file (named *Exchange.pst*) that contains the search results from all source mailboxes included in the search. This option reproduces the mailbox folder structure for each message.
   - **One PST file containing all messages in a single folder:** Exports search results to a single PST file where all messages are located in a single, top-level folder. This option lets reviewers review items in chronological order (items are sorted by sent date) without having to navigate the original mailbox folder structure for each item.
   - **Individual messages:** Exports search results as individual email messages, using the .msg format. If you select this option, email search results are exported to a folder in the file system. The folder path for individual messages is the same as the one used if you exported the results to PST files.

     > **IMPORTANT**
     > To decrypt RMS-protected messages when they're exported, you must export email search results as individual messages. Encrypted messages will remain encrypted if you export the search results as a PST file. For more information, see Decrypting RMS-protected email messages and encrypted file attachments in this article.

8. Click the **Enable de-duplication** checkbox to exclude duplicate messages. This option appears only if the content sources of the search include Exchange mailboxes or public folders.

   If you select this option, only one copy of a message will be exported even if multiple copies of the same message are found in the mailboxes that were searched. The export results report (Results.csv) will contain a row for every copy of a duplicate message so that you can identify the mailboxes (or public folders) that contain a copy of the duplicate message. For more information about de-duplication and how duplicate items are identified, see De-duplication in eDiscovery search results.

9. Click the **Include versions for SharePoint documents** checkbox to export all versions of SharePoint documents. This option appears only if the content sources of the search include SharePoint or OneDrive for Business sites.

10. Click the **Export files in a compressed (zipped) folder** checkbox to export search results to compressed folders. This option is available only when you choose to export Exchange items as individual messages and when the search results include SharePoint or OneDrive documents. This option is primarily used to work around the 260 character limit in Windows file path names when items are exported. See the "Filenames of exported items" in the More information section.

11. Click **Start export**. The search results are prepared for downloading, which means they're being uploaded to an Azure Storage location in the Microsoft cloud. This may take several minutes.

See the next section for instructions to download the exported search results.

Step 2: Download the search results

The next step is to download the search results from the Azure Storage location to your local computer.

1. On the **Content search** page, click the **Exports** tab.

You may have to click **Refresh** to update the list of export jobs so that it shows the export job you created. Export jobs have the same name as the corresponding search with **_Export** appended to the search name.

2. Select the export job that you created in Step 1.
3. On the flyout page under **Export key**, click **Copy to clipboard**. You use this key in step 6 to download the search results.
4. Click **Download results**.
5. If you're prompted to install the **eDiscovery Export Tool**, click **Install**.
6. In the **eDiscovery Export Tool**, do the following:

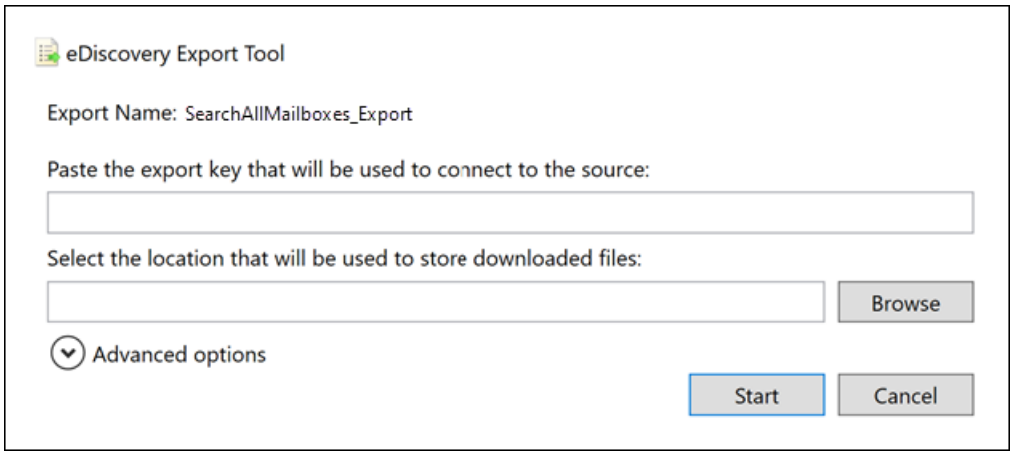

   a. Paste the export key that you copied in step 3 in the appropriate box.

   b. Click **Browse** to specify the location where you want to download the search result files.

   > **IMPORTANT**
   > Due to high network activity during download, you should download search results only to a location on an internal drive on your local computer. For the best download experience, follow these guidelines:
   >
   > - Don't download search results to a UNC path, a mapped network drive, an external USB drive, or a synched OneDrive for Business account.
   > - Disable anti-virus scanning for the folder that you download the search result to.
   > - Download search results to different folders for concurrent download jobs.

7. Click **Start** to download the search results to your computer.

   The **eDiscovery Export Tool** displays status information about the export process, including an estimate of the number (and size) of the remaining items to be downloaded. When the export process is complete, you can access the files in the location where they were downloaded.

More information

Here's more information about exporting search results.

Export limits

Export reports

Exporting partially indexed items

Exporting individual messages or PST files

Exporting results from more than 100,000 mailboxes

Decrypting RMS-protected email messages and encrypted file attachments

Filenames of exported items

Miscellaneous

Export limits

For information about limits when exporting content search results, see the "Export limits" section in Limits for content search.

Export reports

- When you export search results, the following reports are included in addition to the search results.
    - **Export Summary** An Excel document that contains a summary of the export. This includes information such as the number of content sources that were searched, the estimated and downloaded sizes of the search results, and the estimated and downloaded number of items that were exported.
    - **Manifest** A manifest file (in XML format) that contains information about each item included in the search results.
    - **Results** An Excel document that contains information about each item that is download as a search result. For email, the result log contains information about each message, including:
        - The location of the message in the source mailbox (including whether the message is in the primary or archive mailbox).
        - The date the message was sent or received.
        - The Subject line from the message.
        - The sender and recipients of the message.
        - Whether the message is a duplicate message if you enabled the de-duplication option when exporting the search results. Duplicate messages have a value in the **Duplicate to Item** column that identifies the message as a duplicate. The value in the **Duplicate to Item** column contains the item identity of the message that was exported. For more information, see De-duplication in eDiscovery search results.

          For documents from SharePoint and OneDrive for Business sites, the result log contains information about each document, including:
            - The URL for the document.
            - The URL for the site collection where the document is located.
            - The date that the document was last modified.
            - The name of the document (which is located in the Subject column in the result log).
    - **Unindexed Items** An Excel document that contains information about any partially indexed items that would be included in the search results. If you don't include partially indexed items when you generate the search results report, this report will still be downloaded, but will be empty.
    - **Errors and Warnings** Contains errors and warnings for files encountered during export. See the Error Details column for information specific to each individual error or warning.
    - **Skipped Items** When you export search results from SharePoint and OneDrive for Business sites, the export will usually include a skipped items report (SkippedItems.csv). The items cited in this

report are typically items that won't be downloaded, such as a folder or a document set. Not exporting these types of items is by design. For other items that were skipped, the 'Error Type' and 'Error Details' field in the skipped items report show the reason the item was skipped and wasn't downloaded with the other search results.

- **Trace Log** Contains detailed logging information about the export process and can help uncover issues during export.

> **NOTE**
>
> You can just export these documents without having to export the actual search results. See Export a Content Search report.

Exporting partially indexed items

- If you're exporting mailbox items from a content search that returns all mailbox items in the search results (because no keywords where included in the search query), partially indexed items won't be copied to the PST file that contains the unindexed items. This is because all items, including any partially indexed items, are automatically included in the regular search results. This means that partially indexed items will be included in a PST file (or as individual messages) that contains the other, indexed items.

  If you export both the indexed and partially indexed items or if you export only the indexed items from a content search that returns all items, the same number of items will be downloaded. This happens even though the estimated search results for the content search (displayed in the search statistics in the Security & Compliance Center) will still include a separate estimate for the number of partially indexed items. For example, let's say that the estimate for a search that includes all items (no keywords in the search query) shows that 1,000 items were found and that 200 partially indexed items were also found. In this case, the 1,000 items include the partially indexed items because the search returns all items. In other words, there are 1,000 total items returned by the search, and not 1,200 items (as you might expect). If you export the results of this search and choose to export indexed and partially indexed items (or export only partially indexed items), then 1,000 items will be downloaded. Again, that's because partially indexed items are included with the regular (indexed) results when you use a blank search query to return all items. In this same example, if you choose to export only partially indexed items, then only the 200 unindexed items would be downloaded.

  Also note that in the previous example (when you export indexed and partially indexed items or you export only indexed items), the **Export Summary** report included with the exported search results would list 1,000 items estimated items and 1,000 downloaded items for the same reasons as previously described.

- If the search that you're exporting results from was a search of specific content locations or all content locations in your organization, only the partial items from content locations that contain items that match the search criteria will be exported. In other words, if no search results are found in a mailbox or site, then any partially indexed items in that mailbox or site won't be exported. The reason for this is that exporting partially indexed items from lots of locations in the organization might increase the likelihood of export errors and increase the time it takes to export and download the search results.

  To export partially indexed items from all content locations for a search, configure the search to return all items (by removing any keywords from the search query) and then export only partially indexed items when you export the search results.

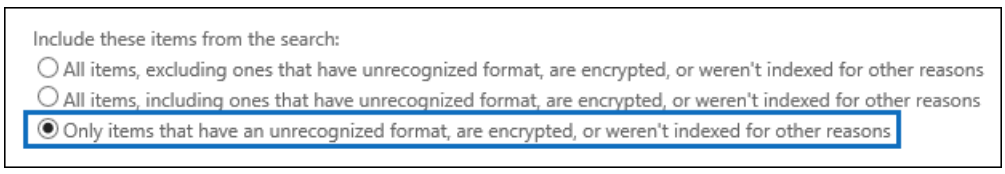

- When exporting search results from SharePoint or OneDrive for Business sites, the ability to export unindexed items also depends on the export option that you select and whether a site that was searched contains an indexed item that matches the search criteria. For example, if you search specific SharePoint or OneDrive for Business sites and no search results are found, then no unindexed items from those sites will be exported if you choose the second export option to export both indexed and unindexed items. If an indexed item from a site does match the search criteria, then all unindexed items from that site will be exported when exporting both indexed and unindexed items. The following illustration describes the export options based on whether a site contains an indexed item that matches the search criteria.

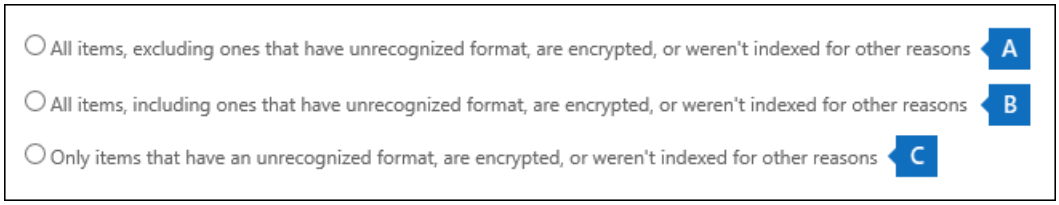

  1. Only indexed items that match the search criteria are exported. No partially indexed items are exported.
  2. If no indexed items from a site match the search criteria, then partially indexed items from that same site aren't exported. If indexed items from a site are returned in the search results, then the partially indexed items from that site are exported. In other words, only the partially indexed items from sites that contain items that match the search criteria are exported.
  3. All partially indexed items from all sites in the search are exported, regardless of whether a site contains items that match the search criteria.

  If you choose to export partially indexed items, partially indexed mailbox items are exported in a separate PST file regardless of the option that you choose under **Export Exchange content as**.

- If partially indexed items are returned in the search results (because other properties of partially indexed items matched the search criteria), then those partially indexed are exported with the regular search results. So, if you choose to export both indexed items and partially indexed items (by selecting the **All items, including ones that have unrecognized format, are encrypted, or weren't indexed for other reasons** export option), the partially indexed items exported with the regular results will be listed in the Results.csv report. They will not be listed in the Unindexed items.csv report.

Exporting individual messages or PST files

- If the file path name of a message exceeds the maximum character limit for Windows, the file path name is truncated. But the original file path name will be listed in the Manifest and ResultsLog.
- As previously explained, email search results are exported to a folder in the file system. The folder path for individual messages would replicate the folder path in the user's mailbox. For example, for a search named "ContosoCase101" messages in a user's inbox would be located in the folder path `~ContosoCase101\\<date of export\Exchange\user@contoso.com (Primary)\Top of Information Store\Inbox`.
- If you choose to export email messages in one PST file containing all messages in a single folder, a **Deleted Items** folder and a **Search Folders** folder are included in the top level of the PST folder. These folders are empty.
- As previously stated, you must export email search results as individual messages to decrypt RMS-protected messages when they're exported. Encrypted messages will remain encrypted if you export email search results as a PST file.

Exporting results from more than 100,000 mailboxes

- As previously explained, you have to use Security & Compliance Center PowerShell to download the search results from more than 100,000 mailboxes. You can run the following script in this section to

download these search results. Using this script assumes that you have already exported the search results (the export job is displayed on the **Exports** tab in the Content Search tool) and now want to download them.

```
$export=Get-ComplianceSearchAction SEARCHNAME_Export -IncludeCredential;
$exportUrl=   [System.Uri]::EscapeDataString(($export.Results.Split(";") | ?{$_ -like '*Container
url*'} | %{$_.Split(":",2)} | select -last 1).Trim());
$exportToken=($export.Results.Split(";") | ?{$_ -like '*SAS Token*'} | %{$_.Split(":",2)} | select -
last 1).Trim();
."$env:ProgramFiles\Internet Explorer\IEXPLORE.EXE"
"https://complianceclientsdf.blob.core.windows.net/v16/Microsoft.Office.Client.Discovery.UnifiedExpor
tTool.application?name=$($export.Name)&source=$exportUrl&zip=allow&trace=1";
$exportToken | clip;
```

In the script, you have to specify the name of the search that you want to export results for. For example, for a search named, `SearchAllMailboxes` replace SEARCHNAME_Export with `SearchAllMailboxes_Export`.

After you add the name of the search to the script, you can copy the script text and then paste it into a Windows PowerShell window that's connected to Security & Compliance Center PowerShell. After you paste the script, the eDiscovery Export Tool is displayed (like it is when you download search results using the UI):

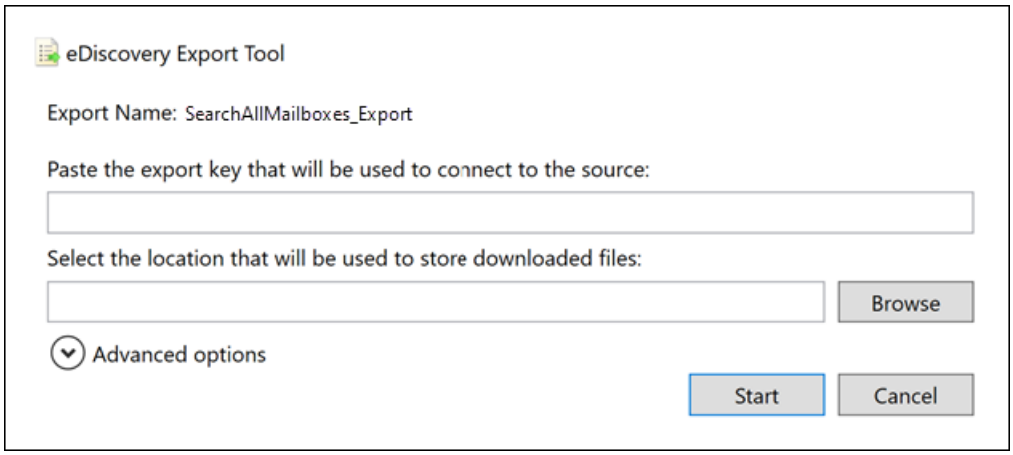

Click in the export key box and then press `CTRL + V` to paste the export key (the script copies the export key to the clipboard). Click **Browse** to specify the location where you want to download the files, and then start the download.

As previously stated, we recommend that you download search results to a local disk drive due to the high amount of disk activity (reads and writes). Don't download search results to a mapped network drive or other network location.

Decrypting RMS-protected email messages and encrypted file attachments

Any rights-protected (RMS-protected) email messages included in the results of a Content Search will be decrypted when you export them. Additionally, any file that's encrypted with a Microsoft encryption technology and is attached to an email message that's included in the search results will also be decrypted when it's exported. This decryption capability is enabled by default for members of the eDiscovery Manager role group. This is because the RMS Decrypt management role is assigned to this role group by default. Keep the following things in mind when exporting encrypted email messages and attachments:

- As previously explained, to decrypt RMS-protected messages when you export them, you have to export the search results as individual messages. If you export search results to a PST file, RMS-protected messages remain encrypted.
- Messages that are decrypted are identified in the **ResultsLog** report. This report contains a column named **Decode Status**, and a value of **Decoded** in this column identifies the messages that were

decrypted.

- In addition to decrypting file attachments when exporting search results, you can also preview the decrypted file when previewing search results. You can only view the rights-protected email message after you export it.
- At this time, the decryption capability when exporting search results doesn't include encrypted content from SharePoint and OneDrive for Business sites. However, support is coming soon for documents encrypted with Microsoft encryption technologies and stored in SharePoint Online and OneDrive for Business.
- If you need to prevent someone from decrypting RMS-protect messages and encrypted file attachments, you have to create a custom role group (by copying the built-in eDiscovery Manager role group) and then remove the RMS Decrypt management role from the custom role group. Then add the person who you don't want to decrypt messages as a member of the custom role group.

Filenames of exported items

- There is a 260-character limit (imposed by the operating system) for the full path name for email messages and site documents exported to your local computer. The full path name for exported items includes the item's original location and the folder location on the local computer where the search results are downloaded to. For example, if you specify to download the search results to `C:\Users\Admin\Desktop\SearchResults` in the eDiscovery Export tool, then the full pathname for a downloaded email item would be `C:\Users\Admin\Desktop\SearchResults\ContentSearch1\03.15.2017-1242PM\Exchange\sarad@contoso.com (Primary)\Top of Information Store\Inbox\Insider trading investigation.msg`.

  If the 260-character limit is exceeded, the full path name for an item will be truncated.

  - If the full path name is longer than 260 characters, the file name will be shortened to get under the limit; note that the truncated filename (excluding the file extension) won't be fewer than eight characters.
  - If the full path name is still too long after shortening the file name, the item is moved from its current location to the parent folder. If the pathname is still too long, then the process is repeated: shorten the filename, and if necessary move again to the parent folder. This process is repeated until the full pathname is under the 260-character limit.
  - If a truncated full path name already exists, a version number is added to the end of the filename; for example, `statusmessage(2).msg`.

    To help mitigate this issue, consider downloading search results to a location with a short path name; for example, downloading search results to a folder named `C:\Results` would add fewer characters to the path names of exported items than downloading them to a folder named `C:\Users\Admin\Desktop\Results`.

- When you export site documents, it's also possible that the original file name of a document will be modified. This happens specifically for documents that have been deleted from a SharePoint or OneDrive for Business site that's been placed on hold. After a document that's on a site that's on hold is deleted, the deleted document is automatically moved to the Preservation Hold library for the site (which was created when the site was placed on hold). When the deleted document is moved to the Preservation Hold library, a randomly generated and unique ID is appended to the original filename of the document. For example, if the filename for a document is `FY2017Budget.xlsx` and that document is later deleted and moved to the Preservation Hold library, the filename of the document that is moved to the Preservation Hold library is modified to something like `FY2017Budget_DEAF727D-0478-4A7F-87DE-5487F033C81A2000-07-05T10-37-55.xlsx`. If a document in the

Preservation Hold library matches the query of a Content Search and you export the results of that search, the exported file has the modified filename; in this example, the filename of the exported document would be `FY2017Budget_DEAF727D-0478-4A7F-87DE-5487F033C81A2000-07-05T10-37-55.xlsx`.

When a document on a site that's on hold is modified (and versioning for the document library in the site has been enabled), a copy of the file is automatically created in the Preservation Hold library. In this case, a randomly generated and unique ID is also appended to the filename of the document that's copied to the Preservation Hold library.

The reason why filenames of documents that are moved or copied to the Preservation Hold library is to prevent conflicting filenames. For more information about placing a hold on sites and the Preservation Hold library, see Overview of in-place hold in SharePoint Server 2016.

Miscellaneous

- When downloading search results using the eDiscovery Export Tool, it's possible you might receive the following error:

  ```
  System.Net.WebException: The remote server returned an error: (412) The condition specified using HTTP conditional header(s) is not met.
  ```

  This is transient error, which typically occurs in the Azure Storage location. To resolve this issue, retry downloading the search results, which will restart the eDiscovery Export Tool.

- All search results and the export reports are included in a folder that has the same name as the Content Search. The email messages that were exported are located in a folder named **Exchange**. Documents are located in a folder named **SharePoint**.

- The file system metadata for documents on SharePoint and OneDrive for Business sites is maintained when documents are exported to your local computer. That means document properties, such as created and last modified dates, aren't changed when documents are exported.

- If your search results include a list item from SharePoint that matches the search query, all rows in the list will be exported in addition to the item that matches the search query and any attachments in the list. The reason for this behavior is to provide a context for list items that are returned in the search results. Also note that the additional list items and attachments may cause the count of exported items to be different than the original estimate of search results.

Export a Content Search report

11/2/2020 • 7 minutes to read • Edit Online

Instead of exporting the full set of search results from a Content Search in the Security & Compliance Center (and from a Content Search that's associated with an eDiscovery case), you can export the same reports that are generated when you export search results.

When you export a report, it's downloaded to a folder that has the same name as the Content Search, but that's appended with *_ReportsOnly*. For example, if the Content Search is named *ContosoCase0815*, then the report is downloaded to a folder named *ContosoCase0815_ReportsOnly*. For a list of documents that are included in the report, see What's included in the report.

Assign roles and check system requirements

- To export a Content Search report, you have to be assigned the Compliance Search management role in the Security & Compliance Center. This role is assigned by default to the built-in eDiscovery Manager and Organization Management role groups. For more information, see Assign eDiscovery permissions.
- When you export a report, the data is temporarily stored in a unique Azure Storage area in the Microsoft cloud before it's downloaded to your local computer. Be sure that your organization can connect to the endpoint in Azure, which is ***.blob.core.windows.net** (the wildcard represents a unique identifier for your export). The search results data is deleted from the Azure Storage area two weeks after it's created.
- The computer you use to export the search results has to meet the following system requirements:
  - 32-bit or 64-bit versions of Windows 7 and later versions
  - Microsoft .NET Framework 4.7
- You have to use one of the following supported browsers to run the eDiscovery Export Tool[1]:
  - Microsoft Edge [2]

    OR
  - Microsoft Internet Explorer 10 and later versions

> **NOTE**
>
> [1] Microsoft doesn't manufacture third-party extensions or add-ons for ClickOnce applications. Exporting search results using an unsupported browser with third-party extensions or add-ons isn't supported.
>
> [2] As a result of recent changes to Microsoft Edge, ClickOnce support is no longer enabled by default. For instructions on enabling ClickOnce support in Edge, see Use the eDiscovery Export Tool in Microsoft Edge.

- If the estimated total size of the results returned by a Content Search exceeds 2 TB, exporting the report fails. To successfully export the report, try to narrow the scope and rerun the search so the estimated size of the results is less than 2 TB.
- Exporting Content Search reports counts against the maximum number of exports running at the same time and the maximum number of exports that a single user can run. For more information about export limits, see Export Content Search results.

Generate and download a Content Search report

The steps to generate and download a Content Search report are similar to actually exporting search results.

Step 1: Generate the report for export

The first step is to prepare the report for downloading to your computer exporting. When you the report, the report documents are uploaded to an Azure Storage area in the Microsoft cloud.

1. Go to https://protection.office.com.
2. Sign in using your work or school account.
3. In the left pane of the Security & Compliance Center, click **Search** > **Content search**.
4. On the **Content search** page, select a search.
5. In the details pane, under **Export report to a computer**, click **Generate report**.

   > **NOTE**
   > If the results for a search are older than 7 days, you are prompted to update the search results. If this happens, cancel the export, click **Update search results** in the details pane for the selected search, and then start the report export again after the results are updated.

6. On the **Export a report** page, under **Include these items from the search**, choose one of the following options:

   - Export only indexed items
   - Export indexed and unindexed items
   - Export only unindexed items

   For more information about unindexed items, see Partially indexed items in Content Search.

7. Choose to include search statistics for all versions of SharePoint documents. This option appears only if the content sources of the search include SharePoint or OneDrive for Business sites.
8. Click **Generate report**.

   The search results report is prepared for downloading, which means the report documents will be uploaded to the Azure Storage area in the Microsoft cloud. When the report is ready for download, the **Download report** link is displayed under **Export report to a computer** in the details pane.

> **NOTE**
> You can also export a report for a Content Search that's associated with an eDiscovery case. To do this, go to **eDiscovery** > **eDiscovery**, select a case, and click **Edit**. On the **Searches** page, select a search, and then click **Export** > **Export a report**.

Step 2: Download the report

The next step is to download the report from the Azure Storage area to your local computer.

1. In the details pane for the search that you generated the report for, under **Export report to a computer**, click **Download report**.

   The **Download report** page is displayed and contains the following information about the report that's downloaded to your computer.

- The number of items that will be downloaded.
- The estimated total size of the items that will be downloaded.
- Whether indexed or unindexed will be exported. Unindexed items are items that have a recognized format, are encrypted, or weren't indexed for other reasons.
- Whether versions of SharePoint documents will be downloaded.
- The status of the report export process. You can start downloading the report even if the preparation of the report isn't complete.

2. Under **Export key**, click **Copy to clipboard**. You use this key in step 5 to download the report.

   > **IMPORTANT**
   > Because anyone can install and start the eDiscovery Export tool, and then use this key to download the search report, be sure to take precautions to protect this key just like you would protect passwords or other security-related information.

3. Click **Download report**.
4. If you're prompted to install the **eDiscovery Export Tool**, click **Install**.
5. In the **eDiscovery Export Tool**, paste the export key that you copied in step 2 in the appropriate box.
6. Click **Browse** to specify the location where you want to download the report.
7. Click **Start** to download the search results to your computer.

   The **eDiscovery Export Tool** displays status information about the export process, including an estimate of the number (and size) of the remaining items to be downloaded. When the export process is complete, you can access the files in the location where they were downloaded.

> **NOTE**
> You can download the report for a Content Search that's associated with an eDiscovery case. To do this, go to **eDiscovery** > **eDiscovery**, select a case, and click **Edit** . On the **Exports** page, select an report export, and then click **Download report** in the details pane.

What's included in the report

When you generate and export a report about the results of a Content Search, the following documents are downloaded:

- **Export Summary:** An Excel document that contains a summary of the export. This includes information such as the number of content sources that were searched, the number of search results from each content location, the estimated number of items, the actual number of items that would be exported, and the estimated and actual size of items that would be exported.

> **NOTE**
>
> If you include unindexed items when exporting the report, the number of unindexed items are included in the total number of estimated search results and in the total number of downloaded search results (if you were to export the search results) that are listed in the Export Summary report. In other words, the total number of items that would be downloaded is equal to the total number of estimated results and the total number of unindexed items.

- **Manifest:** A manifest file (in XML format) that contains information about each item included in the search results.
- **Results:** An Excel document that contains a row with information about each indexed item that would be exported with the search results. For email, the result log contains information about each message, including:
  - The location of the message in the source mailbox (including whether the message is in the primary or archive mailbox).
  - The date the message was sent or received.
  - The Subject line from the message.
  - The sender and recipients of the message.

    For documents from SharePoint and OneDrive for Business sites, the Results log contains information about each document, including:
  - The URL for the document.
  - The URL for the site collection where the document is located.
  - The date that the document was last modified.
  - The name of the document (which is located in the Subject column in the result log).

    > **NOTE**
    >
    > The number of rows in the **Results** report should be equal to the total number of search results minus the total number of items listed in the **Unindexed Items** report.
- **Unindexed Items:** An Excel document that contains information about any unindexed items included in the search results. If you don't include unindexed items when you generate the search results report, this report will still be downloaded, but will be empty.

Use the eDiscovery Export Tool in Microsoft Edge

11/2/2020 • 2 minutes to read • Edit Online

As a result of recent changes to the newest version of Microsoft Edge, ClickOnce support is no longer enabled by default. To continue using the eDiscovery Export Tool to download Content Search or eDiscovery search results, you either need to use Microsoft Internet Explorer or enable ClickOnce support in the newest version of Microsoft Edge.

Enable ClickOnce support in Microsoft Edge

1. In Microsoft Edge, go to **edge://flags/#edge-click-once**.
2. If the existing value is set to **Default** or **Disabled** in the dropdown list, change it to **Enabled**.

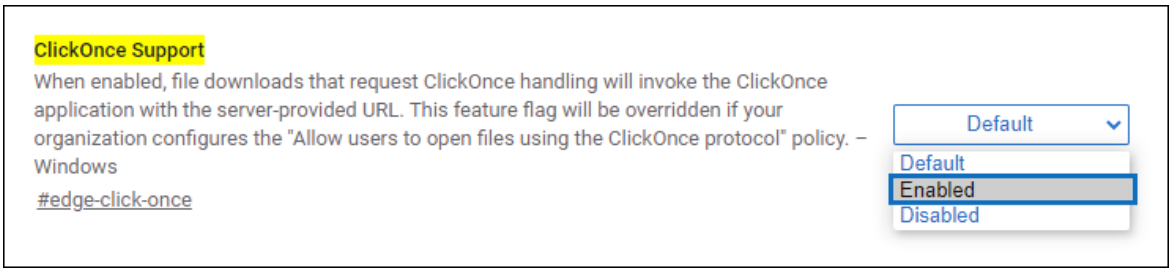

3. Scroll down to the bottom of the browser window and click **Restart** to restart Edge.

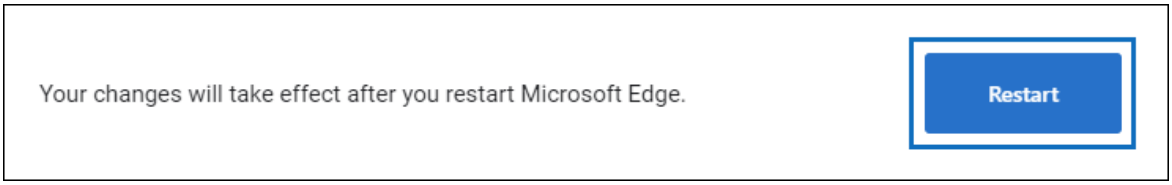

Note: Organizations can use Group Policy to disable ClickOnce support. To check if there is an organizational policy for ClickOnce support, go to **edge://policy**. The following screenshot shows that ClickOnce is enabled across the entire organization. If this policy value is set to **false**, you will need to contact an admin in your organization.

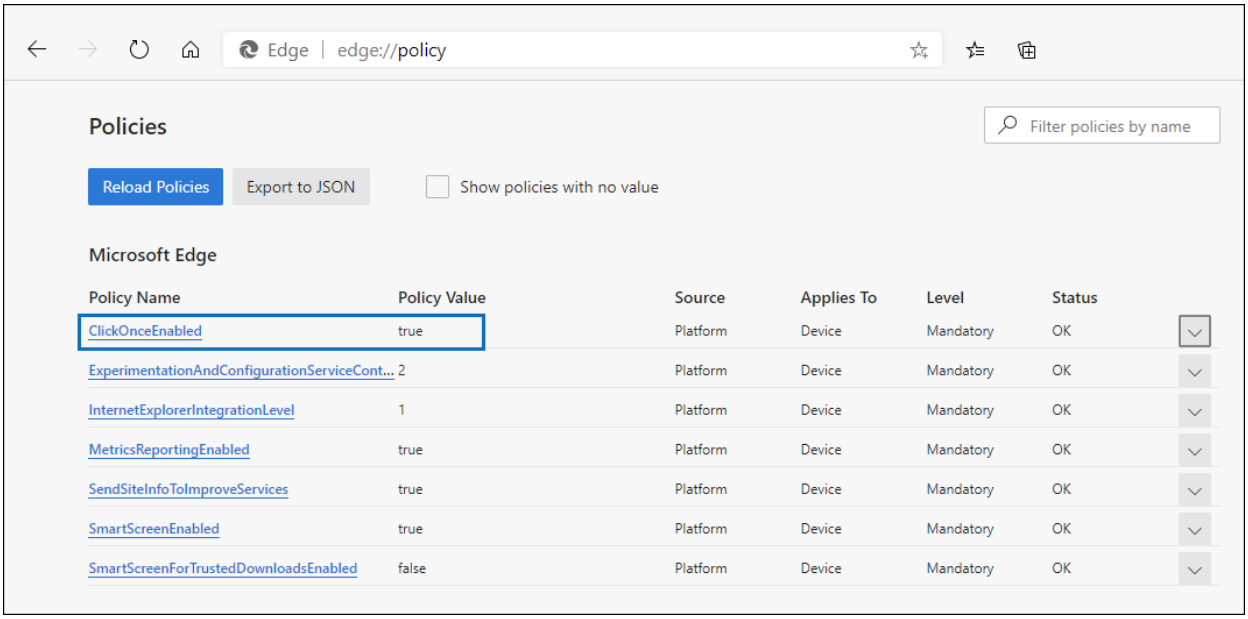

Install and run the eDiscovery Export Tool

1. Click **Download results** on the flyout page of an export in Content Search or an eDiscovery case.

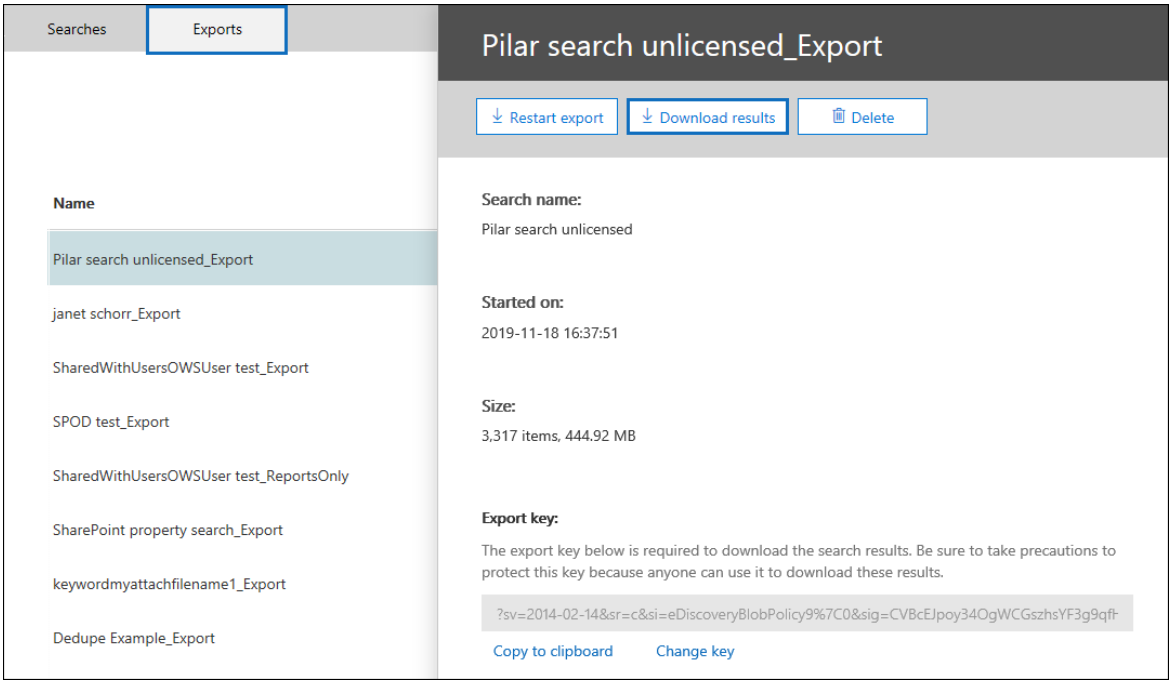

2. You will be prompted with a confirmation to launch the tool, Click **Open**.

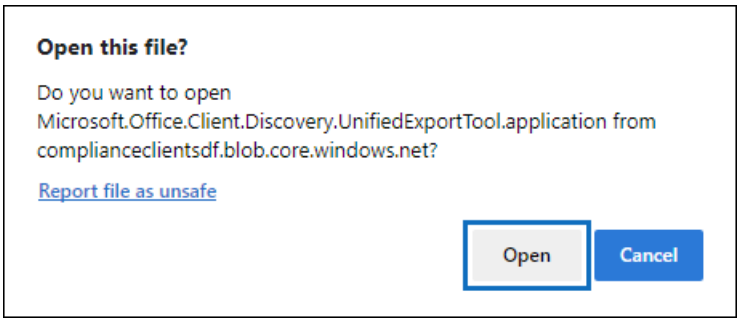

If the eDiscovery Export Tool isn't installed, you will be prompted with a Security Warning,

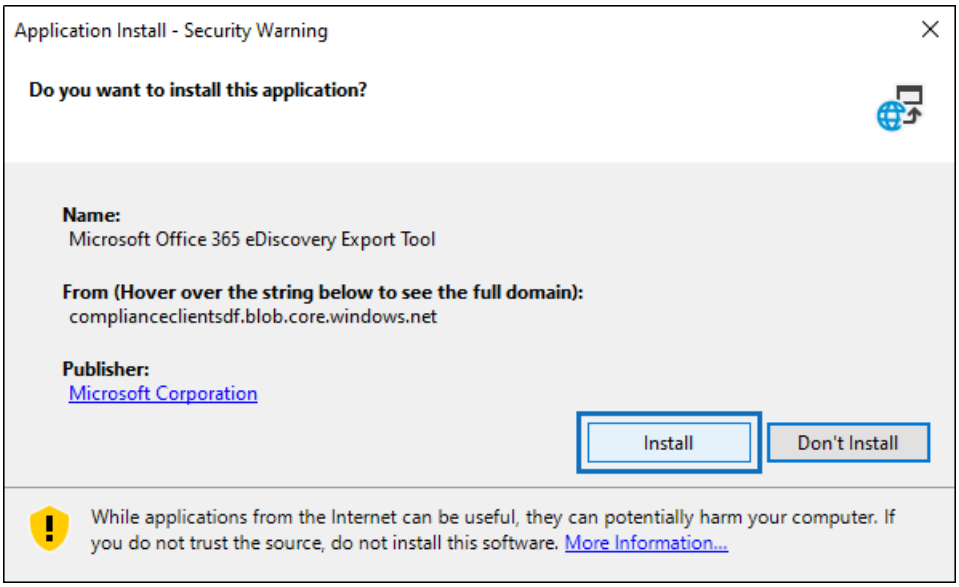

3. Click **Install**. After it's installed, the export tool will launch automatically.

For more information, see the following topics:

- Export Content Search results

- How to enable experiment flags in Microsoft Edge

Search for and delete email messages

11/2/2020 • 9 minutes to read • Edit Online

This article is for administrators. Are you trying to find items in your mailbox that you want to delete? See Find a message or item with Instant Search.

You can use the Content Search feature to search for and delete an email message from all mailboxes in your organization. This can help you find and remove potentially harmful or high-risk email, such as:

- Messages that contain dangerous attachments or viruses
- Phishing messages
- Messages that contain sensitive data

Caution

Search and purge is a powerful feature that allows anyone that is assigned the necessary permissions to delete email messages from mailboxes in your organization.

Before you begin

- To create and run a Content Search, you have to be a member of the **eDiscovery Manager** role group or be assigned the **Compliance Search** management role. To delete messages, you have to be a member of the **Organization Management** role group or be assigned the **Search And Purge** management role. For information about adding users to a role group, see Assign eDiscovery permissions in the Security & Compliance Center.
- You have to use Security & Compliance Center PowerShell to delete messages. See Step 2 for instructions about how to connect.
- A maximum of 10 items per mailbox can be removed at one time. Because the capability to search for and remove messages is intended to be an incident-response tool, this limit helps ensure that messages are quickly removed from mailboxes. This feature isn't intended to clean up user mailboxes.
- The maximum number of mailboxes in a Content Search that you can delete items in by doing a search and purge action is 50,000. If the Content Search (that you create in Step 1) has more than 50,000 source mailboxes, the purge action (that you create in Step 3) will fail. See the More information section for a tip on performing a search and purge operation on more than 50,000 mailboxes.
- The procedure in this article can only be used to delete items in Exchange Online mailboxes and public folders. You can't use it to delete content from SharePoint or OneDrive for Business sites.
- Email items in a review set in an Advanced eDiscovery case can't be deleted by using the procedures in this article. That's because items in a review set are stored in an Azure Storage location, and not in the live service. This means they won't be returned by the content search that you create in Step 1. To delete items in a review set, you have to delete the Advanced eDiscovery case that contains the review set. For more information, see Close or delete an Advanced eDiscovery case.

Step 1: Create a Content Search to find the message to delete

The first step is to create and run a Content Search to find the message that you want to remove from mailboxes in your organization. You can create the search by using the Security & Compliance Center or by running the **New-ComplianceSearch** and **Start-ComplianceSearch** cmdlets. The messages that match the query for this

search will be deleted by running the **New-ComplianceSearchAction -Purge** command in Step 3. For information about creating a Content Search and configuring search queries, see the following topics:

- Content Search in Office 365
- Keyword queries for Content Search
- New-ComplianceSearch
- Start-ComplianceSearch

> **NOTE**
> The content locations that are searched in the Content Search that you create in this step can't include SharePoint or OneDrive for Business sites. You can include only mailboxes and public folders in a Content Search that will be used to email messages. If the Content Search includes sites, you'll receive an error in Step 3 when you run the **New-ComplianceSearchAction** cmdlet.

Tips for finding messages to remove

The goal of the search query is to narrow the results of the search to only the message or messages that you want to remove. Here are some tips:

- If you know the exact text or phrase used in the subject line of the message, use the **Subject** property in the search query.
- If you know that exact date (or date range) of the message, include the **Received** property in the search query.
- If you know who sent the message, include the **From** property in the search query.
- Preview the search results to verify that the search returned only the message (or messages) that you want to delete.
- Use the search estimate statistics (displayed in the details pane of the search in the Security & Compliance Center or by using the Get-ComplianceSearch cmdlet) to get a count of the total number of results.

Here are two examples of queries to find suspicious email messages.

- This query returns messages that were received by users between April 13, 2016 and April 14, 2016 and that contain the words "action" and "required" in the subject line.

  ```
  (Received:4/13/2016..4/14/2016) AND (Subject:'Action required')
  ```

- This query returns messages that were sent by chatsuwloginsset12345@outlook.com and that contain the exact phrase "Update your account information" in the subject line.

  ```
  (From:chatsuwloginsset12345@outlook.com) AND (Subject:"Update your account information")
  ```

Here's an example of using a query to create and start a search by running the **New-ComplianceSearch** and **Start-ComplianceSearch** cmdlets to search all mailboxes in the organization:

```
$Search=New-ComplianceSearch -Name "Remove Phishing Message" -ExchangeLocation All -ContentMatchQuery
'(Received:4/13/2016..4/14/2016) AND (Subject:"Action required")'
Start-ComplianceSearch -Identity $Search.Identity
```

Step 2: Connect to Security & Compliance Center PowerShell

The next step is to connect to Security & Compliance Center PowerShell for your organization. For step-by-step instructions, see Connect to Security & Compliance Center PowerShell.

After you've connected to Security & Compliance Center PowerShell, run the **New-ComplianceSearch** and **Start-ComplianceSearch** cmdlets that you prepared in the previous step.

Step 3: Delete the message

After you've created and refined a Content Search to return the message that you want to remove and are connected to Security & Compliance Center PowerShell, the final step is to run the **New-ComplianceSearchAction** cmdlet to delete the message. You can soft- or hard-delete the message. A soft-deleted message is moved to a user's Recoverable Items folder and retained until the deleted item retention period expires. Hard-deleted messages are marked for permanent removal from the mailbox and will be permanently removed the next time the mailbox is processed by the Managed Folder Assistant. If single item recovery is enabled for the mailbox, hard-deleted items will be permanently removed after the deleted item retention period expires. If a mailbox is placed on hold, deleted messages are preserved until the hold duration for the item expires or until the hold is removed from the mailbox.

In the following example, the command soft-deletes the search results returned by a Content Search named "Remove Phishing Message".

```
New-ComplianceSearchAction -SearchName "Remove Phishing Message" -Purge -PurgeType SoftDelete
```

To hard-delete the items returned by the "Remove Phishing Message" content search, you would run this command:

```
New-ComplianceSearchAction -SearchName "Remove Phishing Message" -Purge -PurgeType HardDelete
```

When you run the previous command to soft- or hard-delete messages, the search specified by the *SearchName* parameter is the Content Search that you created in Step 1.

For more information, see New-ComplianceSearchAction.

More information

- **How do you get status on the search and remove operation?**

  Run the **Get-ComplianceSearchAction** to get the status on the delete operation. The object that is created when you run the **New-ComplianceSearchAction** cmdlet is named using this format: `<name of Content Search>_Purge`.

- **What happens after you delete a message?**

  A message that's deleted with the `New-ComplianceSearchAction -Purge -PurgeType HardDelete` command is moved to the Purges folder and can't be accessed by the user. After the message is moved to the Purges folder, the message is retained for the duration of the deleted item retention period if single item recovery is enabled for the mailbox. (In Microsoft 365, single item recovery is enabled by default when a new mailbox is created.) After the deleted item retention period expires, the message is marked for permanent deletion and will be purged from Microsoft 365 the next time the mailbox is processed by the Managed Folder assistant.

  If you use the `New-ComplianceSearchAction -Purge -PurgeType SoftDelete` command, messages are moved to the Deletions folder in the user's Recoverable Items folder. It isn't immediately purged from Microsoft

365. The user can recover messages in the Deleted Items folder for the duration based on the deleted item retention period configured for the mailbox. After this retention period expires (or if user purges the message before it expires), the message is moved to the Purges folder and can no longer be accessed by the user. Once in the Purges folder, the message is retained for the duration based on the deleted item retention period configured for the mailbox if single items recovery is enabled for the mailbox. (In Microsoft 365, single item recovery is enabled by default when a new mailbox is created.) After the deleted item retention period expires, the message is marked for permanent deletion and will be purged from Microsoft 365 the next time that the mailbox is processed by the Managed Folder assistant.

- **What if you have to delete a message from more than 50,000 mailboxes?**

  As previously stated, you can perform a search and purge operation on a maximum of 50,000 mailboxes. If you have to do a search and purge operation on more than 50,000 mailboxes, consider creating temporary search permissions filters that would reduce the number of mailboxes that would be searched to less than 50,000 mailboxes. For example, if your organization contains mailboxes in different departments, states, or countries, you can create a mailbox search permissions filter based on one of those mailbox properties to search a subset of mailboxes in your organization. After you create the search permissions filter, you would create the search (described in Step 1) and then delete the message (described in Step 3). Then you can edit the filter to search for and purge messages in a different set of mailboxes. For more information about creating search permissions filters, see Configure permissions filtering for Content Search.

- **Will unindexed items included in the search results be deleted?**

  No, the `New-ComplianceSearchAction -Purge command doesn't delete unindexed items.

- **What happens if a message is deleted from a mailbox that has been placed on In-Place Hold or Litigation Hold or is assigned to an Microsoft 365 retention policy?**

  After the message is purged and moved to the Purges folder, the message is retained until the hold duration expires. If the hold duration is unlimited, then items are retained until the hold is removed or the hold duration is changed.

- **Why is the search and remove workflow divided among different security and compliance center role groups?**

  As previously explained, a person has to be a member of the eDiscovery Manager role group or be assigned the Compliance Search management role to search mailboxes. To delete messages, a person has to be a member of the Organization Management role group or be assigned the Search And Purge management role. This makes it possible to control who can search mailboxes in the organization and who can delete messages.

Search for Teams chat data for on-premises users

2/18/2021 • 7 minutes to read • Edit Online

If your organization has an Exchange hybrid deployment (or your organization synchronizes an on-premises Exchange organization with Office 365) and has enabled Microsoft Teams, on-premises users can use the Teams chat application for instant messaging. For a cloud-based user, Teams chat data (also called *1x1 or 1xN chats*) is saved to their primary cloud-based mailbox. When an on-premises user uses the Teams chat application, their chat messages can't be stored in their primary mailbox, which is located on-premises. To get around this limitation, Microsoft has released a new feature where a cloud-based storage area is created so that you use eDiscovery tools to search for and export Teams chat data for on-premises users.

Here are the requirements and limitations for enabling cloud-based storage for on-premises users:

- The user accounts in your on-premises directory service (such as Active Directory) must be synchronized with Azure Active Directory, the directory service in Microsoft 365. This means that a mail user account is created in Microsoft 365 and is associated with a user whose primary mailbox is located in the on-premises organization.
- The user whose primary mailbox is located in the on-premises organization must be assigned a Microsoft Teams license and a minimum of an Exchange Online Plan 1 license.
- If your organization doesn't have an Exchange hybrid deployment, you must synchronize your on-premises Exchange schema to Azure Active Directory. If you don't do this, you might risk creating duplicate cloud-based mailboxes in Exchange Online for users that have a mailbox in your on-premises Exchange organization.
- Only Teams chat data associated with an on-premises user is stored in the cloud-based storage area. An on-premises user can't access this storage area in any way.

> **NOTE**
> Teams channel conversations are always stored in the cloud-based mailbox that's associated with the Team, which means you can search for channel conversations. For more information about searching Teams channel conversations, see Searching Microsoft Teams and Microsoft 365 Groups.

How it works

If a Microsoft Teams-enabled user has an on-premises mailbox and their user account/identity has been synched to the cloud, Microsoft creates cloud-based storage to associate the on-premises user's 1xN Teams chat data with. Teams chat data for on-premises users is indexed for search. This lets you Use Content Search (and searches associated with Core and Advanced eDiscovery cases) to search, preview, and export Teams chat data for on-premises users. You can also use ***ComplianceSearch** cmdlets in the Security & Compliance Center PowerShell to search for Teams chat data for on-premises users.

The following graphic shows the workflow of how Teams chat data for on-premises users is available to search, preview, and export.

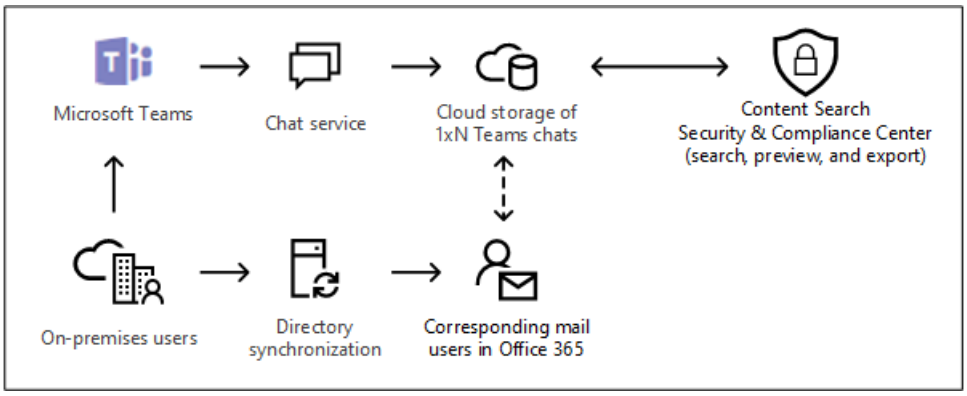

In addition to this new capability, you can still use Content Search to search, preview, and export Teams content in the cloud-based SharePoint site and Exchange mailbox associated with each Microsoft Team and 1xN Teams chat data in the Exchange Online mailbox for cloud-based users.

What happens after this feature is enabled?

After this feature is deployed in your organization, the following changes are made in Content Search and in searches associated with an eDiscovery case in the Security & Compliance Center:

- The **Add Office app content for on-premises users** checkbox is added under the **Locations** in Content Search.

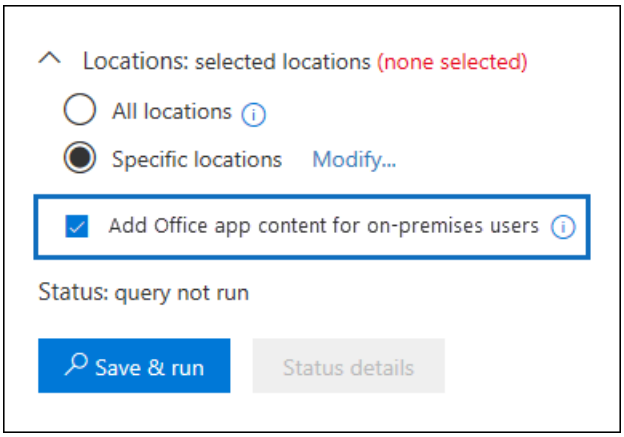

- On-premises users are displayed in the content locations picker that you use to select user mailboxes to search.

Searching for Teams chat content for on-premises users

After the feature has been enabled, you can use Content Search in the Security & Compliance Center to search for Teams chat data for on-premises users.

1. In the Security & Compliance Center, go to **Search** > **Content search**
2. On the **Search** page, click **+ New search**.

   As previously explained, the **Add Office app content for on-premises users** checkbox is displayed under **Locations**. It's selected by default.

3. Create the keyword query and add conditions to the search query if necessary. To only search for Team chats data, you can add the following query in the **Keywords** box:

   ```
   kind:im
   ```

4. At this point, you can choose one of the following options under **Locations**:

- **All locations:** Select this option to search the mailboxes of all users in your organization. When the checkbox is selected, all cloud-based storage of Teams chat data for on-premises users will also be searched.
- **Specific locations:** Select this option and then click **Modify** > Choose user, groups, or teams to search specific mailboxes. As previously explained, the locations picker lets you search for Teams chat data for on-premises users.

5. Save and run the search. Any search results for on-premises users can be previewed like any other search results. You can also export the search results (including any Teams chat data) to a PST file. For more information, see:
   - Create a search
   - Preview search results
   - Export Content Search results

Using PowerShell to search for Teams chat data for on-premises users

You can use the **New-ComplianceSearch** and **Set-ComplianceSearch** cmdlets in the Security & Compliance Center PowerShell to search for Teams chat data for on-premises users. As previously explained, you don't have to submit a support request to use PowerShell to search for Teams chat data for on-premises users.

1. Connect to Security & Compliance Center PowerShell.
2. Run the following PowerShell command to create a content search that searches for Teams chat data for on-premises users.

   ```
   New-ComplianceSearch <name of new search> -ContentMatchQuery <search query> -ExchangeLocation <on-premises user> -IncludeUserAppContent $true -AllowNotFoundExchangeLocationsEnabled $true
   ```

   The *IncludeUserAppContent* parameter is used to specify the cloud-based storage for the user or users who are specified by the *ExchangeLocation* parameter. The *AllowNotFoundExchangeLocationsEnabled* allows you to search the cloud-based storage for on-premises users. When you use the `$true` value for this parameter, the search doesn't try to validate the existence of the mailbox before it runs. This is required to search the cloud-based storage for on-premises users because this cloud-based storage doesn't resolve as a regular cloud-based mailbox.

   The following example searches for Teams chats (which are instant messages) that contain keyword "redstone" in the cloud-based storage for Sara Davis, who is an on-premises user in the Contoso organization.

   ```
   New-ComplianceSearch "Redstone_Search" -ContentMatchQuery "redstone AND kind:im" -ExchangeLocation sarad@contoso.com -IncludeUserAppContent $true -AllowNotFoundExchangeLocationsEnabled $true
   ```

   After you create a search, be sure to use the **Start-ComplianceSearch** cmdlet to run the search.

For more information using these cmdlets, see:

- New-ComplianceSearch
- Set-ComplianceSearch
- Start-ComplianceSearch

Known issues

- Currently, you can search, preview, and export Teams chat data for on-premises users. You can also place the Teams chat data for an on-premises user on a hold associated with a Core or Advanced eDiscovery case, and apply a retention policy for Teams chats or channel messages for on-premises users. However at this time, you can't apply a retention policy for other content locations (such as Exchange mailboxes and SharePoint sites) for on-premises users.

Frequently asked questions

Where is the cloud-based storage for on-premises users located?

Teams chat data is stored in the Preferred Data Location (PDL) for an on-premises user. The PDL is honored in both Single-Geo and Multi-Geo environments. For more information, see Microsoft 365 Multi-Geo.

Are there any other requirements other than submitting a support request?

As previously explained, the identities of users with on-prem mailboxes must be synchronized to your cloud-based organization so that a corresponding mail user account is created for each on-premises user account in Office 365. Your organization must also have an Office 365 enterprise subscription, such as an Office 365 Enterprise E1, E3, or E5 subscription.

Is there a risk of losing the Teams chat data if the user's on-premises mailbox is migrated to the cloud?

No. When you migrate the primary mailbox of an on-premises user to the cloud, the Teams chat data for that user will be migrated to their new cloud-based primary mailbox.

Can I apply an eDiscovery hold or retention policies to on-premises users?

Yes. You can apply eDiscovery holds or retention policies for Teams chats and channel messages of on-premises users.

Can Content Search find older Teams chat data for on-premises users before the time my organization submitted the request to enable this feature?

Microsoft started storing the Teams chat data for on-premises users on January 31, 2018. So, if the identity of an on-premises Teams user has been synched between Active Directory and Azure Active Directory since this date, then their Teams chat data is stored in the cloud and is searchable using Content Search. Microsoft is also working on storing Teams chat data from prior to January 31, 2018 in the cloud-based storage for on-premises users. More information about this will be available soon.

Do on-premises users need a license to store their Teams chat data in the cloud?

Yes. To store Teams chat data for an on-premises user in a cloud-based storage, the user must be assigned a Microsoft Teams license and an Exchange Online Plan license in Office 365 (or Microsoft 365).

Bulk edit Content Searches

5/5/2020 • 6 minutes to read • Edit Online

You can use the Bulk Search Editor in the Content Search tool to edit multiple searches at the same time. Using this tool lets you quickly change the query and content locations for one or more searches. Then you can rerun the searches and get new estimated search results for the revised searches. The editor also lets you copy and paste queries and content locations from a Microsoft Excel file or text file. This means you can use the Search Statistics tool to view the statistics of one or more searches, export the statistics to a CSV file, where you can edit the queries and content locations in Excel. Then you use the Bulk Search Editor to add the revised queries and content locations to the searches. After you've revised one or more searches, you can restart them and get new estimated search results.

For more information about using the Search Statistics tool, see View keyword statistics for Content Search results.

Use the Bulk Search Editor to change queries

1. Go to https://protection.office.com, and then select **Search** > **Content search**.
2. In the list of searches, select one or more searches, and then select **Bulk Search Editor**.

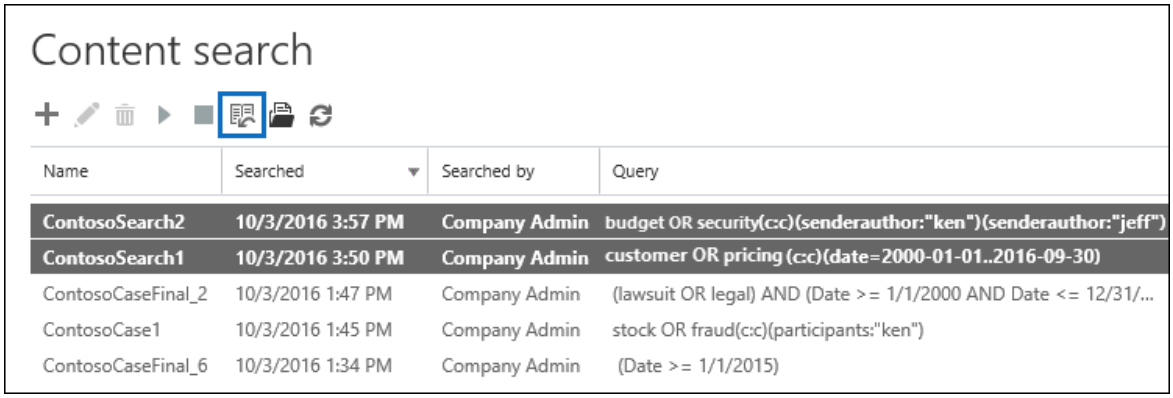

The following information is displayed on the **Queries** page of the Bulk Search Editor.

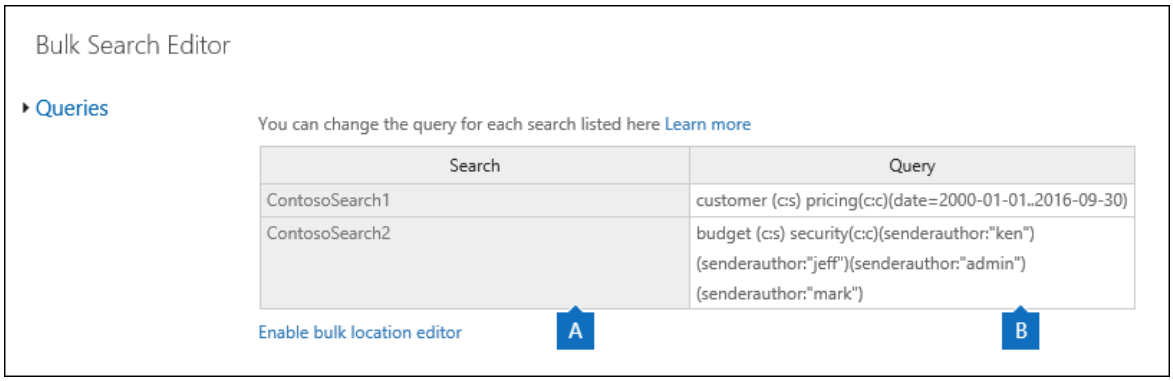

a. The **Search** column displays the name of the Content Search. As previously stated, you can edit the query for multiple searches.

b. The **Query** column displays the query for the Content Search listed in the **Search** column. If the query was created using the keyword list feature, the keywords are separated by the text ** `(c:s)` **. This indicates that the keywords are connected by the OR operator. Additionally, if the query includes conditions, the keywords and the conditions are separated by the text ** `(c:c)` **.** This

indicates that the keywords (or keyword phases) are connected to the conditions by the **AND** operator. For example, in the previous screenshot the for search ContosoSearch1, the KQL query that is equivalent to `customer (c:s) pricing(c:c)(date=2000-01-01..2016-09-30)` would be `(customer OR pricing) AND (date=2002-01-01..2016-09-30)`.

3. To edit a query, select in the cell of the query that you want to change and doing one of the following things. The cell is bordered by a blue box when you select it.

   - Type the new query in the cell. You can't edit a portion of the query. You have to type the entire query.

     Or

   - Paste a new query in the cell. This assumes that you've copied the query text from a file, such as a text file or an Excel file.

4. After you've edited one or more queries on the **Queries** page, select **Save**.

   The revised query is displayed in the **Query** column for the selected search.

5. Select **Close** to close the Bulk Search Editor.

6. On the **Content search** page, select the search that you edited, and select **Start** search to restart the search using the revised query.

Here are some tips for editing queries using the Bulk Search Editor:

- Copy the existing query (by using **Ctrl C**) to a text file. Edit the query in the text file, and then copy the revised query and paste it (using **Ctrl V**) back into the cell on the **Queries** page.
- You can also copy queries from other applications (such as Microsoft Word or Microsoft Excel). However, you might inadvertently add unsupported characters to a query using the Bulk Search Editor. The best way to prevent unsupported characters is to just type the query in a cell on the **Queries** page. Or you can copy a query from Word or Excel and then paste it to file in a plain text editor, such as Microsoft Notepad. Then save the text file and select **ANSI** in the **Encoding** drop-down list. This removes any formatting and unsupported characters. Then you can copy and paste the query from the text file to the **Queries** page.

Use the Bulk Search Editor to change content locations

1. In the Bulk Search Editor for one or more selected searches, select **Enable bulk location editor**, and then select the **Locations** link that is displayed on the page.

   The following information is displayed on the **Locations** page of the Bulk Search Editor.

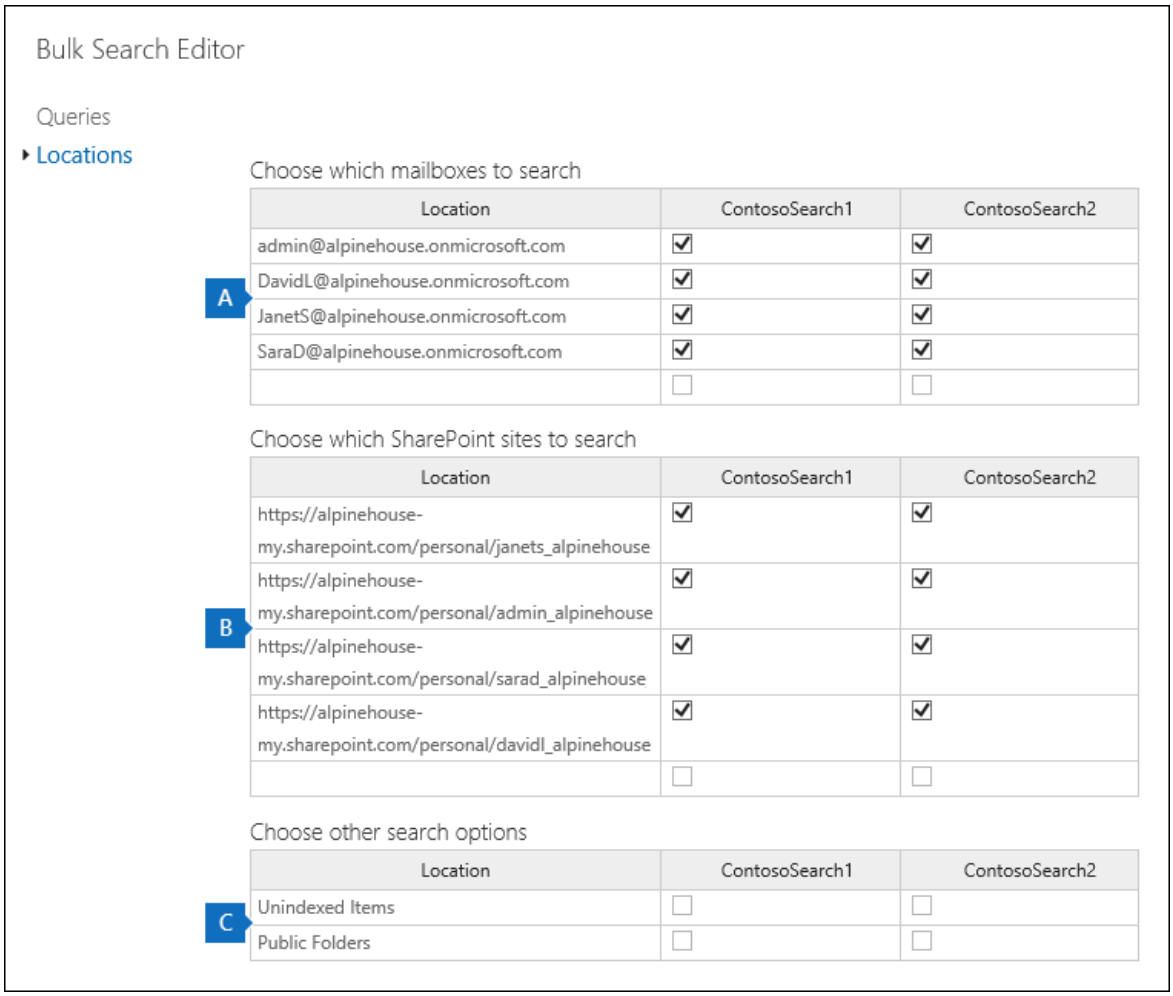

a. **Mailboxes to search** This section displays a column for each selected Content Search and a row for each mailbox that's included in the search. A check mark indicates that the mailbox is included in the search. You can add mailboxes to a search by typing the email address of the mailbox in a blank row and then selecting the check box for the Content Search that you want to add it to. Or you can remove a mailbox from a search by clearing the check box.

b. **SharePoint sites to search** This section displays a row for each SharePoint and OneDrive site that's included in each selected Content Search. A check mark indicates that the site is included in the search. You can add sites to a search by typing the URL for the site in a blank row and then selecting the check box for the Content Search that you want to add it to. Or you can remove a site from a search by clearing the check box.

c. **Other search options** This section indicates whether unindexed items and public folders are included in the search. To include them, make sure the check box is selected. To remove them, clear the check box.

2. After you've edited one or more of the sections on the **Locations** page, select **Save**.

   The revised content locations are displayed in the appropriate section for the selected searches.

3. Select **Close** to close the Bulk Search Editor.

4. On the **Content search** page, select the search that you edited, and select **Start** search to restart the search using the revised content locations.

Here are some tips for editing content locations using the Bulk Search Editor:

- You can edit Content Searches to search all mailboxes or sites in the organization by typing **All** in a blank row in the **Mailboxes to search** or **SharePoint sites to search** section and then selecting the check box.

- You can add multiple content locations to one or more searches by copying multiple rows from a text file or an Excel file and then pasting them in a section on the **Locations** page. After you add new locations, be sure to select the check box for each search that you want add the location to.

  > **TIP**
  >
  > To generate a list of email addresses for all the users in your organization, run the PowerShell command in Step 2 in Step 2: Generate a list of users. Or follow the steps in Get a list of all user OneDrive URLs in your organization to generate a list of all OneDrive for Business sites in your organization. Note that you'll have to append the URL for your organization's MySite domain (for example, https://contoso-my.sharepoint.com) to the OneDrive for Business sites that's created by the script. After you have list of email addresses or OneDrive for Business sites, you can copy and paste them to the **Locations** page in the Bulk Search Editor.

- After you select **Save** to save changes in Bulk Search Editor, the email address for mailboxes that you added to a search will be validated. If the email address doesn't exist, an error message is displayed saying the mailbox can't be located. URLs for sites aren't validated.

Prepare a CSV file for an ID list Content Search

11/2/2020 • 5 minutes to read • Edit Online

You can search for specific mailbox email messages and other mailbox items using a list of Exchange IDs. To create an ID list search (formally called a targeted search), you submit a comma separated value (CSV) file that identifies the specific mailbox items to search for. For this CSV file you use the **Results.csv** file or the **Unindexed Items.csv** file that are included when you export the Content Search results or export a Content Search report from and existing Content Search. Then you edit one of these files to indicate the specific items to search for, and then create a new ID list search and submit the CSV file.

Here's a quick overview of the process for creating an ID list search.

1. Create and run a new or guided Content Search in the Security & Compliance Center.
2. Export the content search results or export the content search report. For more information, see:
   - Export Content Search results
   - Export a Content Search report
3. Edit the **Results.csv** file or the **Unindexed Items.csv** and identify the specific mailbox items that you want to include in the ID list search. See the instructions for preparing a CSV file for an ID list search.
4. Create a new ID list search (see the instructions) and submit the CSV file that you prepared. The search query that's created will only search for the items selected in the CSV file.

> **NOTE**
> ID list searches are only supported for mailbox items. You can't search for SharePoint and OneDrive documents in an ID list search.

Why create an ID list search? If you're unable to determine if an item is responsive to an eDiscovery request based on the metadata in the **Results.csv** or **Unindexed Items.csv** files, you can use an ID list search to find, preview, and then export that item to determine if it's responsive to the case you're investigating. ID list searches are typically used to search for and return a specific set of unindexed items.

Prepare the CSV file for an ID list search

After you export the search results or report for a content search, you can perform the following steps to prepare the CSV file for an ID list search. This CSV file will identify every item in the ID list search.

Note that you can use a CSV file from a search that included SharePoint sites and OneDrive accounts, but you can select *only* mailbox items for an ID list search. If you select a document in SharePoint or OneDrive, the CSV file will fail validation when you create an ID list search.

1. Open the **Results.csv** or **Unindexed Items.csv** file in Excel.
2. In the **Selected** column, type **Yes** in the cell that corresponds to the item that you want to search for. Repeat this step for every item that you want to search for.

> **IMPORTANT**
>
> When you open the CSV file in Excel, the data format for the **Document ID** column is changed to **General**. This results in displaying the document ID for an item in scientific notation. For example, the document ID of "481037338205" is displayed as "4.81037E+11" You have to perform the next steps to change the data format of the **Document ID** column to **Number** to restore the correct format for the document ID. If you don't do this, the ID list search that uses the CSV file will fail.

3. Right-click the entire **Document ID** column and select **Format Cells**.
4. In the **Category** box, click **Number**.
5. Change the number of decimal places to **0**, and then click **OK** to save your changes. Notice that the values in the Document ID column are changed to numbers.

   Here's an example of the a CSV file that's ready to be submitted for a ID list content search.

	A	B	C	D	E	F
1	Selected	ExportItem Id	Item Identity	Document ID	Duplicate	Original Path
2	Yes	F9DEB9A6DD43EE	sarad@contoso	1529009361325		sarad@contoso.onmicrosoft.com, Primary, 571e1954-ce91-449a-9f0f-b4888c2a35b6\
3		D10D05C6E21AB6	sarad@contoso	1529009360704		sarad@contoso.onmicrosoft.com, Primary, 571e1954-ce91-449a-9f0f-b4888c2a35b6\
4		BF45C56B430C3B	sarad@contoso	1529009360700		sarad@contoso.onmicrosoft.com, Primary, 571e1954-ce91-449a-9f0f-b4888c2a35b6\
5		BC37D89AA2E6BE	sarad@contoso	1529009360698		sarad@contoso.onmicrosoft.com, Primary, 571e1954-ce91-449a-9f0f-b4888c2a35b6\
6	Yes	5A4AF6B8E19650	sarad@contoso	1529009360696		sarad@contoso.onmicrosoft.com, Primary, 571e1954-ce91-449a-9f0f-b4888c2a35b6\
7		F959A81553F305F	sarad@contoso	1529009360694		sarad@contoso.onmicrosoft.com, Primary, 571e1954-ce91-449a-9f0f-b4888c2a35b6\
8		0DB064D9D61FED	sarad@contoso	1529009360684		sarad@contoso.onmicrosoft.com, Primary, 571e1954-ce91-449a-9f0f-b4888c2a35b6\
9		970B6230C9A29F8	sarad@contoso	1529009360682		sarad@contoso.onmicrosoft.com, Primary, 571e1954-ce91-449a-9f0f-b4888c2a35b6\
10	Yes	1EC17499A9F8AC	sarad@contoso	1529009360680		sarad@contoso.onmicrosoft.com, Primary, 571e1954-ce91-449a-9f0f-b4888c2a35b6\
11		ED0E648C4630A2	sarad@contoso	1529009360678		sarad@contoso.onmicrosoft.com, Primary, 571e1954-ce91-449a-9f0f-b4888c2a35b6\
12		BC77177C4158B5	sarad@contoso	1529009360674		sarad@contoso.onmicrosoft.com, Primary, 571e1954-ce91-449a-9f0f-b4888c2a35b6\
13	Yes	03987A330074CC	sarad@contoso	1529009360520		sarad@contoso.onmicrosoft.com, Primary, 571e1954-ce91-449a-9f0f-b4888c2a35b6\
14		AF3043CCB08D37	sarad@contoso	1529009360518		sarad@contoso.onmicrosoft.com, Primary, 571e1954-ce91-449a-9f0f-b4888c2a35b6\
15		C84B4A000C7C63	sarad@contoso	1529009360516		sarad@contoso.onmicrosoft.com, Primary, 571e1954-ce91-449a-9f0f-b4888c2a35b6\
16	Yes	DB771AD57C87A0	sarad@contoso	1529009360514		sarad@contoso.onmicrosoft.com, Primary, 571e1954-ce91-449a-9f0f-b4888c2a35b6\

6. Save the CSV file or use **Save As** to the save the file with different file name. In both cases, be sure to save the file with the CSV format.

Create an ID list search

The next step is to create a new ID list Content Search and submit the CSV file that you prepared in the previous step.

> **IMPORTANT**
>
> You should create an ID list search no more than 2 days after exporting the results or report from a Content Search. If the search results or report where exported more than 2 days ago, you should re-export the search results or report to generate updated CSV files. Then you can prepare one of the updated CSV files and use it to create an ID list search.

1. In the Security & Compliance Center, go to **Search** > **Content search**.
2. On the **Search** page, click the arrow next to **+ New search**, and then click **Search by ID List**.

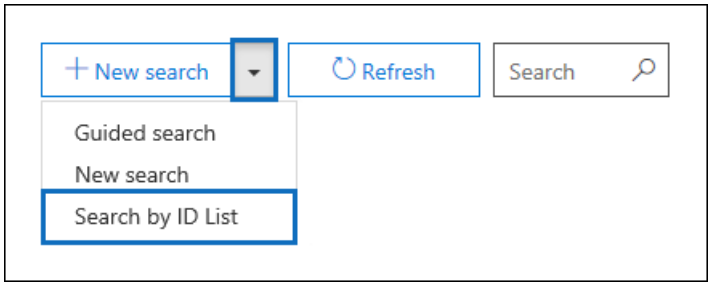

3. On the **Search by ID List** flyout, name the search (and optionally describe it) and then click **Browse** and

select the CSV file that you prepared in the previous step.

Microsoft 365 attempts to validate the CSV file. If the validation is unsuccessful, an error message is displayed that might help you troubleshoot the validation errors. The CSV file has to be successfully validated to create an ID list search.

4. After the CSV file is successfully validated, click **Search** to create the ID list search.

   Here's an example of the estimated search results and the query that's generated for an ID list search.

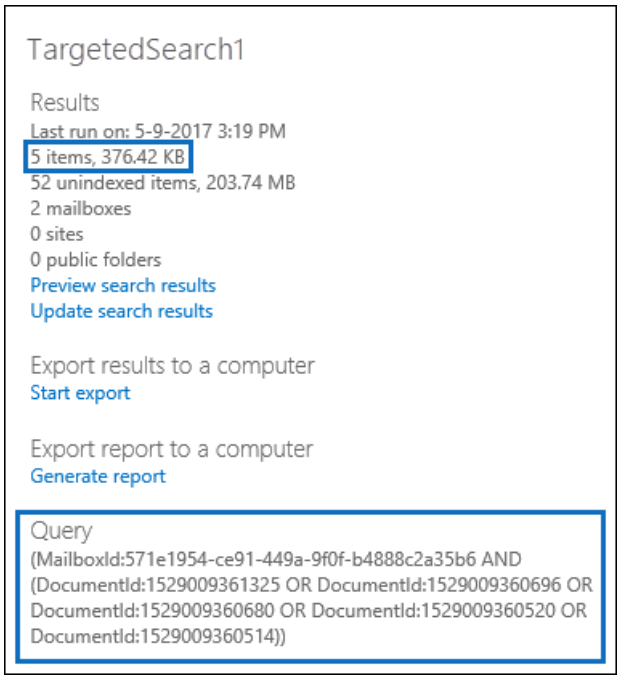

   Note that the number of estimated items displayed in statistics for the ID search should match the number of items that you selected in the CSV file.

5. Preview or export the items returned by the ID list search.

> **NOTE**
>
> If you move a mailbox after creating an ID list search, the query for the search won't return the specified items. That's because the **DocumentId** property for mailbox items are changed when a mailbox is moved. In the rare instance when a mailbox is moved after you create an ID list search, you should create a new content search (or update the search results for the existing content search) and then export the search results or report to generate updated CSV files that can be used to create a new ID list search.

Check your Content Search query for errors

11/2/2020 • 2 minutes to read • Edit Online

When you create or edit a Content Search, you can have Microsoft 365 check your query for unsupported characters and lowercase Boolean operators. How? Just click **Check query for typos** on the query page of a Content Search.

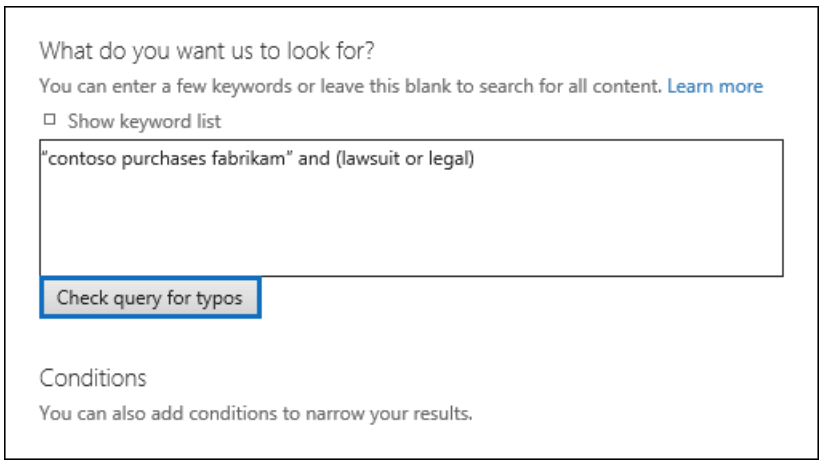

Here's a list of the unsupported characters that we check for. Unsupported characters are often hidden, and they typically cause a search error or return unintended results.

- **Smart quotation marks** - Smart single and double quotation marks (also called curly quotes) aren't supported. Only straight quotation marks can be used in a search query.
- **Non-printable and control characters** - Non-printable and control characters don't represent a written symbol, such as an alpha-numeric character. Examples of non-printable and control characters include characters that format text or separate lines of text.
- **Left-to-right and right-to-left marks** - These marks are control characters used to indicate text direction for left-to-right languages (such as English and Spanish) and right-to-left languages (such as Arabic and Hebrew).
- **Lowercase Boolean operators** - If you use a Boolean operator, such as **AND**, **OR**, and **NOT** in a search query, it must be uppercase. When we check a query for typos, the query syntax will often indicate that a Boolean operator is being used even though lowercase operators might be used; for example, `(WordA or WordB) and (WordC or WordD)`.

What happens if a query has an unsupported character?

If unsupported characters are found in your query, a warning message is displayed that says unsupported characters were found and suggests an alternative. You then have the option keep the original query or replace it with the suggested revised query. Here's an example of the warning message that's displayed after you click **Check query for typos** for the search query in the previous screenshot. Notice that the original query contains smart quotes and lowercase Boolean operators.

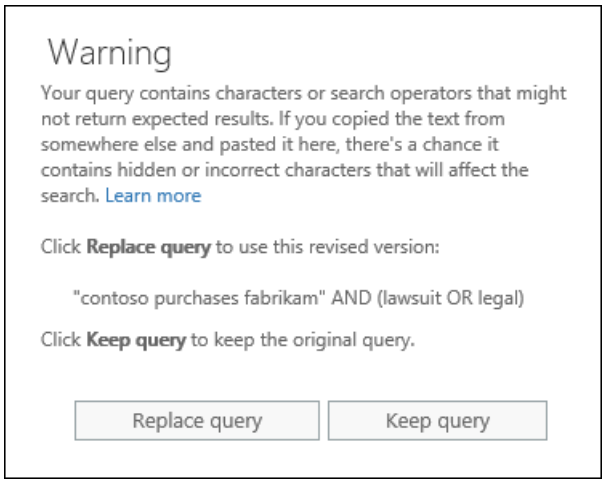

How to prevent unsupported characters in your search queries

Unsupported characters are typically added to a query when you copy the query or parts of the query from other applications (such as Microsoft Word or Microsoft Excel) and paste them in the keyword box on the query page of a Content Search. The best way to prevent unsupported characters is to just type the query in the keyword box. Or you can copy a query from Word or Excel, and then paste it in a plain text editor, such as Microsoft Notepad. Save the text file and select **ANSI** in the **Encoding** drop-down list. This will remove any formatting and unsupported characters. Then you can copy and paste the query from the text file to the keyword query box.

Investigate, troubleshoot, and resolve common eDiscovery issues

2/18/2021 • 5 minutes to read • Edit Online

This topic covers basic troubleshooting steps you can take to identify and resolve issues you may encounter during an eDiscovery search or elsewhere in the eDiscovery process. Resolving some of these scenarios requires help from Microsoft Support. Information on when to contact Microsoft Support is included in the resolution steps.

Error/issue: Ambiguous location

If you try to add user's mailbox location to search and there are duplicate or conflicting objects with the same userID in the Exchange Online Protection (EOP) directory, you receive this error:

```
The compliance search contains the following invalid location(s):useralias@contoso.com. The location
"useralias@contoso.com" is ambiguous
```

.

Resolution

Check for duplicate users or distribution list with the same user ID.

1. Connect to Security & Compliance Center PowerShell.
2. Run the following command to retrieve all instances of the username:

   ```
   Get-Recipient <username>
   ```

   The output for 'useralias@contoso.com' would be similar to the following:

NAME	RECIPIENTTYPE
Alias, User	MailUser
Alias, User	User

3. If multiple users are returned, locate and fix the conflicting object.

Error/issue: Search fails on specific locations

An eDiscovery or content search may yield the following error:

```
This search completed with (#) errors. Would you like to retry the search on the failed locations?
```

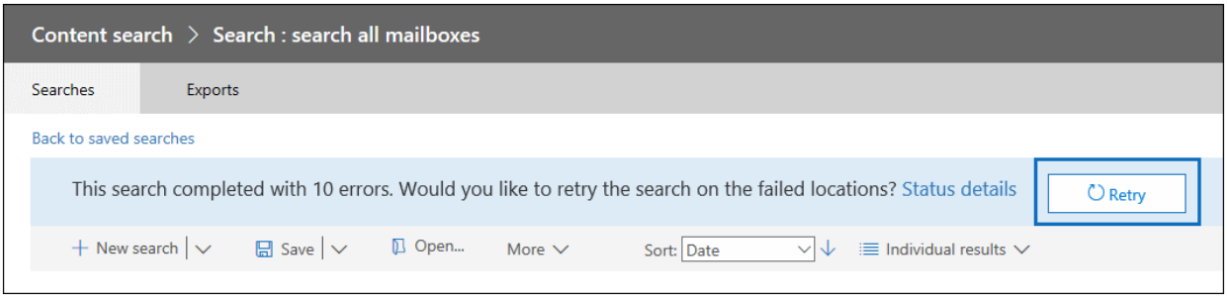

Resolution

If you receive this error, we recommend that you verify the locations that failed in the search then rerun the search only on the failed locations.

1. Connect to Security & Compliance Center PowerShell and then run the following command:

```
Get-ComplianceSearch <searchname> | FL
```

2. From the PowerShell output, view the failed locations in the errors field or from the status details in the error from the search output.
3. Retry the eDiscovery search on the failed locations only.
4. If you continue to receive these errors, see Retry failed locations for more troubleshooting steps.

Error/issue: File not found

When running an eDiscovery search that includes SharePoint Online and One Drive For Business locations, you may receive the error `File Not Found` although the file is located on the site. This error will be in the export warnings and errors.csv or skipped items.csv. This may occur if the file can't be found on the site or if the index is out of date. Here's the text of an actual error (with emphasis added).

> 28.06.2019 10:02:19_FailedToExportItem_Failed to download content. Additional diagnostic info : Microsoft.Office.Compliance.EDiscovery.ExportWorker.Exceptions.ContentDownloadTemporaryFailure: Failed to download from content 6ea52149-91cd-4965-b5bb-82ca6a3ec9be of type Document. Correlation Id: 3bd84722-937b-4c23-b61b-08d6fba9ec32. ServerErrorCode: -2147024894 ---> Microsoft.SharePoint.Client.ServerException: ***File Not Found***. at Microsoft.SharePoint.Client.ClientRequest.ProcessResponseStream(Stream responseStream) at Microsoft.SharePoint.Client.ClientRequest.ProcessResponse() --- End of inner exception stack trace ---

Resolution

1. Check location identified in the search to ensure the that the location of the file is correct and added in the search locations.
2. Use the procedures at Manually request crawling and re-indexing of a site, a library, or a list to reindex the site.

Error/issue: Search fails because recipient is not found

An eDiscovery search fails with error the `recipient not found`. This error may occur if the user object cannot be found in Exchange Online Protection (EOP) because the object has not synced.

Resolution

1. Connect to Exchange Online PowerShell.
2. Run the following command to check if the user is synced to Exchange Online Protection:

```
Get-Recipient <userId> | FL
```

3. There should be a mail user object for the user question. If nothing is returned, investigate the user object. Contact Microsoft Support if the object can't be synced.

Error/issue: Exporting search results is slow

When exporting search results from eDiscovery or Content Search in the Security and Compliance center, the download takes longer than expected. You can check to see the amount of data to be download and possibly increase the export speed.

Resolution

1. Connect to Security & Compliance Center PowerShell and then run the following command:

```
Get-ComplianceSearch <searchname> | FL
```

2. Find the amount of data to be downloaded in the SearchResults and SearchStatistics parameters.
3. Run the following command:

```
Get-ComplianceSearchAction | FL
```

4. In the results field, find the data that has been exported and view any errors encountered.
5. Check the trace.log file located in the directory that you exported the content to for any errors.
6. If you still have issues, consider dividing searches that return a large set of results into smaller searches. For example, you can use date ranges in search queries to return a smaller set of results that can be downloaded faster.

Error/issue: "Internal server error (500) occurred"

When running an eDiscovery search, if the search continually fails with error similar to "Internal server error (500) occurred", you may need rerun the search only on specific mailbox locations.

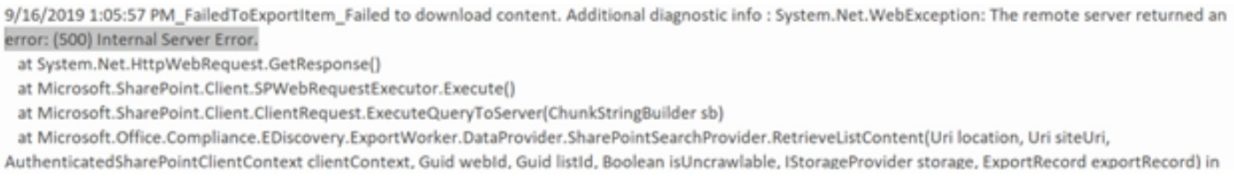

Resolution

1. Break the search into smaller searches and run the search again. Try using a smaller date range or limit the number of locations being searched.
2. Connect to Security & Compliance Center PowerShell and then run the following command:

```
Get-ComplianceSearch <searchname> | FL
```

3. Examine the output for results and errors.
4. Examine the trace.log file. It's located in the same folder that you exported the search results to.
5. Contact Microsoft Support.

Error/issue: Holds don't sync

eDiscovery Case Hold Policy Sync Distribution error. The error reads:

> "Resources: It's taking longer than expected to deploy the policy. It might take an additional 2 hours to update the final deployment status, so check back in a couple hours."

Resolution

1. Connect to Security & Compliance Center PowerShell and then run the following command for an eDiscovery case hold:

   ```
   Get-CaseHoldPolicy <policyname> - DistributionDetail | FL
   ```

   For a retention policy, run the following command:

   ```
   Get-RetentionCompliancePolicy <policyname> - DistributionDetail | FL
   ```

2. Examine the value in the DistributionDetail parameter for errors like the following:

   > Error: Resources: It's taking longer than expected to deploy the policy. It might take an additional 2 hours to update the final deployment status, so check back in a couple hours."

3. Try running the RetryDistribution parameter on the policy in question:

   For eDiscovery case holds:

   ```
   Set-CaseHoldPolicy <policyname> -RetryDistribution
   ```

   For retention policies:

   ```
   Set-RetentionCompliancePolicy <policyname> -RetryDistribution
   ```

4. Contact Microsoft Support.

Error: "The condition specified using HTTP conditional header(s) is not met"

When downloading search results using the eDiscovery Export Tool, it's possible you might receive the following error:

```
System.Net.WebException: The remote server returned an error: (412) The condition specified using HTTP conditional header(s) is not met.
```

This is transient error, which typically occurs in the Azure Storage location.

Resolution

To resolve this issue, retry downloading the search results, which will restart the eDiscovery Export Tool.

Error/issue: Downloaded export shows no results

After a successful export, the completed download via the export tool shows zero files in the results.

Resolution

This is a client-side issue and in order to remediate it, please attempt the following steps:

1. Try using another client/machine to download.
2. Make sure to download to a local drive.
3. Make sure the virus scanner is not running.
4. Make sure that no other export is downloading to the same folder or any parent folder.

5. If the previous steps did not work, disable zipping and de-duplication.
6. If this works then the issue is due to a local virus scanner or a disk issue.

Retry a Content Search to resolve a content location error

11/2/2020 • 2 minutes to read • Edit Online

When you use Content Search in the security and compliance center to search a large number of mailboxes, you may get search errors that are similar to the error:

```
Error


The search on the following locations failed:

User1@contoso.com: Problem in processing the request. Please try again later. If you keep getting this
error, contact your admin. (CS008-009)

User2@contoso.com: Application error occurred. Please try again later. (CS012-002)
```

These errors (with error codes of CS001-002, CS003-002, CS008-009, CS012-002, and other errors of the form CS0XX-0XX) indicate that Content Search failed to search specific content locations; in this example, two mailboxes weren't searched. These errors are displayed on the status details flyout page of the Content Search.

Cause of content location errors

When searching a large number of mailboxes, the search is distributed across thousands of servers in a Microsoft datacenter. At any one time, specific servers could be in reboot state or in the process of failing over to redundant copies. In either of these cases, the Content Search's request to retrieve data will time out. In the previous example, the errors for the mailboxes that failed were the result of the search timing out.

Resolving content location errors

Restarting the search will often result in similar errors on different servers. Instead of restarting the search, click the **Retry** button that is displayed at the top of the search results page.

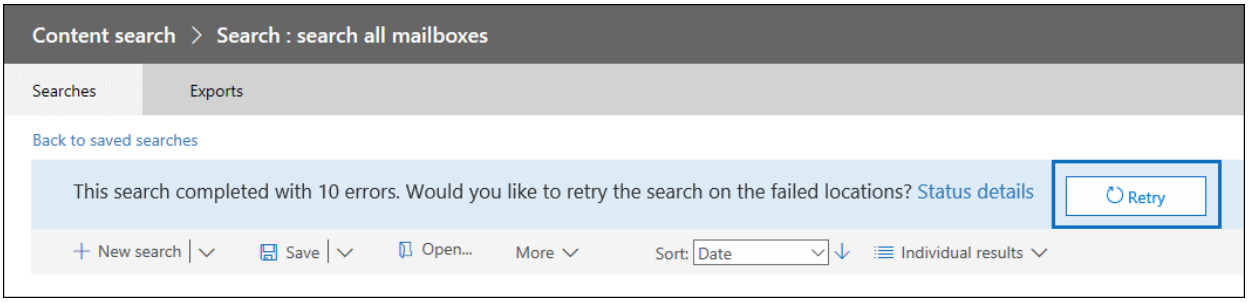

This will result in the retrying the search only for the mailboxes that failed. When you retry the search, the other results that were successfully returned are retained.

Tips to avoid content location errors

Here are some additional causes of content location errors and some tips to help you avoid them when searching large numbers of mailboxes.

- The mailbox being searched might be busy due to user activity. In this case, the search service might

throttle itself to prevent the mailbox from becoming unavailable. To avoid this, try running searches during non-business hours.

- The search query might be retrieving too much content from the mailbox. If possible, try to narrow the scope of the search by using keywords, date ranges, and search conditions.
- Too many keywords or keyword phrases when you create a search query using the keywords list. When you run a search query that uses the keywords list, the service essentially runs a separate search for each row in the keyword list so that statistics can be generated. If you're using the keywords list in search queries, minimize the number of rows in the keyword list or divide the number keywords into smaller lists and create a different search for each keyword list.

  > **NOTE**
  > To help reduce issues caused by large keyword lists, you're now limited to a maximum of 20 rows in the keyword list of a search query.

- Too many searches are being performed on the same mailbox at the same time. If possible, try to run one search at a time on any one mailbox.
- Searching too many mailboxes in a single search. The probability of content location errors increases when searching a large number of mailboxes. If possible, try to run multiple searches so that each search includes a subset of mailboxes in your organization.
- Required maintenance is being performed on the mailbox. Though this cause probably occurs infrequently, wait a little while after receiving the content location error and then retry the search.

Preserve Bcc and expanded distribution group recipients for eDiscovery

2/18/2021 • 5 minutes to read • Edit Online

In-Place Hold, Litigation Hold, and Microsoft 365 retention policies (created in the Security & Compliance Center) allow you to preserve mailbox content to meet regulatory compliance and eDiscovery requirements. Information about recipients directly addressed in the To and Cc fields of a message is included in all messages by default. But your organization may require the ability to search for and reproduce details about all recipients of a message. This includes:

- **Recipients addressed using the Bcc field of a message:** Bcc recipients are stored in the message in the sender's mailbox, but not included in headers of the message delivered to recipients.
- **Expanded distribution group recipients:** Recipients who receive the message because they're members of a distribution group to which the message was addressed, either in the To, Cc or Bcc fields.

Exchange Online and Exchange Server 2013 (Cumulative Update 7 and later versions) retain information about Bcc and expanded distribution group recipients. You can search for this information by using an In-Place eDiscovery search in the Exchange admin center (EAC) or a Content Search in the Security & Compliance Center.

How Bcc recipients and expanded distribution group recipients are preserved

As stated earlier, information about Bcc'ed recipients is stored with the message in the sender's mailbox. This information is indexed and available to eDiscovery searches and holds.

Information about expanded distribution group recipients is stored with the message after you place a mailbox on In-Place Hold or Litigation Hold. In Office 365, this information is also stored when a Microsoft 365 retention policy is applied to a mailbox. Distribution group membership is determined at the time the message is sent. The expanded recipients list stored with the message is not impacted by changes to membership of the group after the message is sent.

INFORMATION ABOUT...	IS STORED IN...	IS STORED BY DEFAULT?	IS ACCESSIBLE TO...
To and Cc recipients	Message properties in the sender and recipients' mailboxes.	Yes	Sender, recipients, and compliance officers
Bcc recipients	Message property in the sender's mailbox.	Yes	Sender and compliance officers
Expanded distribution group recipients	Message properties in the sender's mailbox.	No. Expanded distribution group recipient information is stored after a mailbox is placed on In-Place Hold or Litigation Hold, or assigned to a Microsoft 365 retention policy.	Compliance officers

Searching for messages sent to Bcc and expanded distribution group

recipients

When searching for messages sent to a recipient, eDiscovery search results now include messages sent to a distribution group that the recipient is a member of. The following table shows the scenarios where messages sent to Bcc and expanded distribution group recipients are returned in eDiscovery searches.

Scenario 1: John is a member of the US-Sales distribution group. This table shows eDiscovery search results when Bob sends a message to John directly or indirectly via a distribution group.

WHEN YOU SEARCH BOB'S MAILBOX FOR MESSAGES SENT...	AND THE MESSAGE IS SENT WITH...	RESULTS INCLUDE MESSAGE?
To:John	John on TO	Yes
To:John	US-Sales on TO	Yes
To:US-Sales	US-Sales on TO	Yes
Cc:John	John on CC	Yes
Cc:John	US-Sales on CC	Yes
Cc:US-Sales	US-Sales on CC	Yes

Scenario 2: Bob sends an email to John (To/Cc) and Jack (Bcc directly, or indirectly via a distribution group). The table below shows eDiscovery search results.

WHEN YOU SEARCH...	FOR MESSAGES SENT...	RESULTS INCLUDE MESSAGE?	NOTES
Bob's mailbox	To/Cc:John	Yes	Presents an indication that Jack was Bcc'ed.
Bob's mailbox	Bcc:Jack	Yes	Presents an indication that Jack was Bcc'ed.
Bob's mailbox	Bcc:Jack (via distribution group)	Yes	List of members of the Bcc'ed distribution group, expanded when the message was sent, is visible in eDiscovery search preview, export, and logs.
John's mailbox	To/Cc:John	Yes	No indication of Bcc recipients.
John's mailbox	Bcc:Jack (directly or via distribution group)	No	Bcc information is not stored in the message delivered to recipients. You must search the sender's mailbox.
Jack's mailbox	To/Cc:John (directly or via distribution group)	Yes	To/Cc information is included in message delivered to all recipients.

WHEN YOU SEARCH...	FOR MESSAGES SENT...	RESULTS INCLUDE MESSAGE?	NOTES
Jack's mailbox	Bcc:Jack (directly or via distribution group)	No	Bcc information is not stored in the message delivered to recipients. You must search the sender's mailbox.

Frequently asked questions

Q. When and where is Bcc recipient information stored?

A. Bcc recipient information is preserved by default in the original message in sender's mailbox. If the Bcc recipient is a distribution group, distribution group membership is only expanded if the sender's mailbox is on hold or assigned to a Microsoft 365 retention policy.

Q. When and where is the list of expanded distribution group recipients stored?

A. Group membership is expanded at the time the message is sent. The list of expanded distribution group members is stored in the original message in the sender's mailbox. The sender's mailbox must be on In-Place Hold, Litigation Hold, or assigned to a Microsoft 365 retention policy.

Q. Can the To/Cc recipients see which recipients were Bcc'ed?

A. No. This information is not included in message headers, and isn't visible to To/Cc recipients. The sender can see the Bcc field stored in the original message stored in their mailbox. Compliance officers can see this information when searching the sender's mailbox.

Q. How can I ensure that expanded distribution group recipients are always preserved?

A. To ensure that expanded distribution group members are always preserved with a message, Place all mailboxes on hold or create an organization-wide Microsoft 365 retention policy.

Q. Which types of groups are supported?

A. Distribution groups, mail-enabled security groups, and dynamic distribution groups are supported.

Q. Is there a limit on the number of distribution group recipients that are expanded and stored in the message?

A. Up to 10,000 members of a distribution group is preserved.

Q. Are nested distribution groups supported?

A. Yes, 25 levels of nested distribution groups are expanded.

Q. Where is the Bcc and expanded distribution group recipient information visible?

A. Bcc and expanded distribution group recipients information is visible to Compliance officers when performing an eDiscovery search. Bcc and expanded distribution group recipients are included in search results copied to a Discovery mailbox or exported to a PST file and in the eDiscovery log included in search results. Bcc recipient information is also available in search preview.

Q. What happens if a member of a distribution group is hidden from the organization's global address list (GAL)?

A. There's no impact. If recipients are hidden from the GAL, they are still included in the list of recipients for the expanded distribution group.

Configure permissions filtering for Content Search

2/18/2021 • 18 minutes to read • Edit Online

You can use search permissions filtering to let an eDiscovery manager search only a subset of mailboxes and sites in your organization. You can also use permissions filtering to let that same eDiscovery manager search only for mailbox or site content that meets a specific search criteria. For example, you might let an eDiscovery manager search only the mailboxes of users in a specific location or department. You do this by creating a filter that uses a supported recipient filter to limit which mailboxes a specific user or group of users can search. You can also create a filter that specifies what mailbox content a user can search for. This is done by creating a filter that uses a searchable message property. Similarly, you can let an eDiscovery manager search only specific SharePoint sites in your organization. You do this by creating a filter that limits which site can be searched. You can also create a filter that specifies what site content can be searched. This is done by creating a filter that uses a searchable site property.

You can also use search permissions filtering to create logical boundaries (called *compliance boundaries*) within an organization that control the user content locations (such as mailboxes, SharePoint sites, and OneDrive accounts) that specific eDiscovery managers can search. For more information, see Set up compliance boundaries for eDiscovery investigations in Office 365.

Search permissions filtering is supported by the Content Search feature in the Security & Compliance Center. These four cmdlets let you configure and manage search permissions filters:

New-ComplianceSecurityFilter

Get-ComplianceSecurityFilter

Set-ComplianceSecurityFilter

Remove-ComplianceSecurityFilter

Requirements to configure permissions filtering

- To run the compliance security filter cmdlets, you have to be a member of the Organization Management role group in the Security & Compliance Center. For more information, see Permissions in the Security & Compliance Center.
- You have to connect to both Exchange Online and Security & Compliance Center PowerShell to use the compliance security filter cmdlets. This is necessary because these cmdlets require access to mailbox properties, which is why you have to connect to Exchange Online PowerShell. See the steps in the next section.
- See the More information section for additional information about search permissions filters.
- Search permissions filtering is applicable to inactive mailboxes, which means you can use mailbox and mailbox content filtering to limit who can search an inactive mailbox. See the More information section for additional information about permissions filtering and inactive mailboxes.
- Search permissions filtering can't be used to limit who can search public folders in Exchange.
- There is no limit to the number of search permissions filters that can be created in an organization. But search performance will be impacted when there are more than 100 search permissions filters. To keep the number of search permissions filters in your organization as small as possible, create filters that combine rules for Exchange, SharePoint, and OneDrive in a single filter whenever possible.

Connect to Exchange Online and Security & Compliance Center PowerShell in a single session

Before you can successfully run the script in this section, you have to download and install the Exchange Online PowerShell V2 module. For information, see About the Exchange Online PowerShell V2 module.

1. Save the following text to a Windows PowerShell script file by using a filename suffix of **.ps1**. For example, you could save it to a file named **ConnectEXO-SCC.ps1**.

```
Import-Module ExchangeOnlineManagement
$UserCredential = Get-Credential
Connect-ExchangeOnline -Credential $UserCredential -ShowBanner:$false
Connect-IPPSSession -Credential $UserCredential
$Host.UI.RawUI.WindowTitle = $UserCredential.UserName + " (Exchange Online + Compliance Center)"
```

2. On your local computer, open Windows PowerShell, go to the folder where the script that you created in the previous step is located, and then run the script; for example:

```
.\ConnectEXO-SCC.ps1
```

How do you know if this worked? After you run the script, cmdlets from Exchange Online and Security & Compliance PowerShell are imported to your local Windows PowerShell session. If you don't receive any errors, you connected successfully. A quick test is to run an Exchange Online and Security & Compliance Center cmdlet. For example, you can run and **Get-Mailbox** and **Get-ComplianceSearch**.

For troubleshooting PowerShell connection errors, see:

- Connect to Exchange Online PowerShell
- Connect to Security & Compliance Center PowerShell

New-ComplianceSecurityFilter

The **New-ComplianceSecurityFilter** is used to create a search permissions filter. The following table describes the parameters for this cmdlet. All parameters are required to create a compliance security filter.

PARAMETER	DESCRIPTION
Action	The *Action* parameter specifies that type of search action that the filter is applied to. The possible Content Search actions are: **Export:** The filter is applied when exporting search results. **Preview:** The filter is applied when previewing search results. **Purge:** The filter is applied when purging search results. **Search:** The filter is applied when running a search. **All:** The filter is applied to all search actions.
FilterName	The *FilterName* parameter specifies the name of the permissions filter. This name is used to identity a filter when using the **Get-ComplianceSecurityFilter**, **Set-ComplianceSecurityFilter,** and **Remove-ComplianceSecurityFilter** cmdlets.
Filters	The *Filters* parameter specifies the search criteria for the compliance security filter. You can create three different types

PARAMETER	DESCRIPTION

	of filters:
	Mailbox filtering: This type of filter specifies the mailboxes the assigned users (specified by the *Users* parameter) can search. The syntax for this type of filter is **Mailbox_** *MailboxPropertyName*, where *MailboxPropertyName* specifies a mailbox property used to scope the mailboxes that can be searched. For example, the mailbox filter `"Mailbox_CustomAttribute10 -eq 'OttawaUsers'"` would allow the user assigned this filter to search only the mailboxes that have the value "OttawaUsers" in the CustomAttribute10 property.
	Any supported filterable recipient property can be used for the *MailboxPropertyName* property. For a list of supported properties, see Filterable properties for the -RecipientFilter parameter.
	
	Mailbox content filtering: This type of filter is applied on the content that can be searched. It specifies the mailbox content the assigned users can search for. The syntax for this type of filter is **MailboxContent_** *SearchablePropertyName:value*, where *SearchablePropertyName* specifies a Keyword Query Language (KQL) property that can be specified in a Content Search. For example, the mailbox content filter `MailboxContent_recipients:contoso.com` would allow the user assigned this filter to only search for messages sent to recipients in the contoso.com domain.
	For a list of searchable message properties, see Keyword queries and search conditions for Content Search.
	
	Important: A single search filter can't contain a mailbox filter and a mailbox content filter. To combine these in a single filter, you have to use a filters list. But a filter can contain a more complex query of the same type. For example,
	`"Mailbox_CustomAttribute10 -eq 'FTE' -and Mailbox_MemberOfGroup -eq '$($DG.DistinguishedName)'"`
	
	Site and site content filtering: There are two SharePoint and OneDrive for Business site-related filters that you can use to specify what site or site content the assigned users can search:
	
	- **Site_** *SearchableSiteProperty*
	- **SiteContent_** *SearchableSiteProperty*
	
	These two filters are interchangeable. For example,
	`"Site_Path -like 'https://contoso.sharepoint.com/sites/doctors*'"`
	and
	`"SiteContent_Path -like 'https://contoso.sharepoint.com/sites/doctors*'"`
	return the same results. But to help you identify what a filter does, you can use `Site_` to specify site-related properties (such as a site URL) and `SiteContent_` to specify content-related properties (such as document types. For example, the filter
	`"Site_Path -like 'https://contoso.sharepoint.com/sites/doctors*'"`
	would allow the user assigned this filter to only search for content in the https://contoso.sharepoint.com/sites/doctors site collection. The filter
	`"SiteContent_FileExtension -eq 'docx'"` would allow the user assigned this filter to only search for Word documents (Word 2007 and later).

PARAMETER	DESCRIPTION
	For a list of searchable site properties, see Overview of crawled and managed properties in SharePoint. Properties marked with a **Yes** in the **Queryable** column can be used to create a site or site content filter. **Important:** You have to create a search permissions filter to explicitly prevent users from searching content locations in a specific service (such as preventing a user from searching any Exchange mailbox or any SharePoint site). In other words, creating a search permissions filter that allows a user to search all SharePoint sites in the organization doesn't prevent that user from searching mailboxes. For example, to allow SharePoint admins to only search SharePoint sites, you have to create a filter that prevents them from searching mailboxes. Similarly, to allow Exchange admins to only search mailboxes, you have to create a filter that prevents them from searching sites.
Users	The *Users* parameter specifies the users who get this filter applied to their Content Searches. Identify users by their alias or primary SMTP address. You can specify multiple values separated by commas, or you can assign the filter to all users by using the value **All**. You can also use the *Users* parameter to specify a Security & Compliance Center role group. This lets you create a custom role group and then assign that role group a search permissions filter. For example, let's say you have a custom role group for eDiscovery managers for the U.S. subsidiary of a multi-national corporation. You can use the *Users* parameter to specify this role group (by using the Name property of the role group) and then use the *Filter* parameter to allow only mailboxes in the U.S. to be searched. You can't specify distribution groups with this parameter.

Using a filters list to combine filter types

A *filters list* is a filter that includes a mailbox filter and a site filter separated by a comma. Using a filters list is the only supported method for combining different types of filters. In the following example, notice that a comma separates the **Mailbox** and **Site** filters:

```
-Filters "Mailbox_CustomAttribute10 -eq 'OttawaUsers'", "Site_Path -like
'https://contoso.sharepoint.com/sites/doctors*'"
```

When a filter that contains a filters list is processed during the running of a content search, two search permissions filters are created from the filters list: One for each filter that's separated by a comma. So in the previous example, one mailbox search permissions filter and one site search permissions filter would be created.

An alternative to using a filters list would be to create two separate search permissions filters. So in the previous example, you'd create one filter for the mailbox attribute and one filter for the site attribute. In either case, the results are the same. Using a filters list or creating separate search permissions filters is a matter of preference.

Keep the following things in mind about using a filters list:

- You have to use a filters list to create a filter that includes a **Mailbox** filter and a **MailboxContent** filter.
- As previously suggested, you don't have to use a filters list to include a **Site** and a **SiteContent** filter in a single search permissions filter. For example, you can combine **Site** and a **SiteContent** filters using an **-or** operator.

```
-Filters "Site_ComplianceAttribute -eq 'FourthCoffee' -or Site_Path -like
'https://contoso.sharepoint.com/sites/FourthCoffee*'"
```

- Each component of a filters list can contain a complex filter syntax. For example, the mailbox and site filters can contain multiple filters separated by an **-or** operator:

```
-Filters "Mailbox_Department -eq 'CohoWinery' -or Mailbox_CustomAttribute10 -eq 'CohoUsers'",
"Site_ComplianceAttribute -eq 'CohoWinery' -or Site_Path -like
'https://contoso.sharepoint.com/sites/CohoWinery*'"
```

Examples of creating search permissions filters

Here are examples of using the **New-ComplianceSecurityFilter** cmdlet to create a search permissions filter.

This example allows the user annb@contoso.com to perform all Content Search actions only for mailboxes in Canada. This filter contains the three-digit numeric country code for Canada from ISO 3166-1.

```
New-ComplianceSecurityFilter -FilterName CountryFilter  -Users annb@contoso.com -Filters
"Mailbox_CountryCode  -eq '124'" -Action All
```

This example allows the users donh and suzanf to search only the mailboxes that have the value 'Marketing' for the CustomAttribute1 mailbox property.

```
New-ComplianceSecurityFilter -FilterName MarketingFilter  -Users donh,suzanf -Filters
"Mailbox_CustomAttribute1  -eq 'Marketing'" -Action Search
```

This example allows members of the "US Discovery Managers" role group to perform all Content Search actions only on mailboxes in the United States. This filter contains the three-digit numeric country code for the United States from ISO 3166-1.

```
New-ComplianceSecurityFilter -FilterName USDiscoveryManagers  -Users "US Discovery Managers" -Filters
"Mailbox_CountryCode  -eq '840'" -Action All
```

This example allows members of the eDiscovery Manager role group to search only the mailboxes of members of the Ottawa Users distribution group. The Get-DistributionGroup cmdlet in Exchange Online PowerShell is used to find the members of the Ottawa Users group.

```
$DG = Get-DistributionGroup "Ottawa Users"
```

```
New-ComplianceSecurityFilter -FilterName DGFilter  -Users eDiscoveryManager -Filters "Mailbox_MemberOfGroup
-eq '$($DG.DistinguishedName)'" -Action Search
```

This example prevents any user from deleting content from the mailboxes of members of the Executive Team distribution group. The Get-DistributionGroup cmdlet in Exchange Online PowerShell is used to find the members of the Executive Team group.

```
$DG = Get-DistributionGroup "Executive Team"
```

```
New-ComplianceSecurityFilter -FilterName NoExecutivesPreview  -Users All -Filters "Mailbox_MemberOfGroup -ne
'$($DG.DistinguishedName)'" -Action Purge
```

This example allows members of the OneDrive eDiscovery Managers custom role group to only search for content in OneDrive for Business locations in the organization.

```
New-ComplianceSecurityFilter -FilterName OneDriveOnly  -Users "OneDrive eDiscovery Managers" -Filters
"Site_Path -like 'https://contoso-my.sharepoint.com/personal*'" -Action Search
```

NOTE

To restrict users to searching specific sites, use the filter `Site_Path`, as shown in the previous example. Using `Site_Site` will not work.

This example restricts the user to performing all Content Search actions only on email messages sent during the calendar year 2015.

```
New-ComplianceSecurityFilter -FilterName EmailDateRestrictionFilter -Users donh@contoso.com -Filters
"MailboxContent_Received -ge '01-01-2015' -and MailboxContent_Received -le '12-31-2015'" -Action All
```

Similar to the previous example, this example restricts the user to performing all Content Search actions on documents that were last changed sometime in the calendar year 2015.

```
New-ComplianceSecurityFilter -FilterName DocumentDateRestrictionFilter -Users donh@contoso.com -Filters
"SiteContent_LastModifiedTime -ge '01-01-2015' -and SiteContent_LastModifiedTime -le '12-31-2015'" -Action
All
```

This example prevents members of the "OneDrive Discovery Managers" role group from performing content search actions on any mailbox in the organization.

```
New-ComplianceSecurityFilter -FilterName NoEXO -Users "OneDrive Discovery Managers" -Filters "Mailbox_Alias
-notlike '*'"  -Action All
```

This example prevents anyone in the organization from searching for email messages that were sent or received by janets or sarad.

```
New-ComplianceSecurityFilter -FilterName NoSaraJanet -Users All -Filters "MailboxContent_Participants -
notlike 'janets@contoso.onmicrosoft.com' -and MailboxContent_Participants -notlike
'sarad@contoso.onmicrosoft.com'" -Action Search
```

This example uses a filters list to combine mailbox and site filters.

```
New-ComplianceSecurityFilter -FilterName "Coho Winery Security Filter" -Users "Coho Winery eDiscovery
Managers", "Coho Winery Investigators" -Filters "Mailbox_Department -eq 'CohoWinery'",
"Site_ComplianceAttribute -eq 'CohoWinery' -or Site_Path -like
'https://contoso.sharepoint.com/sites/CohoWinery*'" -Action ALL
```

Get-ComplianceSecurityFilter

The **Get-ComplianceSecurityFilter** is used to return a list of search permissions filters. Use the *FilterName*

parameter to return information for a specific search filter.

Set-ComplianceSecurityFilter

The **Set-ComplianceSecurityFilter** is used to modify an existing search permissions filter. The only required parameter is *FilterName*.

PARAMETER	DESCRIPTION
Action	The *Action* parameter specifies that type of search action that the filter is applied to. The possible Content Search actions are: **Export:** The filter is applied when exporting search results. **Preview:** The filter is applied when previewing search results. **Purge:** The filter is applied when purging search results. **Search:** The filter is applied when running a search. **All:** The filter is applied to all search actions.
FilterName	The *FilterName* parameter specifies the name of the permissions filter.
Filters	The *Filters* parameter specifies the search criteria for the compliance security filter. You can create two different types of filters: **Mailbox filtering:** This type of filter specifies the mailboxes the assigned users (specified by the *Users* parameter) can search. The syntax for this type of filter is **Mailbox_** *MailboxPropertyName*, where *MailboxPropertyName* specifies a mailbox property used to scope the mailboxes that can be searched. For example, the mailbox filter `"Mailbox_CustomAttribute10 -eq 'OttawaUsers'"` would allow the user assigned this filter to search only the mailboxes that have the value "OttawaUsers" in the CustomAttribute10 property. Any supported filterable recipient property can be used for the *MailboxPropertyName* property. For a list of supported properties, see Filterable properties for the -RecipientFilter parameter. **Mailbox content filtering:** This type of filter is applied on the content that can be searched. It specifies the mailbox content the assigned users can search for. The syntax for this type of filter is **MailboxContent_** *SearchablePropertyName:value*, where *SearchablePropertyName* specifies a Keyword Query Language (KQL) property that can be specified in a Content Search. For example, the mailbox content filter `MailboxContent_recipients:contoso.com` would allow the user assigned this filter to only search for messages sent to recipients in the contoso.com domain. For a list of searchable message properties, see Keyword queries for Content Search. **Site and site content filtering:** There are two SharePoint and OneDrive for Business site-related filters that you can use to specify what site or site content the assigned users can search: - **Site_** *SearchableSiteProperty* - **SiteContent_** *SearchableSiteProperty*

PARAMETER	DESCRIPTION
	- SiteContent_*SearchableSiteProperty* These two filters are interchangeable. For example, `"Site_Path -like 'https://contoso.spoppe.com/sites/doctors*'"` and `"SiteContent_Path -like 'https://contoso.spoppe.com/sites/doctors*'"` returns the same results. But to help you identify what a filter does, you can use `Site_` to specify site-related properties (such as a site URL) and `SiteContent_` to specify content-related properties (such as document types. For example, the filter `"Site_Path -like 'https://contoso.spoppe.com/sites/doctors*'"` would allow the user assigned this filter to only search for content in the https://contoso.spoppe.com/sites/doctors site collection. The filter `"SiteContent_FileExtension -eq 'docx'"` would allow the user assigned this filter to only search for Word documents (Word 2007 and later). For a list of searchable site properties, see Overview of crawled and managed properties in SharePoint. Properties marked with a **Yes** in the **Queryable** column can be used to create a site or site content filter.
Users	The *Users* parameter specifies the users who get this filter applied to their Content Searches. Because this is a multi-value property, specifying a user or group of users with this parameter overwrite the existing list of users. See the following examples for the syntax to add and remove selected users. You can also use the *Users* parameter to specify a Security & Compliance Center role group. This lets you create a custom role group and then assign that role group a search permissions filter. For example, let's say you have a custom role group for eDiscovery managers for the U.S. subsidiary of a multi-national corporation. You can use the *Users* parameter to specify this role group (by using the Name property of the role group) and then use the *Filter* parameter to allow only mailboxes in the U.S. to be searched. You can't specify distribution groups with this parameter.

Examples of changing search permissions filters

These examples show how to use the **Get-ComplianceSecurityFilter** and **Set-ComplianceSecurityFilter** cmdlets to add or remove a user to the existing list of users that the filter is assigned to. When you add or remove users from a filter, specify the user by using their SMTP address.

This example adds a user to the filter.

```
$filterusers = Get-ComplianceSecurityFilter -FilterName OttawaUsersFilter
```

```
$filterusers.users.add("pilarp@contoso.com")
```

```
Set-ComplianceSecurityFilter -FilterName OttawaUsersFilter -Users $filterusers.users
```

This example removes a user from the filter.

```
$filterusers = Get-ComplianceSecurityFilter -FilterName OttawaUsersFilter
```

```
$filterusers.users.remove("annb@contoso.com")
```

```
Set-ComplianceSecurityFilter -FilterName OttawaUsersFilter -Users $filterusers.users
```

Remove-ComplianceSecurityFilter

The **Remove-ComplianceSecurityFilter** is used to delete a search filter. Use the *FilterName* parameter to specify the filter you want to delete.

More information

- **How does search permissions filtering work?** The permissions filter is added to the search query when a Content Search is run. The permissions filter is joined to the search query by the **AND** Boolean operator. For example, you have a permissions filter that allows Bob to perform all search actions on the mailboxes of members of the Workers distribution group. Then Bob runs a Content Search on all mailboxes in the organization with the search query `sender:jerry@adatum.com`. Because the permissions filter and the search query are logically combined by an **AND** operator, the search returns any message sent by jerry@adatum.com to any member of the Workers distribution group.

- **What happens if you have multiple search permissions filters?** 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. Let's take the previous example, where a search filter allows Bob to search only the mailboxes of the members of the Workers distribution group. Then we create another filter that prevents Bob from searching Phil's mailbox ("Mailbox_Alias -ne 'Phil'"). And let's also assume that Phil is a member of the Workers group. When Bob runs a Content Search (from the previous example) on all mailboxes in the organization, search results are returned for Phil's mailbox even though you applied filter to prevent Bob from searching Phil's mailbox. This is because the first filter, which allows Bob to search the Workers group, is true. And because Phil is a member of the Workers group, Bob can search Phil's mailbox.

- **Does search permissions filtering work for inactive mailboxes?** Yes, you can use mailbox and mailbox content filters to limit who can search inactive mailboxes in your organization. Like a regular mailbox, an inactive mailbox has to be configured with the recipient property that's used to create a permissions filter. If necessary, you can use the **Get-Mailbox -InactiveMailboxOnly** command to display the properties of inactive mailboxes. For more information, see Create and manage inactive mailboxes in Office 365.

- **Does search permissions filtering work for public folders?** No. As previously explained, search permissions filtering can't be used to limit who can search public folders in Exchange. For example, items in public folder locations can't be excluded from the search results by a permissions filter.

- **Does allowing a user to search all content locations in a specific service also prevent them from searching content locations in a different service?** No. As previously explained, you have to create a search permissions filter to explicitly prevent users from searching content locations in a specific

service (such as preventing a user from searching any Exchange mailbox or any SharePoint site). In other words, creating a search permissions filter that allows a user to search all SharePoint sites in the organization doesn't prevent that user from searching mailboxes. For example, to allow SharePoint admins to only search SharePoint sites, you have to create a filter that prevents them from searching mailboxes. Similarly, to allow Exchange admins to only search mailboxes, you have to create a filter that prevents them from searching sites.

Change the size of PST files when exporting eDiscovery search results

4/30/2020 • 2 minutes to read • Edit Online

When you use the eDiscovery Export tool to export the email results of an eDiscovery search from the different Microsoft eDiscovery tools, the default size of a PST file that can be exported is 10 GB. If you want to change this default size, you can edit the Windows Registry on the computer that you use to export the search results. One reason to do this is so a PST file can fit on removable media, such a DVD, a compact disc, or a USB drive.

> **NOTE**
> The eDiscovery Export tool is used to export the search results when using the Content Search tool in the Security & Compliance Center, In-Place eDiscovery in Exchange Online, and the eDiscovery Center in SharePoint Online.

Create a registry setting to change the size of PST files when you export eDiscovery search results

Perform the following procedure on the computer that you'll use to export the results of an eDiscovery search.

1. Close the eDiscovery Export tool if it's open.
2. Save the following text to a Window registry file by using a filename suffix of .reg; for example, PstExportSize.reg.

   ```
   Windows Registry Editor Version 5.00
   [HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Exchange\Client\eDiscovery\ExportTool]
   "PstSizeLimitInBytes"="1073741824"
   ```

   In the example above, the `PstSizeLimitInBytes` value is set to 1,073,741,824 bytes or approximately 1 GB. Here are some other sample values for the `PstSizeLimitInBytes` setting.

SIZE IN GB (APPROX.)	SIZE IN BYTES
0.7 GB (700 MB)	751619277
2 GB	2147483648
4 GB	4294967296
8 GB	8589934592

3. Change the `PstSizeLimitInBytes` value to the desired maximum size of a PST file when you export search results, and then save the file.
4. In Windows Explorer, click or double-click the .reg file that you created in the previous steps.
5. In the User Access Control window, click **Yes** to let the Registry Editor make the change.
6. When prompted to continue, click **Yes**.

The Registry Editor displays a message saying that the setting was successfully added to the registry.

7. You can repeat steps 3 - 6 to change the value for the `PstSizeLimitInBytes` registry setting.

Frequently asked questions about changing the default size of PST files when you export eDiscovery search results

Why is the default size 10 GB?

The default size of 10 GB was based on customer feedback; 10 GB is a good balance between the optimal amount of content in a single PST and with a minimum chance of file corruption.

Should I increase or decrease the default size of PST files?

Customers tend to decrease the size limit so that the search results will fit on removable media that they can physically ship to other locations in their organization. We don't recommend that you increase the default size because PST files larger than 10 GB might have corruption issues.

What computer do I have to do this on?

You need to change the registry setting on any local computer that you run the eDiscovery Export tool on.

After I change this setting, do I have to reboot the computer?

No, you don't have to reboot the computer. But, if the eDiscovery Export tool is running, you'll have to close it and the restart it after you change this setting.

Does an existing registry key get edited or does a new key get created?

A new registry key is created the first time you run the .reg file that you created in this procedure. Then the setting is edited each time you change and rerun the .reg edit file.

Disable reports when you export Content Search results

11/2/2020 • 4 minutes to read • Edit Online

When you use the eDiscovery Export tool to export the results of a Content Search in the Security & Compliance Center, the tool automatically creates and exports two reports that contain additional information about the exported content. These reports are the Results.csv file and the Manifest.xml file (see the Frequently asked questions about disabling export reports section in this topic for detailed descriptions of these reports). Because these files can be very large, you can speed up the download time and save disk space by preventing these files from being exported. You can do this by changing the Windows Registry on the computer that you use to export the search results. If you want to include the reports at a later time, you can edit the registry setting.

Create registry settings to disable the export reports

Perform the following procedure on the computer that you'll use to export the results a content search.

1. Close the eDiscovery Export tool if it's open.
2. Perform one or both of the following steps, depending on which export report you want to disable.

   - **Results.csv**

     Save the following text to a Windows registry file by using a filename suffix of .reg; for example, DisableResultsCsv.reg.

     ```
     Windows Registry Editor Version 5.00
     reg add HKLM\SOFTWARE\Microsoft\Exchange\Client\eDiscovery\ExportTool /v ResultCsvEnabled /t
     REG_SZ /d False
     ```

   - **Manifest.xml**

     Save the following text to a Windows registry file by using a filename suffix of .reg; for example, DisableManifestXml.reg.

     ```
     Windows Registry Editor Version 5.00
     reg add HKLM\SOFTWARE\Microsoft\Exchange\Client\eDiscovery\ExportTool /v ResultEdrmEnabled /t
     REG_SZ /d False
     ```

3. In Windows Explorer, click or double-click the .reg file that you created in the previous steps.
4. In the User Access Control window, click **Yes** to let the Registry Editor make the change.
5. When prompted to continue, click **Yes**.

   The Registry Editor displays a message saying that the setting was successfully added to the registry.

Edit registry settings to re-enable the export reports

If you disabled the Results.csv and Manifest.xml reports by creating the .reg files in the previous procedure, you can edit those files to re-enable a report so that it's exported with the search results. Again, perform the following procedure on the computer that you'll use to export the results a content search.

1. Close the eDiscovery Export tool if it's open.
2. Edit one or both of the .reg edit files that you created in the previous procedure.

   - **Results.csv**

     Open the DisableResultsCsv.reg file in Notepad, change the value `False` to `True`, and then save the file. For example, after you edit the file, it looks like this:

     ```
     Windows Registry Editor Version 5.00
     reg add HKLM\SOFTWARE\Microsoft\Exchange\Client\eDiscovery\ExportTool /v ResultCsvEnabled /t
     REG_SZ /d True
     ```

   - **Manifest.xml**

     Open the DisableManifestXml.reg file in Notepad, change the value `False` to `True`, and then save the file. For example, after you edit the file, it looks like this:

     ```
     Windows Registry Editor Version 5.00
     reg add HKLM\SOFTWARE\Microsoft\Exchange\Client\eDiscovery\ExportTool /v ResultEdrmEnabled /t
     REG_SZ /d True
     ```

3. In Windows Explorer, click or double-click a .reg file that you edited in the previous step.
4. In the User Access Control window, click **Yes** to let the Registry Editor make the change.
5. When prompted to continue, click **Yes**.

   The Registry Editor displays a message saying that the setting was successfully added to the registry.

Frequently asked questions about disabling export reports

What are the Results.csv and Manifest.xml reports?

The Results.csv and Manifest.xml files contain additional information about the content that was exported.

- **Results.csv** An Excel document that contains information about each item that is download as a search result. For email, the result log contains information about each message, including:
  - The location of the message in the source mailbox (including whether the message is in the primary or archive mailbox).
  - The date the message was sent or received.
  - The Subject line from the message.
  - The sender and recipients of the message.
  - Whether the message is a duplicate message if you enabled de-duplication when exporting the search results. Duplicate messages will have a value in the **Parent ItemId** column that identifies the message as a duplicate. The value in the **Parent ItemId** column is the same as the value in the **Item DocumentId** column of the message that was exported.

    For documents from SharePoint and OneDrive for Business sites, the result log contains information about each document, including:
  - The URL for the document.
  - The URL for the site collection where the document is located.

- The date that the document was last modified.
- The name of the document (which is located in the Subject column in the result log).

- **Manifest.xml** A manifest file (in XML format) that contains information about each item included in the search results. The information in this report is the same as the Results.csv report, but it's in the format specified by the Electronic Discovery Reference Model (EDRM). For more information about EDRM, go to https://www.edrm.net.

When should I disable exporting these reports?

It depends on your specific needs. Many organizations don't require additional information about search results, and don't need these reports.

What computer do I have to do this on?

You have to change the registry setting on any local computer that you run the eDiscovery Export tool on.

After I change this setting, do I have to restart the computer?

No, you don't have to restart the computer. But if the eDiscovery Export tool is running, you have to close it and then restart it after you change the registry setting.

Does an existing registry key get edited or does a new key get created?

A new registry key is created the first time you run the .reg file that you created in the procedure in this topic. Then the setting is edited each time you change and re-run the .reg edit file.

Limits for Content search

2/18/2021 • 10 minutes to read • Edit Online

Various limits are applied to the Content search tool in the Microsoft 365 compliance center. This includes searches run on the **Content search** page and searches that are associated with an eDiscovery case on the **Core eDiscovery** page. These limits help to maintain the health and quality of services provided to organizations. There are also limits related to the indexing of email messages in Exchange Online for search. You can't modify the limits for Content Search or email indexing, but you should be aware of them so that you can take these limits into consideration when planning, running, and troubleshooting content searches.

Search limits

The following table lists the search limits when using the content search tool in the Microsoft 365 compliance center and for searches that are associated with a Core eDiscovery case.

DESCRIPTION OF LIMIT	LIMIT
The maximum number of mailboxes or sites that can be searched in a single search	No limit [1]
The maximum number of searches that can run at the same time in your organization.	30
The maximum number of searches that a single user can start at the same time. This limit is most likely hit when the user tries to start multiple searches by using the **Get-ComplianceSearch | Start-ComplianceSearch** command in Security & Compliance Center PowerShell.	10
The maximum number of items per user mailbox that are displayed on the preview page when previewing Content Search results.	100
The maximum number of items found in all user mailboxes that are displayed on the preview page when previewing search results. The newest items are displayed.	1,000
The maximum number of user mailboxes that can be previewed for search results. If there are more than 1000 mailboxes that contain content that matches the search query, only the top 1000 mailboxes with the most search results will be available for preview.	1,000
The maximum number of items found in SharePoint and OneDrive for Business sites that are displayed on the preview page when previewing search results. The newest items are displayed.	200
The maximum number of sites (in SharePoint and OneDrive for Business) that can be previewed for search results. If there are more than 200 total sites that contain content that matches the search query, only the top 200 sites with the most search results will be available for preview.	200

DESCRIPTION OF LIMIT	LIMIT
The maximum number of items per public folder mailbox that are displayed on the preview page when previewing content search results.	100
The maximum number of items found in all public folder mailboxes that are displayed on the preview page when previewing content search results.	200
The maximum number of public mailboxes that can be previewed for search results. If there are more than 500 public folder mailboxes that contain content that matches the search query, only the top 500 public folder mailboxes with the most search results will be available for preview.	500
The maximum number of characters for the search query (including operators and conditions) for a search. **Note:** This limit takes effect after the query is expanded, which means the query will get expanded against each of the keywords. For example, if a search query has 15 keywords and additional parameters and conditions, the query gets expanded 15 times, each with the other parameters and conditions in the query. So even though the number of characters in search query may be below the limit, it's the expanded query that may contribute to exceeding this limit.	**Mailboxes:** 10,000 **Sites:** 4,000 when searching all sites or 2,000 when searching up to 20 sites [2]
Maximum number of variants returned when using a prefix wildcard to search for an exact phrase in a search query or when using a prefix wildcard and the **NEAR** Boolean operator.	10,000 [3]
The minimum number of alpha characters for prefix wildcards; for example, `time*`, `one*`, or `set*`.	3
The maximum number of mailboxes in a search that you can delete items in by doing a "search and purge" action (by using the **New-ComplianceSearchAction -Purge** command). If the search that you're doing a purge action for has more source mailboxes than this limit, the purge action will fail. For more information about search and purge, see Search for and delete email messages in your organization.	50,000
The maximum number of locations in a search that you can export items from. If the search that you're exporting has more locations than this limit, the export will fail. For more information, see Export content search results.	100,000

NOTE

[1] Although you can search an unlimited number of mailboxes in a single search, you can only download the exported search results from a maximum of 100,000 mailboxes using the eDiscovery Export Tool in the Microsoft 365 compliance center. To download the search results from more than 100,000 mailboxes, you have to use Security & Compliance Center PowerShell. For more information and a sample script, see Exporting results from more than 100,000 mailboxes.

[2] When searching SharePoint and OneDrive for Business locations, the characters in the URLs of the sites being searched are counted against this limit.

[3] For non-phrase queries (a keyword value that doesn't use double quotation marks) we use a special prefix index. This tells us that a word occurs in a document, but not where it occurs in the document. To do a phrase query (a keyword value with double quotation marks), we need to compare the position within the document for the words in the phrase. This means that we can't use the prefix index for phrase queries. In this case, we internally expand the query with all possible words that the prefix expands to; for example, `"time*"` can expand to `"time OR timer OR times OR timex OR timeboxed OR …"`. 10,000 is the maximum number of variants the word can expand to, not the number of documents matching the query. There is no upper limit for non-phrase terms.

Export limits

The following table lists the limits when exporting the results of a content search. These limits also apply when you export content from a Core eDiscovery case.

DESCRIPTION OF LIMIT	LIMIT
Maximum amount of exportable data from a single search **Note:** If the search results are larger than 2 TB, consider using date ranges or other types of filters to decrease the total size of the search results.	2 TB
Maximum an organization can export in a single day **Note:** This limit is reset daily at 12:00AM UTC	2 TB
Maximum concurrent exports that can be ran at same time within your organization **Note:** Running a **Report Only** export counts against total concurrent exports for your organization. If three users are performing 3 exports each, then only one other export can be performed. Whether it is exporting a report or search results, no other exports can be performed until one has completed.	10
Maximum exports a single user can run at any one time	3
Maximum number of mailboxes for search results that can be downloaded using the eDiscovery Export Tool **Note:** To download the search results from more than 100,000 mailboxes, you have to use Security & Compliance Center PowerShell. For instructions, see Exporting results from more than 100,000 mailboxes.	100,000

DESCRIPTION OF LIMIT	LIMIT
Maximum size of PST file that can be exported **Note:** If the search results from a user's mailbox are larger than 10 GB, the search results for the mailbox will be exported in two (or more) separate PST files. If you choose to export all search results in a single PST file, the PST file will be spilt into additional PST files if the total size of the search results is larger than 10 GB. If you want to change this default size, you can edit the Windows Registry on the computer that you use to export the search results. See Change the size of PST files when exporting eDiscovery search results. The search results from a specific mailbox won't be divided among multiple PST files unless the content from a single mailbox is more than 10 GB. If you chose to export the search results in one PST file for that contains all messages in a single folder and the search results are larger than 10 GB, the items are still organized in chronological order, so they will be spilt into additional PST files based on the sent date.	10 GB
Rate at which search results from mailboxes and sites are uploaded to a Microsoft-provided Azure Storage location.	Maximum of 2 GB per hour

Indexing limits for email messages

The following table describes the indexing limits that might result in an email message being returned as an unindexed item or a partially indexed item in the results of a content search.

INDEXING LIMIT	MAXIMUM VALUE	DESCRIPTION
Maximum attachment size	150 MB	The maximum size of an email attachment that will parse for indexing. Any attachment that's larger than this limit won't be parsed for indexing, and the message with the attachment will be marked as partially indexed. **Note:** Parsing is the process where the indexing service extracts text from the attachment, removes unnecessary characters like punctuation and spaces, and then divides the text into words (in a process called tokenization), that are then stored in the index.
Maximum number of attachments	250	The maximum number of files attached to an email message that will be parsed for indexing. If a message has more than 250 attachments, the first 250 attachments are parsed and indexed, and the message is marked as partially indexed because it had additional attachments that weren't parsed.

INDEXING LIMIT	MAXIMUM VALUE	DESCRIPTION
Maximum attachment depth	30	The maximum number of nested attachments that are parsed. For example, if an email message has another message attached to it and the attached message has an attached Word document, the Word document and the attached message will be indexed. This behavior will continue for up to 30 nested attachments.
Maximum number of attached images	0	An image that's attached to an email message is skipped by the parser and isn't indexed.
Maximum time spent parsing an item	30 seconds	A maximum of 30 seconds is spent parsing an item for indexing. If the parsing time exceeds 30 seconds, the item is marked as partially indexed.
Maximum parser output	2 million characters	The maximum amount of text output from the parser that's indexed. For example, if the parser extracted 8 million characters from a document, only the first 2 million characters are indexed.
Maximum annotation tokens	2 million	When an email message is indexed, each word is annotated with different processing instructions that specify how that word should be indexed. Each set of processing instructions is called an annotation token. To maintain the quality of service in Office 365, there is a limit of 2 million annotation tokens for an email message.
Maximum body size in index	67 million characters	The total number of characters in the body of an email message and all its attachments. When an email message is indexed, all text in the body of the message and in all attachments is concatenated into a single string. The maximum size of this string that is indexed is 67 million characters.

INDEXING LIMIT	MAXIMUM VALUE	DESCRIPTION
Maximum unique tokens in body	1 million	As previously explained, tokens are the result of extracting text from content, removing punctuation and spaces, and then dividing it into words (called tokens) that are stored in the index. For example, the phrase `"cat, mouse, bird, dog, dog"` contains 5 tokens. But only 4 of these are unique tokens. There is a limit of 1 million unique tokens per email message, which helps prevent the index from getting too large with random tokens.

More information

There are additional limits related to different aspects of searching for content, such as content indexing. For more information about these limits, see the following topics:

- Partially indexed items in Content Search
- Investigating partially indexed items in eDiscovery
- Search limits for SharePoint Online

For information about content searches, see:

- Content search in Microsoft 365
- Search for content in a Core eDiscovery case
- Keyword queries and search conditions for content search

For case limits related to Core eDiscovery and Advanced eDiscovery, see:

- Limits in Core eDiscovery
- Limits in Advanced eDiscovery

Partially indexed items in eDiscovery

2/18/2021 • 12 minutes to read • Edit Online

An eDiscovery search that you run from the Microsoft 365 compliance center automatically includes partially indexed items in the estimated search results when you run a search. Partially indexed items are Exchange mailbox items and documents on SharePoint and OneDrive for Business sites that for some reason weren't completely indexed for search. In Exchange, a partially indexed item typically contains a file (of a file type that can't be indexed) that is attached to an email message. Here are some other reasons why items can't be indexed for search and are returned as partially indexed items when you run an eDiscovery search:

- The file type is unrecognized or unsupported for indexing.
- Messages have an attached file without a valid handler, such as image files; this is the most common cause of partially indexed email items.
- The file type is supported for indexing but an indexing error occurred for a specific file.
- Too many files attached to an email message.
- A file attached to an email message is too large.
- A file is encrypted with non-Microsoft technologies.
- A file is password-protected.

> **NOTE**
> Most organizations have less than 1% of content by volume and less than 12% by size that is partially indexed. The reason for the difference between volume and size is that larger files have a higher probability of containing content that can't be completely indexed.

For legal investigations, your organization may be required to review partially indexed items. You can also specify whether to include partially indexed items when you export search results to a local computer or when you prepare the results for analysis with Advanced eDiscovery. For more information, see Investigating partially indexed items in eDiscovery.

File types not indexed for search

Certain types of files, such as Bitmap or MP3 files, don't contain content that can be indexed. As a result, the search indexing servers in Exchange and SharePoint don't perform full-text indexing on these types of files. These types of files are considered to be unsupported file types. There are also file types for which full-text indexing has been disabled, either by default or by an administrator. Unsupported and disabled file types are labeled as unindexed items in Content Searches. As previously stated, partially indexed items can be included in the set of search results when you run a search, export the search results to a local computer, or prepare search results for Advanced eDiscovery.

For a list of supported and disabled file formats, see the following topics:

- **Exchange** - File formats indexed by Exchange Search
- **Exchange** - Get-SearchDocumentFormat
- **SharePoint** - Default crawled file name extensions and parsed file types in SharePoint

Messages and documents with partially indexed file types can be returned in search results

Not every email message with a partially indexed file attachment or every partially indexed SharePoint document is automatically returned as a partially indexed item. That's because other message or document properties, such as the **Subject** property in email messages and the **Title** or **Author** properties for documents are indexed and available to be searched. For example, a keyword search for "financial" will return items with a partially indexed file attachment if that keyword appears in the subject of an email message or in the file name or title of a document. However, if the keyword appears only in the body of the file, the message or document would be returned as a partially indexed item.

Similarly, messages with partially indexed file attachments and documents of a partially indexed file type are included in search results when other message or document properties, which are indexed and searchable, meet the search criteria. Message properties that are indexed for search include sent and received dates, sender and recipient, the file name of an attachment, and text in the message body. Document properties indexed for search include created and modified dates. So even though a message attachment may be a partially indexed item, the message will be included in the regular search results if the value of other message or document properties matches the search criteria.

For a list of email and document properties that you can search for by using the Search feature in the Security & Compliance Center, see Keyword queries and search conditions for eDiscovery.

Partially indexed items included in the search results

Your organization might be required to identify and perform additional analysis on partially indexed items to determine what they are, what they contain, and whether they're relevant to a specific investigation. As previously explained, the partially indexed items in the content locations that are searched are automatically included with the estimated search results. You have the option to include these partially indexed items when you export search results or prepare the search results for Advanced eDiscovery.

Keep the following in mind about partially indexed items:

- When you run an eDiscovery search, the total number and size of partially indexed Exchange items (returned by the search query) are displayed in search statistics in the details pane, and labeled as **Indexed items**. Statistics about partially indexed items displayed in the details pane don't include partially indexed items in SharePoint or OneDrive.
- If the search that you're exporting results from was a search of specific content locations or all content locations in your organization, only the unindexed items from content locations that contain items that match the search criteria will be exported. In other words, if no search results are found in a mailbox or site, then any unindexed items in that mailbox or site won't be exported. The reason for this is that exporting partially indexed items from lots of locations in the organization might increase the likelihood of export errors and increase the time it takes to export and download the search results.

  To export partially indexed items from all content locations for a search, configure the search to return all items (by removing any keywords from the search query) and then export only partially indexed items when you export the search results (by clicking **Only items that have an unrecognized format, are encrypted, or weren't indexed for other reasons** under **Output options**).
- If you choose to include all mailbox items in the search results, or if a search query doesn't specify any keywords or only specifies a date range, partially indexed items might not be copied to the PST file that contains the partially indexed items. This is because all items, including any partially indexed items, will be automatically included in the regular search results.
- Partially indexed items aren't available to be previewed. You have to export the search results to view

partially indexed items returned by the search.

Additionally, when you export search results and include partially indexed items in the export, partially indexed items from SharePoint items are exported to a folder named **Uncrawlable**. When you export partially indexed Exchange items, they are exported differently depending on whether or not the partially indexed items matched the search query and the configuration of the export settings.

The following table shows the export behavior of indexed and partially indexed items and whether or not each is included for the different export configuration settings.

EXPORT CONFIGURATION	INDEXED ITEMS THAT MATCH SEARCH QUERY	PARTIALLY INDEXED ITEMS THAT MATCH SEARCH QUERY	PARTIALLY INDEXED ITEMS THAT DON'T MATCH SEARCH QUERY
Export only indexed items	Exported	Exported (included with the indexed items that are exported)	Not exported
Export only partially indexed items	Not exported	Exported (as partially indexed items)	Exported (as partially indexed items)
Export indexed and partially indexed items	Exported	Exported (included with the indexed items that are exported)	Exported (as partially indexed items)

Partially indexed items excluded from the search results

If an item is partially indexed but it doesn't meet the search query criteria, it won't be included as a partially indexed item in the search results. In other words, the item is excluded from the search results. For example, let's say you run a search and don't include any keywords or properties because you want to include all content. But you include a date range condition for the query. If a partially indexed item falls outside of that date range, it won't be included as a partially indexed item. Date ranges are an effective way to exclude partially indexed items from your search results.

Similarly, if you choose to include partially indexed items when you export the results of a search, partially indexed items that were excluded from the search results won't be exported.

One exception to this rule is when you create a query-based hold that's associated with an eDiscovery case. If you create a query-based eDiscovery hold, all partially indexed items are placed on hold. This includes partially indexed items that don't match the search query criteria and partially indexed items that might fall outside of a date range condition. For more information about creating query-based eDiscovery holds, see Create an eDiscovery hold.

Indexing limits for messages

The following table describes the indexing limits that might result in an email message being returned as a partially indexed item in an eDiscovery search in Microsoft 365.

For a list of indexing limits for SharePoint documents, see Search limits for SharePoint Online.

INDEXING LIMIT	MAXIMUM VALUE	DESCRIPTION

INDEXING LIMIT	MAXIMUM VALUE	DESCRIPTION
Maximum attachment size (excluding Excel files)	150 MB	The maximum size of an email attachment that will parse for indexing. Any attachment that's larger than this limit won't be parsed for indexing, and the message with the attachment will be marked as partially indexed. **Note:** Parsing is the process where the indexing service extracts text from the attachment, removes unnecessary characters like punctuation and spaces, and then divides the text into words (in a process called tokenization), that are then stored in the index.
Maximum size of Excel files	4 MB	The maximum size of an Excel file located on a site or attached to an email message that will be parsed for indexing. Any Excel file that's larger than this limit won't be parsed, and the file or the email the message with the file attachment will be marked as unindexed.
Maximum number of attachments	250	The maximum number of files attached to an email message that will be parsed for indexing. If a message has more than 250 attachments, the first 250 attachments are parsed and indexed, and the message is marked as partially indexed because it had additional attachments that weren't parsed.
Maximum attachment depth	30	The maximum number of nested attachments that are parsed. For example, if an email message has another message attached to it and the attached message has an attached Word document, the Word document and the attached message will be indexed. This behavior will continue for up to 30 nested attachments.
Maximum number of attached images	0	An image that's attached to an email message is skipped by the parser and isn't indexed.
Maximum time spent parsing an item	30 seconds	A maximum of 30 seconds is spent parsing an item for indexing. If the parsing time exceeds 30 seconds, the item is marked as partially indexed.
Maximum parser output	2 million characters	The maximum amount of text output from the parser that's indexed. For example, if the parser extracted 8 million characters from a document, only the first 2 million characters are indexed.

INDEXING LIMIT	MAXIMUM VALUE	DESCRIPTION
Maximum annotation tokens	2 million	When an email message is indexed, each word is annotated with different processing instructions that specify how that word should be indexed. Each set of processing instructions is called an annotation token. To maintain the quality of service in Office 365, there is a limit of 2 million annotation tokens for an email message.
Maximum body size in index	67 million characters	The total number of characters in the body of an email message and all its attachments. When an email message is indexed, all text in the body of the message and in all attachments is concatenated into a single string. The maximum size of this string that is indexed is 67 million characters.
Maximum unique tokens in body	1 million	As previously explained, tokens are the result of extracting text from content, removing punctuation and spaces, and then dividing it into words (called tokens) that are stored in the index. For example, the phrase `"cat, mouse, bird, dog, dog"` contains 5 tokens. But only 4 of these are unique tokens. There is a limit of 1 million unique tokens per email message, which helps prevent the index from getting too large with random tokens.

More information about partially indexed items

- As previously stated, because message and document properties and their metadata are indexed, a keyword search might return results if that keyword appears in the indexed metadata. However, that same keyword search might not return the same item if the keyword only appears in the content of an item with an unsupported file type. In this case, the item would be returned as a partially indexed item.
- If a partially indexed item is included in the search results because it met the search query criteria (and wasn't excluded), then it won't be included as a partially indexed item in the estimated search statistics. Also, it won't be included with partially indexed items when you export search results.
- Although a file type is supported for indexing and is indexed, there can be indexing or search errors that will cause a file to be returned as a partially indexed item. For example, searching a very large Excel file might be partially successful (because the first 4 MB are indexed), but then fails because the file size limit is exceeded. In this case, it's possible that the same file is returned with the search results and as a partially indexed item.
- Files that are encrypted with Microsoft encryption technologies and are attached to an email message that matches the criteria of a search can be previewed and will be decrypted when exported. At this time, files that are encrypted with Microsoft encryption technologies (and stored in SharePoint or OneDrive for Business) are partially indexed.
- Email messages encrypted with S/MIME are partially indexed. This includes encrypted messages with or

without file attachments.

- Email messages protected using Azure Rights Management are indexed and will be included in the search results if they match the search query. Rights-protected email messages are decrypted and can be previewed and exported. This functionality requires that you are assigned the RMS Decrypt role, which is assigned by default to the eDiscover Manager role group.

See also

Investigating partially indexed items in eDiscovery

Investigating partially indexed items in eDiscovery

2/18/2021 • 8 minutes to read • Edit Online

An eDiscovery search that you run from the Microsoft 365 compliance center automatically includes partially indexed items in the estimated search results when you run a search. Partially indexed items are Exchange mailbox items and documents on SharePoint and OneDrive for Business sites that for some reason weren't completely indexed for search. Most email messages and site documents are successfully indexed because they fall within the Indexing limits for email messages. However, some items may exceed these indexing limits, and will be partially indexed. Here are other reasons why items can't be indexed for search and are returned as partially indexed items when you run an eDiscovery search:

- Email messages have an attached file without a valid handler, such as image files; this is the most common cause of partially indexed email items.
- Too many files attached to an email message.
- A file attached to an email message is too large.
- The file type is supported for indexing but an indexing error occurred for a specific file.

Although it varies, most organizations customers have less than 1% of content by volume and less than 12% of content by size that is partially indexed. The reason for the difference between the volume versus size is that larger files have a higher probability of containing content that can't be completely indexed.

Why does the partially indexed item count change for a search?

After you run an eDiscovery search, the total number and size of partially indexed items in the locations that were searched are listed in the search result statistics that are displayed in the detailed statistics for the search. Note these are called *unindexed items* in the search statistics. Here are a few things that will affect the number of partially indexed items that are returned in the search results:

- If an item is partially indexed and matches the search query, it's included in both the count (and size) of search result items and partially indexed items. However, when the results of that same search are exported, the item is included only with set of search results; it's not included as a partially indexed item.
- If you specify a date range for a search query (by including it in the keyword query or by using a condition), any partially indexed item that doesn't match the date range isn't included in the count of partially indexed items. Partially indexed items that fall within date range are included in the count of indexed items.

> **NOTE**
> Partially indexed items located in SharePoint and OneDrive sites *are not* included in the estimate of partially indexed items that's displayed in the detailed statistics for the search. However, partially indexed items can be exported when you export the results of an eDiscovery search. For example, if you only search sites, the estimated number partially indexed items will be zero.

Calculating the ratio of partially indexed items in your organization

To understand your organization's exposure to partially indexed items, you can run a search for all content in all mailboxes (by using a blank keyword query). In the following example below, there are 56,208 (4,830 MB) fully indexed items and 470 (316 MB) partially indexed items.

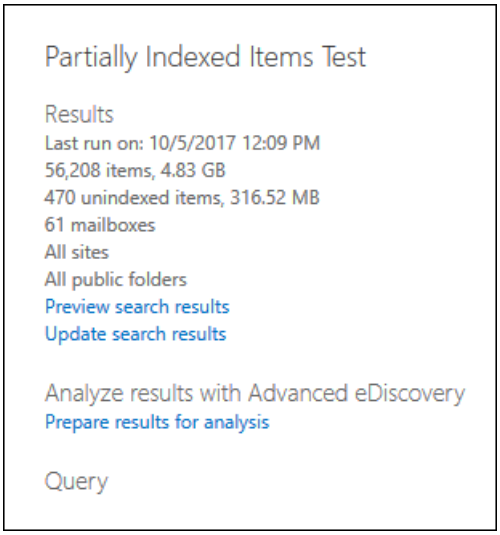

You can determine the percentage of partially indexed items by using the following calculations.

To calculate the ratio of partially indexed items in your organization:

```
(Total number of partially indexed items/Total number of items) x 100
```

```
(470/56,208) x 100 = 0.84%
```

By using the search results from the previous example, .84% of all mailboxes items are partially indexed.

To calculate the percentage of the size of partially indexed items in your organization:

```
(Size of all partially indexed items/Size of all items) x 100
```

```
(316 MB/4830 MB) x 100 = 6.54%
```

So in the previous example, 6.54% of the total size of mailbox items are from partially indexed items. As previously stated, most organizations customers have less than 1% of content by volume and less than 12% of content by size that is partially indexed.

Working with partially indexed items

In cases when you need to examine partially items to validate that they don't contain relevant information, you can export a content search report that contains information about partially indexed items. When you export a content search report, be sure to choose one of the export options that includes partially indexed items.

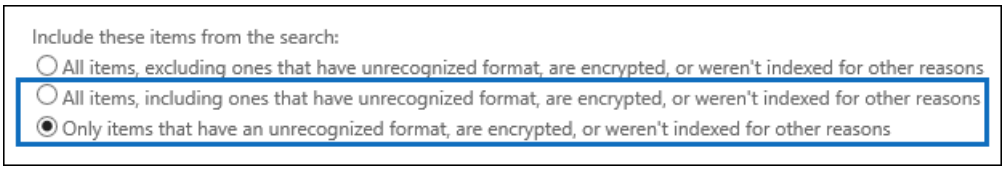

When you export eDiscovery search results or a search report using one of these options, the export includes a report named Unindexed Items.csv. This report includes most of the same information as the ResultsLog.csv file; however, the Unindexed Items.csv file also includes two fields related to partially indexed items: **Error Tags** and **Error Properties**. These fields contain information about the indexing error for each partially indexed item. Using the information in these two fields can help you determine whether or not the indexing error for a particular impacts your investigation. If it does, you can perform a targeted search and retrieve and export specific email messages and SharePoint or OneDrive documents so that you can examine them to determine if they're relevant to your investigation. For step-by-step instructions, see Prepare a CSV file for a targeted search in Office 365.

NOTE

The Unindexed Items.csv file also contains fields named **Error Type** and **Error Message**. These are legacy fields that contain information that is similar to the information in the **Error Tags** and **Error Properties** fields, but with less detailed information. You can safely ignore these legacy fields.

Errors related to partially indexed items

Error tags are made up of two pieces of information, the error and the file type. For example, in this error/file-type pair:

```
parseroutputsize_xls
```

`parseroutputsize` is the error and `xls` is the file type of the file the error occurred on. In cases where the file type wasn't recognized or the file type didn't apply to the error, you will see the value `noformat` in place of the file type.

The following is a list of indexing errors and a description of the possible cause of the error.

ERROR TAG	DESCRIPTION
`attachmentcount`	An email message had too many attachments, and some of these attachments weren't processed.
`attachmentdepth`	The content retriever and document parser found too many levels of attachments nested inside other attachments. Some of these attachments were not processed.
`attachmentrms`	An attachment failed decoding because it was RMS-protected.
`attachmentsize`	A file attached to an email message was too large and couldn't be processed.
`indexingtruncated`	When writing the processed email message to the index, one of the indexable properties was too large and was truncated. The truncated properties are listed in Error Properties field.
`invalidunicode`	An email message contained text that couldn't be processed as valid Unicode. Indexing for this item may be incomplete.
`parserencrypted`	The content of attachment or email message is encrypted, and Microsoft 365 couldn't decode the content.
`parsererror`	An unknown error occurred during parsing. This typically results from a software bug or a service crash.
`parserinputsize`	An attachment was too large for the parser to handle, and the parsing of that attachment didn't happen or wasn't completed.

ERROR TAG	DESCRIPTION
`parsermalformed`	An attachment was malformed and couldn't be handled by the parser. This result can be due to old file formats, files created by incompatible software, or viruses pretending to be something other than claimed.
`parseroutputsize`	The output from the parsing of an attachment was too large and had to be truncated.
`parserunknowntype`	An attachment had a file type that Microsoft 365 couldn't detect.
`parserunsupportedtype`	An attachment had a file type that Office 365 could detect, but parsing that file type isn't supported.
`propertytoobig`	The value of an email property in Exchange Store was too large to be retrieved and the message couldn't be processed. This typically only happens to the body property of an email message.
`retrieverrms`	The content retriever failed to decode an RMS-protected message.
`wordbreakertruncated`	Too many words were identified in the document during indexing. Processing of the property stopped when reaching the limit, and the property is truncated.

Error fields describe which fields are affected by the processing error listed in the Error Tags field. If you're searching a property such as `subject` or `participants`, errors in the body of the message won't impact the results of your search. This can be useful when determining exactly which partially indexed items you might need to further investigate.

Using a PowerShell script to determine your organization's exposure to partially indexed email items

The following steps show you how to run a PowerShell script that searches for all items in all Exchange mailboxes, and then generates a report about your organization's ratio of partially indexed email items (by count and by size) and displays the number of items (and their file type) for each indexing error that occurs. Use the error tag descriptions in the previous section to identify the indexing error.

1. Save the following text to a Windows PowerShell script file by using a filename suffix of .ps1; for example, `PartiallyIndexedItems.ps1`.

```
  write-host "************************************************"
  write-host "     Security & Compliance Center      " -foregroundColor yellow -backgroundcolor
darkgreen
  write-host "   eDiscovery Partially Indexed Item Statistics   " -foregroundColor yellow -
backgroundcolor darkgreen
  write-host "************************************************"
  " "
  # Create a search with Error Tags Refinders enabled
  Remove-ComplianceSearch "RefinerTest" -Confirm:$false -ErrorAction 'SilentlyContinue'
  New-ComplianceSearch -Name "RefinerTest" -ContentMatchQuery "size>0" -RefinerNames ErrorTags -
ExchangeLocation ALL
  Start-ComplianceSearch "RefinerTest"
  # Loop while search is in progress
  do{
      Write-host "Waiting for search to complete..."
      Start-Sleep -s 5
      $complianceSearch = Get-ComplianceSearch "RefinerTest"
  }while ($complianceSearch.Status -ne 'Completed')
  $refiners = $complianceSearch.Refiners | ConvertFrom-Json
  $errorTagProperties = $refiners.Entries | Get-Member -MemberType NoteProperty
  $partiallyIndexedRatio = $complianceSearch.UnindexedItems / $complianceSearch.Items
  $partiallyIndexedSizeRatio = $complianceSearch.UnindexedSize / $complianceSearch.Size
  " "
  "===== Partially indexed items ====="
  "         Total          Ratio"
  "Count    {0:N0}{1:P2}" -f $complianceSearch.Items.ToString("N0").PadRight(15, " "),
$partiallyIndexedRatio
  "Size(GB) {0:N2}{1:P2}" -f ($complianceSearch.Size / 1GB).ToString("N2").PadRight(15, " "),
$partiallyIndexedSizeRatio
  " "
  Write-Host ===== Reasons for partially indexed items =====
  foreach($errorTagProperty in $errorTagProperties)
  {
      $name = $refiners.Entries.($errorTagProperty.Name).Name
      $count = $refiners.Entries.($errorTagProperty.Name).TotalCount
      $frag = $name.Split("{_}")
      $errorTag = $frag[0]
      $fileType = $frag[1]
      if ($errorTag -ne $lastErrorTag)
      {
          $errorTag
      }
      "    " + $fileType + " => " + $count
      $lastErrorTag = $errorTag
  }
```

2. Connect to Security & Compliance Center PowerShell.

3. In Security & Compliance Center PowerShell, go to the folder where you saved the script in step 1, and then run the script; for example:

```
.\PartiallyIndexedItems.ps1
```

Here's an example fo the output returned by the script.

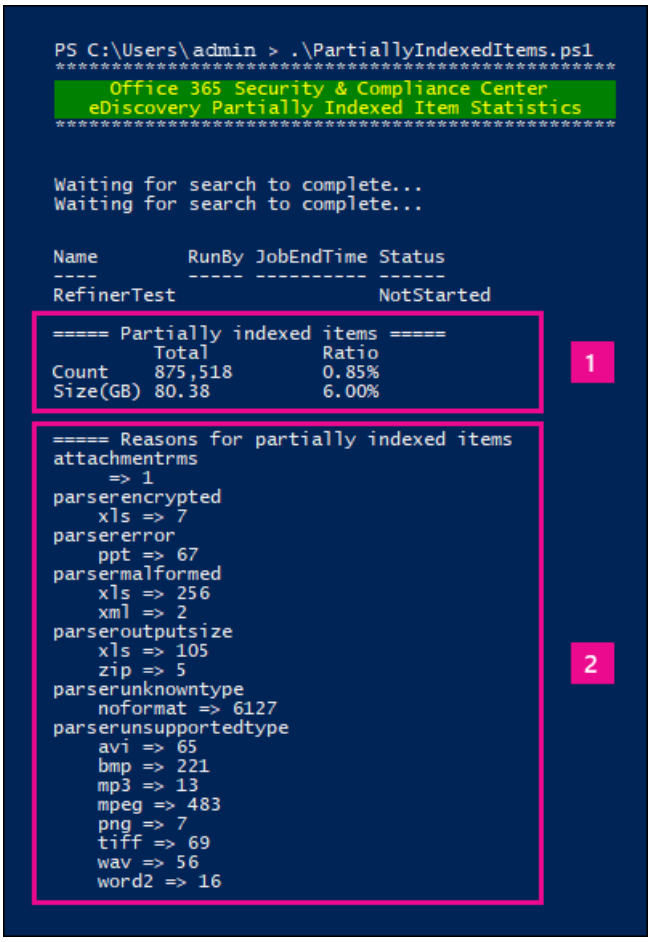

NOTE

Note the following:

- The total number and size of email items, and your organization's ratio of partially indexed email items (by count and by size).
- A list error tags and the corresponding file types for which the error occurred.

See also

Partially indexed items in eDiscovery

De-duplication in eDiscovery search results

11/2/2020 • 5 minutes to read • Edit Online

This article describes how de-duplication of eDiscovery search results works and explains the limitations of the de-duplication algorithm.

When using eDiscovery tools to export the results of an eDiscovery search, you have the option to de-duplicate the results that are exported. What does this mean? When you enable de-duplication (by default, de-duplication isn't enabled), only one copy of an email message is exported even though multiple instances of the same message might have been found in the mailboxes that were searched. De-duplication helps you save time by reducing the number of items that you have to review and analyze after the search results are exported. But it's important to understand how de-duplication works and be aware that there are limitations to the algorithm that might cause a unique item to be marked as a duplicate during the export process.

How duplicate messages are identified

eDiscovery tools use a combination of the following email properties to determine whether a message is a duplicate:

- **InternetMessageId** - This property specifies the Internet message identifier of an email message, which is a globally unique identifier that refers to a specific version of a specific message. This ID is generated by the sender's email client program or host email system that sends the message. If a person sends a message to more than one recipient, the Internet message ID will be the same for each instance of the message. Subsequent revisions to the original message will receive a different message identifier.
- **ConversationTopic** - This property specifies the subject of the conversation thread of a message. The value of the **ConversationTopic** property is the string that describes the overall topic of the conversation. A conservation consists of an initial message and all messages sent in reply to the initial message. Messages within the same conversation have the same value for the **ConversationTopic** property. The value of this property is typically the Subject line from the initial message that spawned the conversation.
- **BodyTagInfo** - This is an internal Exchange store property. The value of this property is calculated by checking various attributes in the body of the message. This property is used to identify differences in the body of messages.

During the eDiscovery export process, these three properties are compared for every message that matches the search criteria. If these properties are identical for two (or more) messages, those messages are determined to be duplicates and the result is that only one copy of the message will be exported if de-duplication is enabled. The message that is exported is known as the "source item". Information about duplicate messages is included in the **Results.csv** and **Manifest.xml** reports that are included with the exported search results. In the **Results.csv** file, a duplicate message is identified by having a value in the **Duplicate to Item** column. The value in this column matches the value in the **Item Identity** column for the message that was exported.

The following graphics show how duplicate messages are displayed in the **Results.csv** and **Manifest.xml** reports that are exported with the search results. These reports don't include the email properties previously described, which are used in the de-duplication algorithm. Instead, the reports include the **Item Identity** property that is assigned to items by the Exchange store.

Results.csv report (viewed in Excel)

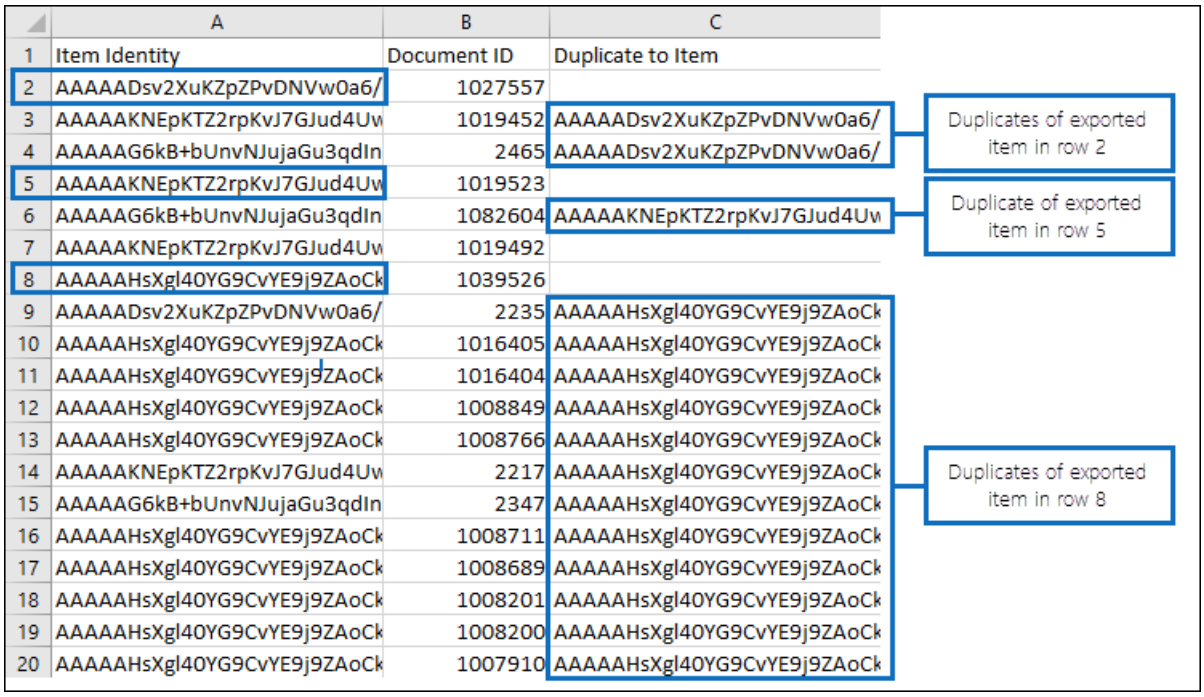

	A	B	C
1	Item Identity	Document ID	Duplicate to Item
2	AAAAADsv2XuKZpZPvDNVw0a6/	1027557	
3	AAAAAKNEpKTZ2rpKvJ7GJud4Uw	1019452	AAAAADsv2XuKZpZPvDNVw0a6/
4	AAAAAG6kB+bUnvNJujaGu3qdIn	2465	AAAAADsv2XuKZpZPvDNVw0a6/
5	AAAAAKNEpKTZ2rpKvJ7GJud4Uw	1019523	
6	AAAAAG6kB+bUnvNJujaGu3qdIn	1082604	AAAAAKNEpKTZ2rpKvJ7GJud4Uw
7	AAAAAKNEpKTZ2rpKvJ7GJud4Uw	1019492	
8	AAAAAHsXgl40YG9CvYE9j9ZAoCk	1039526	
9	AAAAADsv2XuKZpZPvDNVw0a6/	2235	AAAAAHsXgl40YG9CvYE9j9ZAoCk
10	AAAAAHsXgl40YG9CvYE9j9ZAoCk	1016405	AAAAAHsXgl40YG9CvYE9j9ZAoCk
11	AAAAAHsXgl40YG9CvYE9j9ZAoCk	1016404	AAAAAHsXgl40YG9CvYE9j9ZAoCk
12	AAAAAHsXgl40YG9CvYE9j9ZAoCk	1008849	AAAAAHsXgl40YG9CvYE9j9ZAoCk
13	AAAAAHsXgl40YG9CvYE9j9ZAoCk	1008766	AAAAAHsXgl40YG9CvYE9j9ZAoCk
14	AAAAAKNEpKTZ2rpKvJ7GJud4Uw	2217	AAAAAHsXgl40YG9CvYE9j9ZAoCk
15	AAAAAG6kB+bUnvNJujaGu3qdIn	2347	AAAAAHsXgl40YG9CvYE9j9ZAoCk
16	AAAAAHsXgl40YG9CvYE9j9ZAoCk	1008711	AAAAAHsXgl40YG9CvYE9j9ZAoCk
17	AAAAAHsXgl40YG9CvYE9j9ZAoCk	1008689	AAAAAHsXgl40YG9CvYE9j9ZAoCk
18	AAAAAHsXgl40YG9CvYE9j9ZAoCk	1008201	AAAAAHsXgl40YG9CvYE9j9ZAoCk
19	AAAAAHsXgl40YG9CvYE9j9ZAoCk	1008200	AAAAAHsXgl40YG9CvYE9j9ZAoCk
20	AAAAAHsXgl40YG9CvYE9j9ZAoCk	1007910	AAAAAHsXgl40YG9CvYE9j9ZAoCk

Manifest.xml report (viewed in Excel)

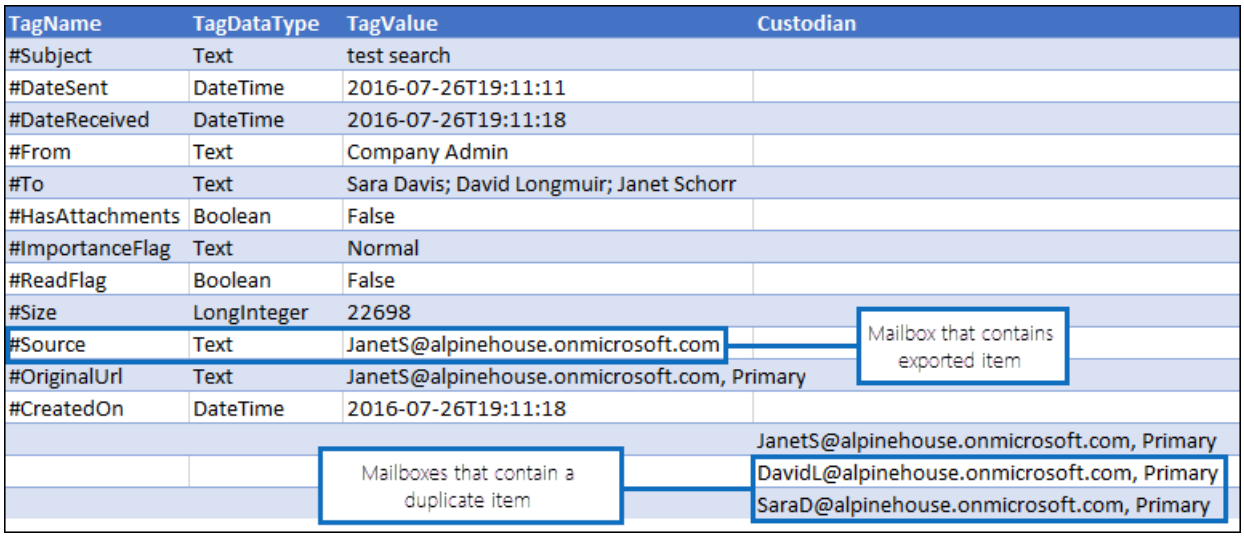

TagName	TagDataType	TagValue	Custodian
#Subject	Text	test search	
#DateSent	DateTime	2016-07-26T19:11:11	
#DateReceived	DateTime	2016-07-26T19:11:18	
#From	Text	Company Admin	
#To	Text	Sara Davis; David Longmuir; Janet Schorr	
#HasAttachments	Boolean	False	
#ImportanceFlag	Text	Normal	
#ReadFlag	Boolean	False	
#Size	LongInteger	22698	
#Source	Text	JanetS@alpinehouse.onmicrosoft.com	
#OriginalUrl	Text	JanetS@alpinehouse.onmicrosoft.com, Primary	
#CreatedOn	DateTime	2016-07-26T19:11:18	
			JanetS@alpinehouse.onmicrosoft.com, Primary
			DavidL@alpinehouse.onmicrosoft.com, Primary
			SaraD@alpinehouse.onmicrosoft.com, Primary

Additionally, other properties from duplicate messages are included in the export reports. This includes the mailbox the duplicate message is located in, whether the message was sent to a distribution group, and whether the message was Cc'd or Bcc'd to another user.

Limitations of the de-duplication algorithm

There are some known limitations of the de-duplication algorithm that might cause unique items to get marked as duplicates. It's important to understand these limitations so you can decide whether or not to use the optional de-duplication feature.

There's one situation where the de-duplication feature might mistakenly identify a message as a duplicate and not export it (but still cite it as a duplicate in the export reports). These are messages that a user edits but doesn't send. For example, let's say a user selects a message in Outlook, copies the contents of the message, and then pastes it in a new message. Then the user changes one of the copies by removing or adding an attachment, or changing the subject line or the body itself. If these two messages match the query of an eDiscovery search, only one of the messages will be exported if de-duplication is enabled when the search results are exported. So even though the original message or the copied message was changed, neither of the revised messages were sent and therefore the values of **InternetMessageId**, **ConversationTopic** and **BodyTagInfo** properties weren't updated. But as previously explained, both messages will be listed in the export reports

Unique messages can also be marked as duplicates when the Copy-on-Write page protection feature is enabled, as in the case of a mailbox being on Litigation Hold or In-Place Hold. The Copy-on-Write feature copies the original message (and saves it in the Versions folder of the user's Recoverable Items folder) before the revision to original item is saved. In this case, the revised copy and the original message (in the Recoverable Items folder) might be considered as duplicate messages and therefore only one of them would be exported.

> **IMPORTANT**
>
> If the limitations of the de-duplication algorithm might impact the quality of your search results, then you shouldn't enable de-duplication when you export items. If the situations described in this section are unlikely to be a factor in your search results, and you want to reduce the number of items most likely to be duplicates, then you should consider enabling de-duplication.

More information

- The information in this article is applicable when exporting search results using one of the following eDiscovery tools:
  - Content search in compliance center in Office 365
  - In-Place eDiscovery in Exchange Online
  - The eDiscovery Center in SharePoint Online
- For more information about exporting search results, see:
  - Export Content Search
  - Export a Content Search report
  - Export In-Place eDiscovery search results to a PST file
  - Export content and create reports in the eDiscovery Center

Differences between estimated and actual eDiscovery search results

2/18/2021 • 5 minutes to read • Edit Online

This topic applies to searches that you can run using one of the following Microsoft 365 eDiscovery tools:

- Content search
- Core eDiscovery

When you run an eDiscovery search, the tool you're using will return an estimate of the number of items (and their total size) that meet the search criteria. For example, when you run a search in the Microsoft 365 compliance center, the estimated search results are displayed on the flyout page for the selected search.

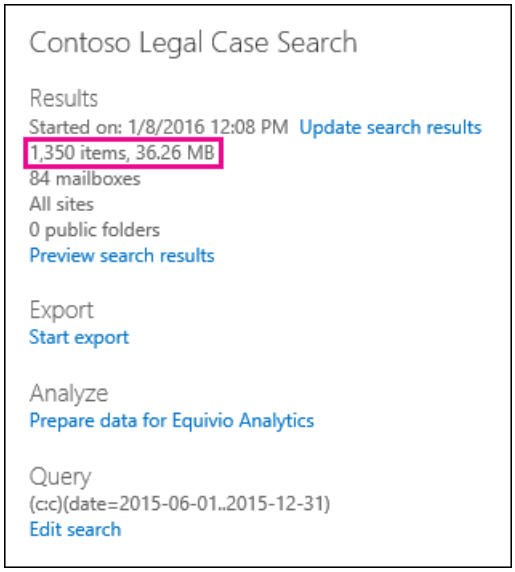

This is the same estimate of total size and number of items that is displayed in the eDiscovery Export Tool when you export results to a local computer and in the Export Summary report that's downloaded with the search results.

Estimated results in the eDiscovery Export tool

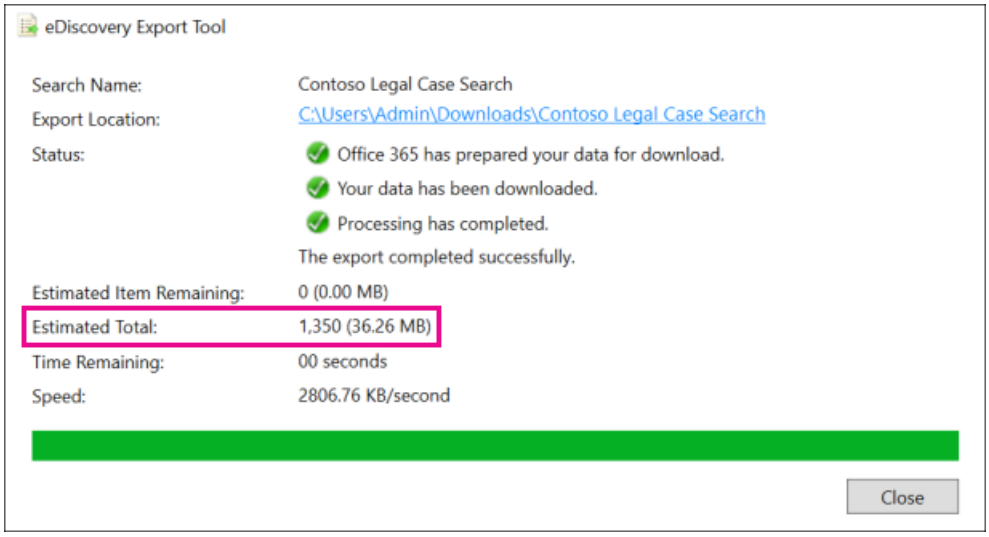

Estimated results in Export Summary report

1	Application launched from	
2	Export Name	Contoso Legal Case Search
3	Download Started	1/8/2016 14:43
4	Download Completed	1/8/2016 14:46
5	Status	The export completed successfully.
6	Locations	
7	Total Locations	22
8	Succeeded Locations	22
9	Failed Locations	0
10	Size	
11	Estimated Size	36.26 MB (38,024,714 bytes)
12	Downloaded Size	37.08 MB (38,882,587 bytes)
13	Items	
14	Estimated Items	1350
15	Not Downloaded Due to Errors	0
16	Total Downloaded Items	1218

However, as you'll notice in the previous screenshot of the Export Summary report, the size and number of actual search results that are downloaded are different than the size and number of estimated search results.

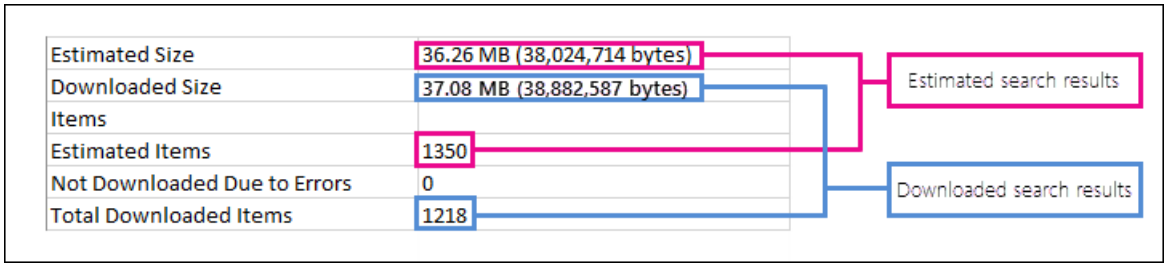

Here are some reasons for these differences:

- **The way results are estimated**. An estimate of the search results is just that, an estimate (and not an actual count) of the items that meet the search query criteria. To compile the estimate of Exchange items, a list of the message IDs that meet the search criteria is requested from the Exchange database by the eDiscovery tool you're using. But when you export the search results, the search is rerun and the actual messages are retrieved from the Exchange database. So these differences might result because of how the estimated number of items and the actual number of items are determined.

- **Changes that happen between the time when estimating and exporting search results**. When you export search results, the search is restarted to collect that most recent items in the search index that meet the search criteria. It's possible there are additional items were created, sent, or received that meet the search criteria in the time between when the estimated search results were collected and when the search results were exported. It's also possible that items that were in the search index when the search results were estimated are no longer there because they were purged from the content location before the search results are exported. One way to mitigate this issue is to specify a date range for an eDiscovery search. Another way is to place a hold on content locations so that items are preserved and can't be purged.

  Although rare, even in the case when a hold is applied, maintenance of built-in calendar items (which aren't editable by the user, but are included in many search results) may be removed from time to time. This periodic removal of calendar items will result in fewer items that are exported.

- **Unindexed items**. Items that are unindexed for search can cause differences between estimated and actual search results. You can include unindexed items when you export the search results. If you include unindexed items when exporting search results, there might be more items that are exported. This will cause a difference between the estimated and exported search results.

  When using the Content search tool, you have the option to include unindexed items in the search estimate. The number of unindexed items returned by the search is listed on the flyout page together with the other estimated search results. Any unindexed items would also be included in the total size of

the estimated search results. When you export search results, you have the option to include or not include unindexed items. How you configure these options might result in differences between estimated and the actual search results that are downloaded.

- **Exporting the results of a Content Search that includes all content locations**. If the search that you're exporting results from was a search of all content locations in your organization, then only the unindexed items from content locations that contain items that match the search criteria will be exported. In other words, if no search results are found in a mailbox or site, then any unindexed items in that mailbox or site won't be exported. However, unindexed items from all content locations (even those that don't contain items that match the search query) will be included in the estimated search results.

  Alternatively, if the search that you're exporting results from included specific content locations, then unindexed items (that aren't excluded by the search criteria) from all the content locations specified in the search will be exported. In this case, the estimated number of unindexed items and the number of unindexed items that are exported should be the same.

  The reason for not exporting unindexed items from every location in the organization is because it might increase the likelihood of export errors and increase the time it takes to export and download the search results.

- **Raw file formats versus exported file formats**. For Exchange items, the estimated size of the search results is calculated by using the raw Exchange message sizes. However, email messages are exported in a PST file or as individual messages (which are formatted as EML files). Both of these export options use a different file format than raw Exchange messages, which results in the total exported file size being different than the estimated file size.

- **Document versions**. For SharePoint documents, multiple versions of a document aren't included in the estimated search results. But you have the option to include all document versions when you export the search results, which will increase the actual number (and total size) of the exported documents.

- **De-duplication**. For Exchange items, de-duplication reduces the number of items that are exported. You have the option to de-duplicate the search results when you export them. For Exchange messages, this means that only a single instance of a message is exported, even though that message might be found in multiple mailboxes. The estimated search results include every instance of a message. So if you choose the de-duplication option when exporting search results, the actual number of items that are exported might be considerably less than the estimated number of items.

  Another thing to keep in mind if you choose the de-duplication option is that all Exchange items are exported in a single PST file and the folder structure from the source mailboxes isn't preserved. The exported PST file just contains the email items. However, a search results report contains an entry for each exported message that identifies the source mailbox where the message is located. This helps you identify all mailboxes that contain a duplicate message. If you don't enable de-duplication, a separate PST file is exported for each mailbox included in the search.

> **NOTE**
>
> If you don't select the **Include items that are encrypted or have an unrecognized format** option when you export search results or just download the reports, the index error reports are downloaded but they don't have any entries. This doesn't mean there aren't any indexing errors. It just means that unindexed items weren't included in the export.

Decryption in Microsoft 365 eDiscovery tools

2/18/2021 • 3 minutes to read • Edit Online

Encryption is an important part of your file protection and information protection strategy. Organizations of all types use encryption technology to protect sensitive content within their organization and ensure that only the right people have access to that content.

To execute common eDiscovery tasks on encrypted content, eDiscovery managers were required to decrypt email message content as it was exported from content searches, Core eDiscovery cases, and Advanced eDiscovery cases. Content encrypted with Microsoft encryption technologies wasn't available for review until after it was exported.

To make it easier to manage encrypted content in the eDiscovery workflow, Microsoft 365 eDiscovery tools now incorporate decryption of encrypted files that are attached to email messages and sent in Exchange Online. Additionally, encrypted documents stored in SharePoint Online and OneDrive for Business are decrypted in Advanced eDiscovery.

Prior to this new capability, only the content of an email message protected by rights management (and not attached files) were decrypted. Encrypted documents in SharePoint and OneDrive couldn't be decrypted during the eDiscovery workflow. Now, if a file that's encrypted with a Microsoft encryption technology is attached to an email message or located on a SharePoint or OneDrive account, those encrypted items are decrypted when the search results are prepared for preview, added to a review set in Advanced eDiscovery, and exported. This allows eDiscovery managers to view the content of encrypted email attachments and site documents when previewing search results, and review them after they have been added to a review set in Advanced eDiscovery.

Supported encryption technologies

Microsoft eDiscovery tools support items encrypted with Microsoft encryption technologies. These technologies include Office Message Encryption, Azure Rights Management, and Microsoft Information Protection (specifically sensitivity labels). For more information about Microsoft encryption technologies, see Encryption. Content encrypted by third-party encryption technologies isn't supported. For example, previewing or exporting content encrypted with non-Microsoft technologies isn't supported.

eDiscovery activities that support encrypted items

The following table identifies the supported tasks that can be performed in Microsoft 365 eDiscovery tools on encrypted files attached to email massages and encrypted documents in SharePoint and OneDrive. These supported tasks can be performed on encrypted files that match the criteria of a search. A value of `N/A` indicates the functionality isn't available in the corresponding eDiscovery tool.

EDISCOVERY TASK	CONTENT SEARCH	CORE EDISCOVERY	ADVANCED EDISCOVERY
Search for content in encrypted files in email and sites	Yes	Yes	Yes
Preview encrypted files attached to email	Yes	Yes	Yes

EDISCOVERY TASK	CONTENT SEARCH	CORE EDISCOVERY	ADVANCED EDISCOVERY
Preview encrypted documents in SharePoint and OneDrive	No	No	Yes
Review encrypted files in a review set	N/A	N/A	Yes
Export encrypted files attached to email	Yes	Yes	Yes
Export encrypted documents in SharePoint and OneDrive	No	No	Yes

Note: eDiscovery doesn't support encrypted files in SharePoint and OneDrive when a sensitivity label that applied the encryption is configured with either of the following settings:

- Users can assign permissions when they manually apply the label to a document. This is sometimes referred to as *user-defined permissions*.
- User access to the document has an expiration setting that is set to a value other than **Never**.

For more information about these settings, see the "Configure encryption settings" section in Restrict access to content by using sensitivity labels to apply encryption.

Documents encrypted with the previous settings can still be returned by an eDiscovery search. This may happen when a document property (such as the title, author, or modified date) matches the search criteria. Although these documents might be included in search results, they can't be previewed or reviewed. These documents will also remain encrypted when they're exported in Advanced eDiscovery.

Requirements for decryption in eDiscovery

You have to be assigned the RMS Decrypt role to preview, review, and export files encrypted with Microsoft encryption technologies. You also have to be assigned this role to review and query encrypted files that are added to a review set in Advanced eDiscovery.

This role is assigned by default to the eDiscovery Manager role group on the **Permissions** page in the Office 365 Security & Compliance Center. For more information about the RMS Decrypt role, see Assign eDiscovery permissions.

Use Content Search for targeted collections

2/18/2021 • 13 minutes to read • Edit Online

The Content Search feature in the Microsoft 365 compliance center doesn't provide a direct way in the UI to search specific folders in Exchange mailboxes or SharePoint and OneDrive for Business sites. However, it's possible to search specific folders (called a *targeted collection*) by specifying the folder ID property for email or path (DocumentLink) property for sites in the actual search query syntax. Using Content Search to perform a targeted collection is useful when you're confident that items responsive to a case or privileged items are located in a specific mailbox or site folder. You can use the script in this article to obtain the folder ID for mailbox folders or the path (DocumentLink) for folders on a SharePoint and OneDrive for Business site. Then you can use the folder ID or path in a search query to return items located in the folder.

> **NOTE**
>
> To return content located in a folder in a SharePoint or OneDrive for Business site, the script in this topic uses the DocumentLink managed property instead of the Path property. The DocumentLink property is more robust than the Path property because it will return all content in a folder, whereas the Path property won't return some media files.

Before you run a targeted collection

- You have to be a member of the eDiscovery Manager role group in the Security & Compliance Center to run the script in Step 1. For more information, see Assign eDiscovery permissions.

  Additionally, you have to be assigned the Mail Recipients role in your Exchange Online organization. This is required to run the **Get-MailboxFolderStatistics** cmdlet, which is included in the script. By default, the Mail Recipients role is assigned to the Organization Management and Recipient Management role groups in Exchange Online. For more information about assigning permissions in Exchange Online, see Manage role group members. You could also create a custom role group, assign the Mail Recipients role to it, and then add the members who need to run the script in Step 1. For more information, see Manage role groups.

- The script in this article supports modern authentication. You can use the script as-is if you are a Microsoft 365 or a Microsoft 365 GCC organization. If you are an Office 365 Germany organization, a Microsoft 365 GCC High organization, or a Microsoft 365 DoD organization, you will have to edit the script to successfully run it. Specifically, you have to edit the line `Connect-ExchangeOnline` and use the *ExchangeEnvironmentName* parameter (and the appropriate value for your organization type) to connect to Exchange Online PowerShell. Also, you have to edit the line `Connect-IPPSSession` and use the *ConnectionUri* and *AzureADAuthorizationEndpointUri* parameters (and the appropriate values for your organization type) to connect to Security & Compliance Center PowerShell. For more information, see the examples in Connect to Exchange Online PowerShell and Connect to Security & Compliance Center PowerShell.

- Each time you run the script, a new remote PowerShell session is created. That means you can use up all the remote PowerShell sessions available to you. To prevent this from happening, run the following command to disconnect your active remote PowerShell sessions.

  ```
  Get-PSSession | Remove-PSSession
  ```

  For more information, see Connect to Exchange Online PowerShell.

- The script includes minimal error handling. The primary purpose of the script is to quickly display a list of mailbox folder IDs or site paths that can be used in the search query syntax of a Content Search to perform a targeted collection.
- The sample script provided in this topic isn't supported under any Microsoft standard support program or service. The sample script is provided AS IS without warranty of any kind. Microsoft further disclaims all implied warranties including, without limitation, any implied warranties of merchantability or of fitness for a particular purpose. The entire risk arising out of the use or performance of the sample script and documentation remains with you. In no event shall Microsoft, its authors, or anyone else involved in the creation, production, or delivery of the scripts be liable for any damages whatsoever (including, without limitation, damages for loss of business profits, business interruption, loss of business information, or other pecuniary loss) arising out of the use of or inability to use the sample scripts or documentation, even if Microsoft has been advised of the possibility of such damages.

Step 1: Run the script to get a list of folders for a mailbox or site

The script that you run in this first step will return a list of mailbox folders or SharePoint and OneDrive for Business folders, and the corresponding folder ID or path for each folder. When you run this script, it will prompt you for the following information.

- **Email address or site URL**: Type an email address of the custodian to return a list of Exchange mailbox folders and folder IDs. Or type the URL for a SharePoint site or a OneDrive for Business site to return a list of paths for the specified site. Here are some examples:
  - **Exchange**: stacig@contoso.onmicrosoft.com
  - **SharePoint**: https://contoso.sharepoint.com/sites/marketing
  - **OneDrive for Business**: https://contoso-my.sharepoint.com/personal/stacig_contoso_onmicrosoft_com
- **Your user credentials**: The script will use your credentials to connect to Exchange Online PowerShell or Security & Compliance Center PowerShell using modern authentication. As previously explained, you have to be assigned the appropriate permissions to successfully run this script.

To display a list of mailbox folders or site documentlink (path) names:

1. Save the following text to a Windows PowerShell script file by using a filename suffix of .ps1; for example, `GetFolderSearchParameters.ps1`.

```
###########################################################################################################
####
# This PowerShell script will prompt you for:                                  #
#    * Admin credentials for a user who can run the Get-MailboxFolderStatistics cmdlet in Exchange
#
#      Online and who is an eDiscovery Manager in the Security & Compliance Center.            #
# The script will then:                                          #
#    * If an email address is supplied: list the folders for the target mailbox.            #
#    * If a SharePoint or OneDrive for Business site is supplied: list the documentlinks (folder
paths) #
#    * for the site.
#
#    * In both cases, the script supplies the correct search properties (folderid: or documentlink:)
#
#      appended to the folder ID or documentlink to use in a Content Search.                #
# Notes:                                                #
#    * For SharePoint and OneDrive for Business, the paths are searched recursively; this means the
#
#      the current folder and all sub-folders are searched.                        #
#    * For Exchange, only the specified folder will be searched; this means sub-folders in the folder
#
```

```
#
#      will not be searched.  To search sub-folders, you need to use the specify the folder ID for
#
#      each sub-folder that you want to search.                                    #
#    * For Exchange, only folders in the user's primary mailbox will be returned by the script.
#
#############################################################################################################
####
# Collect the target email address or SharePoint Url
$addressOrSite = Read-Host "Enter an email address or a URL for a SharePoint or OneDrive for Business
site"
# Authenticate with Exchange Online and the Security & Compliance Center (Exchange Online Protection
- EOP)
if ($addressOrSite.IndexOf("@") -ige 0)
{
   # List the folder Ids for the target mailbox
   $emailAddress = $addressOrSite
   # Connect to Exchange Online PowerShell
   if (!$ExoSession)
   {
    Import-Module ExchangeOnlineManagement
    Connect-ExchangeOnline
   }
   $folderQueries = @()
   $folderStatistics = Get-MailboxFolderStatistics $emailAddress
   foreach ($folderStatistic in $folderStatistics)
   {
    $folderId = $folderStatistic.FolderId;
    $folderPath = $folderStatistic.FolderPath;
    $encoding= [System.Text.Encoding]::GetEncoding("us-ascii")
    $nibbler= $encoding.GetBytes("0123456789ABCDEF");
    $folderIdBytes = [Convert]::FromBase64String($folderId);
    $indexIdBytes = New-Object byte[] 48;
    $indexIdIdx=0;
    $folderIdBytes | select -skip 23 -First 24 | %{$indexIdBytes[$indexIdIdx++]=$nibbler[$_ -shr
4];$indexIdBytes[$indexIdIdx++]=$nibbler[$_ -band 0xF]}
    $folderQuery = "folderid:$($encoding.GetString($indexIdBytes))";
    $folderStat = New-Object PSObject
    Add-Member -InputObject $folderStat -MemberType NoteProperty -Name FolderPath -Value $folderPath
    Add-Member -InputObject $folderStat -MemberType NoteProperty -Name FolderQuery -Value
$folderQuery
    $folderQueries += $folderStat
   }
   Write-Host "-----Exchange Folders-----"
   $folderQueries |ft
}
elseif ($addressOrSite.IndexOf("http") -ige 0)
{
   $searchName = "SPFoldersSearch"
   $searchActionName = "SPFoldersSearch_Preview"
   # List the folders for the SharePoint or OneDrive for Business Site
   $siteUrl = $addressOrSite
   # Connect to Security & Compliance Center PowerShell
   if (!$SccSession)
   {
    Import-Module ExchangeOnlineManagement
    Connect-IPPSSession
   }
   # Clean-up, if the script was aborted, the search we created might not have been deleted.  Try to
do so now.
   Remove-ComplianceSearch $searchName -Confirm:$false -ErrorAction 'SilentlyContinue'
   # Create a Content Search against the SharePoint Site or OneDrive for Business site and only
search for folders; wait for the search to complete
   $complianceSearch = New-ComplianceSearch -Name $searchName -ContentMatchQuery "contenttype:folder"
-SharePointLocation $siteUrl
   Start-ComplianceSearch $searchName
   do{
    Write-host "Waiting for search to complete..."
    Start-Sleep -s 5
```

```
    $complianceSearch = Get-ComplianceSearch $searchName
   }while ($complianceSearch.Status -ne 'Completed')
   if ($complianceSearch.Items -gt 0)
   {
    # Create a Compliance Search Action and wait for it to complete. The folders will be listed in
the .Results parameter
    $complianceSearchAction = New-ComplianceSearchAction -SearchName $searchName -Preview
    do
    {
     Write-host "Waiting for search action to complete..."
     Start-Sleep -s 5
     $complianceSearchAction = Get-ComplianceSearchAction $searchActionName
    }while ($complianceSearchAction.Status -ne 'Completed')
    # Get the results and print out the folders
    $results = $complianceSearchAction.Results
    $matches = Select-String "Data Link:.+[,}]" -Input $results -AllMatches
    foreach ($match in $matches.Matches)
    {
     $rawUrl = $match.Value
     $rawUrl = $rawUrl -replace "Data Link: " -replace "," -replace "}"
     Write-Host "DocumentLink:""$rawUrl"""
    }
   }
   else
   {
    Write-Host "No folders were found for $siteUrl"
   }
   Remove-ComplianceSearch $searchName -Confirm:$false -ErrorAction 'SilentlyContinue'
}
else
{
   Write-Error "Couldn't recognize $addressOrSite as an email address or a site URL"
}
```

2. On your local computer, open Windows PowerShell and go to the folder where you saved the script.
3. Run the script; for example:

```
.\GetFolderSearchParameters.ps1
```

4. Enter the information that the script prompts you for.

   The script displays a list of mailbox folders or site folders for the specified user. Leave this window open so that you can copy a folder ID or documentlink name and paste it in to a search query in Step 2.

> **TIP**
> Instead of displaying a list of folders on the computer screen, you can re-direct the output of the script to a text file. This file will be saved to the folder where the script is located. For example, to redirect the script output to a text file, run the following command in Step 3: `.\GetFolderSearchParameters.ps1 > StacigFolderIds.txt`
> Then you can copy a folder ID or documentlink from the file to use in a search query.

Script output for mailbox folders

If you're getting mailbox folder IDs, the script connects to Exchange Online PowerShell, runs the **Get-MailboxFolderStatisics** cmdlet, and then displays the list of the folders from the specified mailbox. For every folder in the mailbox, the script displays the name of the folder in the **FolderPath** column and the folder ID in the **FolderQuery** column. Additionally, the script adds the prefix of **folderId** (which is the name of the mailbox property) to the folder ID. Because the **folderid** property is a searchable property, you'll use `folderid:<folderid>` in a search query in Step 2 to search that folder. The script displays a maximum of 100 mailbox folders.

IMPORTANT

The script in this article includes encoding logic that converts the 64-character folder Id values that are returned by **Get-MailboxFolderStatistics** to the same 48-character format that is indexed for search. If you just run the **Get-MailboxFolderStatistics** cmdlet in PowerShell to obtain a folder Id (instead of running the script in this article), a search query that uses that folder Id value will fail. You have to run the script to get the correctly-formatted folder Ids that can be used in a Content Search.

Here's an example of the output returned by the script for mailbox folders.

```
FolderPath                                FolderQuery
----------                                -----------
/Top of Information Store                 folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB0000000001080000
/Archive                                  folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB0000328C0E850000
/Calendar                                 folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB000000000010D0000
/Calendar/Birthdays                       folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB000000000014E0000
/Calendar/United States holidays          folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB000000000014F0000
/Contacts                                 folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB000000000010E0000
/Contacts/Companies                       folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB000000000014D0000
/Contacts/GAL Contacts                    folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB000000000012A0000
/Contacts/Organizational Contacts         folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB000000000014C0000
/Contacts/Recipient Cache                 folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB0000000001270000
/Conversation Action Settings             folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB0000000001260000
/Conversation History                     folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB0000000001510000
/Deleted Items                            folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB000000000010A0000
/Drafts                                   folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB000000000010F0000
/ExternalContacts                         folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB0000000001200000
/Files                                    folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB000000000013D0000
/Inbox                                    folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB000000000010C0000
/Journal                                  folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB0000000001100000
/Junk Email                               folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB0000000001250000
/Notes                                    folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB0000000001110000
/Outbox                                   folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB000000000010B0000
/PersonMetadata                           folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB000044E9DCFA0000
/Sent Items                               folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB0000000001090000
/Tasks                                    folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB0000000001120000
/Yammer Root                              folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB0000000001450000
/Yammer Root/Feeds                        folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB0000000001480000
/Yammer Root/Inbound                      folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB0000000001460000
/Yammer Root/Outbound                     folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB0000000001470000
/Recoverable Items                        folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB0000000001140000
/Calendar Logging                         folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB0000000001180000
/Deletions                                folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB0000000001150000
/Purges                                   folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB0000000001170000
/Versions                                 folderid:FDB58AF45BAF2F4A8CFD98F5396C46EB0000000001160000
```

The example in Step 2 shows the query used to search the Purges subfolder in the user's Recoverable Items folder.

Script output for site folders

If you're getting the path of the **documentlink** property from SharePoint or OneDrive for Business sites, the script connects to Security & Compliance PowerShell, creates a new Content Search that searches the site for folders, and then displays a list of the folders located in the specified site. The script displays the name of each folder and adds the prefix of **documentlink** to the folder URL. Because the **documentlink** property is a searchable property, you'll use `documentlink:<path>` property:value pair in a search query in Step 2 to search that folder. The script displays a maximum of 200 site folders. If there are more than 200 site folders, the newest ones are displayed.

Here's an example of the output returned by the script for site folders.

```
Waiting for search to complete...
Waiting for search action to complete...
DocumentLink:"https://alpinehouse-my.sharepoint.com/personal/admin_alpinehouse_onmicrosoft_com/Documents/Shared with Everyone"
DocumentLink:"https://alpinehouse-my.sharepoint.com/personal/admin_alpinehouse_onmicrosoft_com/Documents/Personal Info"
DocumentLink:"https://alpinehouse-my.sharepoint.com/personal/admin_alpinehouse_onmicrosoft_com/Documents/Private"
DocumentLink:"https://alpinehouse-my.sharepoint.com/personal/admin_alpinehouse_onmicrosoft_com/Documents/Naughty Users"
```

Step 2: Use a folder ID or documentlink to perform a targeted collection

After you've run the script to collect a list of folder IDs or document links for a specific user, the next step to go to the Microsoft 365 compliance center and create a new Content Search to search a specific folder. You'll use the `folderid:<folderid>` or `documentlink:<path>` property:value pair in the search query that you configure in the

Content Search keyword box (or as the value for the *ContentMatchQuery* parameter if you use the **New-ComplianceSearch** cmdlet). You can combine the `folderid` or `documentlink` property with other search parameters or search conditions. If you only include the `folderid` or `documentlink` property in the query, the search will return all items located in the specified folder.

1. Go to https://compliance.microsoft.com and sign in using the account and credentials that you used to run the script in Step 1.
2. In the left pane of the compliance center, click **Show all** > **Content search**, and then click **New search**.
3. In the **Keywords** box, paste the `folderid:<folderid>` or `documentlink:<path>` value that was returned by the script in Step 1.

   For example, the query in the following screenshot will search for any item in the Purges subfolder in the user's Recoverable Items folder (the value of the `folderid` property for the Purges subfolder is shown in the screenshot in Step 1):

   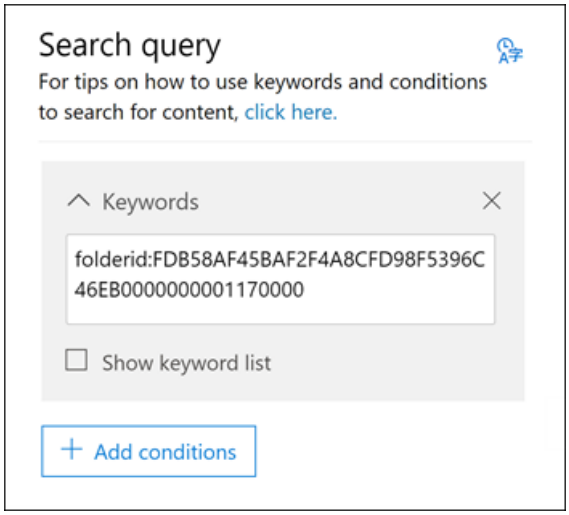

4. Under **Locations**, select **Specific locations** and then click **Modify**.
5. Do one of the following, based on whether you're searching a mailbox folder or a site folder:

   - Next to **Exchange email**, click **Choose users, groups, or teams** and then add the same mailbox that you specified when you ran the script in Step 1.

     Or

   - Next to **SharePoint sites**, click **Choose sites** and then add the same site URL that you specified when you ran the script in Step 1.

6. After you save the content location to search, click **Save & run**, type a name for the Content Search, and then click **Save** to start the targeted collection search.

Examples of search queries for targeted collections

Here are some examples of using the `folderid` and `documentlink` properties in a search query to perform a targeted collection. Placeholders are used for `folderid:<folderid>` and `documentlink:<path>` to save space.

- This example searches three different mailbox folders. You could use similar query syntax to search the hidden folders in a user's Recoverable Items folder.

  ```
  folderid:<folderid> OR folderid:<folderid> OR folderid:<folderid>
  ```

- This example searches a mailbox folder for items that contain an exact phrase.

```
folderid:<folderid> AND "Contoso financial results"
```

- This example searches a site folder (and any subfolders) for documents that contain the letters "NDA" in the title.

```
documentlink:<path> AND filename:nda
```

- This example searches a site folder (and any subfolder) for documents there were changed within a date range.

```
documentlink:<path> AND (lastmodifiedtime>=01/01/2017 AND lastmodifiedtime<=01/21/2017)
```

More information

Keep the following things in mind when using the script in this article to perform targeted collections.

- The script doesn't remove any folders from the results. So some folders listed in the results might be unsearchable (or return zero items) because they contain system-generated content or because they only contain subfolders and not mailbox items.
- This script only returns folder information for the user's primary mailbox. It doesn't return information about folders in the user's archive mailbox. To return information about folders in the user's archive mailbox, you can edit the script. To do this, change the line `$folderStatistics = Get-MailboxFolderStatistics $emailAddress` to `$folderStatistics = Get-MailboxFolderStatistics $emailAddress -Archive` and then save and run the edited script. This change will return the folder IDs for folders and subfolders in the user's archive mailbox. To search the entire archive mailbox, you can connect all folder ID property:value pairs with an `OR` operator in a search query.
- When searching mailbox folders, only the specified folder (identified by its `folderid` property) will be searched; subfolders won't be searched. To search subfolders, you need to use the folder ID for the subfolder that you want to search.
- When searching site folders, the folder (identified by its `documentlink` property) and all subfolders will be searched.
- When exporting the results of a search in which you only specified the `folderid` property in the search query, you can choose the first export option, "All items, excluding ones that have an unrecognized format, are encrypted, or weren't indexed for other reasons." All items in the folder will always be exported regardless of their indexing status because the folder ID is always indexed.

Use Content Search to search the mailbox and OneDrive for Business site for a list of users

11/2/2020 • 7 minutes to read • Edit Online

The Security & Compliance Center provides a number of Windows PowerShell cmdlets that let you automate time-consuming eDiscovery-related tasks. Currently, creating a Content Search in the Security & Compliance Center to search a large number of custodian content locations takes time and preparation. Before you create a search, you have to collect the URL for each OneDrive for Business site and then add each mailbox and OneDrive for Business site to the search. In future releases, this will be easier to do in the Security & Compliance Center. Until then, you can use the script in this article to automate this process. This script prompts you for the name of your organization's MySite domain (for example, **contoso** in the URL `https://contoso-my.sharepoint.com`), a list of user email addresses, the name of the new Content Search, and the search query to use. The script gets the OneDrive for Business URL for each user in the list, and then it creates and starts a Content Search that searches the mailbox and OneDrive for Business site for each user in the list, using the search query that you provide.

Permissions and script information

- You have to be a member of the eDiscovery Manager role group in the Security & Compliance Center and a SharePoint Online global administrator to run the script in Step 3.
- Be sure to save the list of users that you create in Step 2 and the script in Step 3 to the same folder. That will make it easier to run the script.
- The script includes minimal error handling. Its primary purpose is to quickly and easily search the mailbox and OneDrive for Business site of each user.
- The sample scripts provided in this topic aren't supported under any Microsoft standard support program or service. The sample scripts are provided AS IS without warranty of any kind. Microsoft further disclaims all implied warranties including, without limitation, any implied warranties of merchantability or of fitness for a particular purpose. The entire risk arising out of the use or performance of the sample scripts and documentation remains with you. In no event shall Microsoft, its authors, or anyone else involved in the creation, production, or delivery of the scripts be liable for any damages whatsoever (including, without limitation, damages for loss of business profits, business interruption, loss of business information, or other pecuniary loss) arising out of the use of or inability to use the sample scripts or documentation, even if Microsoft has been advised of the possibility of such damages.

Step 1: Install the SharePoint Online Management Shell

The first step is to install the SharePoint Online Management Shell. You don't have to use the shell in this procedure, but you have to install it because it contains pre-requisites required by the script that you run in Step 3. These prerequisites allow the script to communicate with SharePoint Online to get the URLs for the OneDrive for Business sites.

Go to Set up the SharePoint Online Management Shell Windows PowerShell environment and perform Step 1 and Step 2 to install the SharePoint Online Management Shell.

Step 2: Generate a list of users

The script in Step 3 will create a Content Search to search the mailboxes and OneDrive accounts for a list of users. You can just type the email addresses in a text file, or you can run a command in Windows PowerShell to get a list of email addresses and save them to a file (located in same folder that you'll save the script to in Step 3).

Here's an Exchange Online PowerShell command that you can runt to get a list of email addresses for all users in your organization and save it to a text file named `Users.txt`.

```
Get-Mailbox -ResultSize unlimited -Filter { RecipientTypeDetails -eq 'UserMailbox'} | Select-Object
PrimarySmtpAddress > Users.txt
```

After you run this command, be sure to open the file and remove the header that contains the property name, `PrimarySmtpAddress`. The text file should just contain a list of email addresses, and nothing else. Make sure there are no blank rows before or after the list of email addresses.

Step 3: Run the script to create and start the search

When you run the script in this step, it will prompt you for the following information. Be sure to have this information ready before you run the script.

- **Your user credentials** - The script will use your credentials to access SharePoint Online to get the OneDrive for Business URLs and to connect to the Security & Compliance Center with remote PowerShell.
- **Name of your MySite domain** - The MySite domain is the domain that contains all the OneDrive for Business sites in your organization. For example, if the URL for your MySite domain is **https://contoso-my.sharepoint.com**, then you would enter `contoso` when the script prompts you for the name of your MySite domain.
- **Pathname of the text file from Step 2** - The pathname of the text file that you created in Step 2. If the text file and the script are located in the same folder, then enter the name of the text file. Otherwise, enter the complete pathname for the text file.
- **Name of the Content Search** - The name of the Content Search that will be created by the script.
- **Search query** - The search query that will be used with the Content Search is created and run. For more information about search queries, see Keyword queries and search conditions for Content Search.

To run the script:

1. Save the following text to a Windows PowerShell script file by using a filename suffix of .ps1; for example, `SearchEXOOD4B.ps1`. Save the file to the same folder where you saved the list of users in Step 2.

```
# This PowerShell script will prompt you for the following information:
#    * Your user credentials
#    * The name of your organization's MySite domain
#    * The pathname for the text file that contains a list of user email addresses
#    * The name of the Content Search that will be created
#    * The search query string
# The script will then:
#    * Find the OneDrive for Business site for each user in the text file
#    * Create and start a Content Search using the above information
# Get user credentials
if (!$credentials)
{
    $credentials = Get-Credential
}
# Get the user's MySite domain name.  We use this to create the admin URL and root URL for OneDrive for
Business
```

```
$mySiteDomain = Read-Host "What is your organization's MySite domain?  For example,  'contoso' for
'https://contoso-my.sharepoint.com'"
$AdminUrl = "https://$mySiteDomain-admin.sharepoint.com"
$mySiteUrlRoot = "https://$mySiteDomain-my.sharepoint.com"
# Get other required information
$inputfile = read-host "Enter the file name of the text file that contains the email addresses for the users
you want to search"
$searchName = Read-Host "Enter the name for the new search"
$searchQuery = Read-Host "Enter the search query you want to use"
$emailAddresses = Get-Content $inputfile | where {$_ -ne ""}  | foreach{ $_.Trim() }
# Connect to Office 365
if (!$s -or !$a)
{
    $s = New-PSSession -ConfigurationName Microsoft.Exchange -ConnectionUri
"https://ps.compliance.protection.outlook.com/powershell-liveid" -Credential $credentials -Authentication
Basic -AllowRedirection -SessionOption (New-PSSessionOption -SkipCACheck -SkipCNCheck -SkipRevocationCheck)
    $a = Import-PSSession $s -AllowClobber
    if (!$s)
    {
        Write-Error "Could not create PowerShell session."
        return;
    }
}
# Load the SharePoint assemblies from the SharePoint Online Management Shell
# To install, go to https://go.microsoft.com/fwlink/p/?LinkId=255251
if (!$SharePointClient -or !$SPRuntime -or !$SPUserProfile)
{
    $SharePointClient = [System.Reflection.Assembly]::LoadWithPartialName("Microsoft.SharePoint.Client")
    $SPRuntime = [System.Reflection.Assembly]::LoadWithPartialName("Microsoft.SharePoint.Client.Runtime")
    $SPUserProfile =
[System.Reflection.Assembly]::LoadWithPartialName("Microsoft.SharePoint.Client.UserProfiles")
    if (!$SharePointClient)
    {
        Write-Error "SharePoint Online Management Shell isn't installed, please install from:
https://go.microsoft.com/fwlink/p/?LinkId=255251 and then run this script again"
        return;
    }
}
if (!$spCreds)
{
    $spCreds = New-Object Microsoft.SharePoint.Client.SharePointOnlineCredentials($credentials.UserName,
$credentials.Password)
}
# Add the path of the User Profile Service to the SPO admin URL, then create a new webservice proxy to
access it
$proxyaddr = "$AdminUrl/_vti_bin/UserProfileService.asmx?wsdl"
$UserProfileService= New-WebServiceProxy -Uri $proxyaddr -UseDefaultCredential False
$UserProfileService.Credentials = $credentials
# Take care of auth cookies
$strAuthCookie = $spCreds.GetAuthenticationCookie($AdminUrl)
$uri = New-Object System.Uri($AdminUrl)
$container = New-Object System.Net.CookieContainer
$container.SetCookies($uri, $strAuthCookie)
$UserProfileService.CookieContainer = $container
Write-Host "Getting each user's OneDrive for Business URL"
$urls = @()
foreach($emailAddress in $emailAddresses)
{
    try
    {
        $prop = $UserProfileService.GetUserProfileByName("i:0#.f|membership|$emailAddress") | Where-Object {
$_.Name -eq "PersonalSpace" }
        $url = $prop.values[0].value
        $furl = $mySiteUrlRoot + $url
        $urls += $furl
        Write-Host "-$emailAddress => $furl"
    }
    catch
    {
```

```
        Write-Warning "Could not locate OneDrive for $emailAddress"
    }
}
Write-Host "Creating and starting the search"
$search = New-ComplianceSearch -Name $searchName -ExchangeLocation $emailAddresses -SharePointLocation $urls
-ContentMatchQuery $searchQuery
# Finally, start the search and then display the status
if($search)
{
    Start-ComplianceSearch $search.Name
    Get-ComplianceSearch $search.Name
}
```

2. Open Windows PowerShell and go to the folder where you saved the script and the list of users from Step 2.

3. Start the script; for example:

   ```
   .\SearchEXOOD4B.ps1
   ```

4. When prompted for your credentials, enter your email address and password, and then click **OK**.

5. Enter following information when prompted by the script. Type each piece of information and then press **Enter**.

   - The name of your MySite domain.
   - The pathname of the text file that contains the list of users.
   - A name for the Content Search.
   - The search query (leave this blank to return all items in the content locations).

   The script gets the URLs for each OneDrive for Business site and then creates and starts the search. You can either run the **Get-ComplianceSearch** cmdlet in Security & Compliance Center PowerShell to display the search statistics and results, or you can go to the **Content search** page in the Security & Compliance Center to view information about the search.

Create, report on, and delete multiple Content Searches

11/2/2020 • 11 minutes to read • Edit Online

Quickly creating and reporting discovery searches is often an important step in eDiscovery and investigations when you're trying to learn about the underlying data, and the richness and quality of your searches. To help you do this, the Security & Compliance Center PowerShell offers a set of cmdlets to automate time-consuming Content Search tasks. These scripts provide a quick and easy way to create a number of searches, and then run reports of the estimated search results that can help you determine the quantity of data in question. You can also use the scripts to create different versions of searches to compare the results each one produces. These scripts can help you to quickly and efficiently identify and cull your data.

Before you create a Content Search

- You have to be a member of the eDiscovery Manager role group in the Security & Compliance Center to run the scripts that are described in this topic.
- To collect a list of the URLs for the OneDrive for Business sites in your organization that you can add to the CSV file in Step 1, see Create a list of all OneDrive locations in your organization.
- Be sure to save all the files that you create in this topic to the same folder. That will make it easier to run the scripts.
- The scripts include minimal error handling. Their primary purpose is to quickly create, report on, and delete multiple Content Searches.
- The sample scripts provided in this topic aren't supported under any Microsoft standard support program or service. The sample scripts are provided AS IS without warranty of any kind. Microsoft further disclaims all implied warranties including, without limitation, any implied warranties of merchantability or of fitness for a particular purpose. The entire risk arising out of the use or performance of the sample scripts and documentation remains with you. In no event shall Microsoft, its authors, or anyone else involved in the creation, production, or delivery of the scripts be liable for any damages whatsoever (including, without limitation, damages for loss of business profits, business interruption, loss of business information, or other pecuniary loss) arising out of the use of or inability to use the sample scripts or documentation, even if Microsoft has been advised of the possibility of such damages.

Step 1: Create a CSV file that contains information about the searches you want to run

The comma separated value (CSV) file that you create in this step contains a row for each user that want to search. You can search the user's Exchange Online mailbox (which includes the archive mailbox, if it's enabled) and their OneDrive for Business site. Or you can search just the mailbox or the OneDrive for Business site. You can also search any site in your SharePoint Online organization. The script that you run in Step 3 will create a separate search for each row in the CSV file.

1. Copy and paste the following text into a .txt file using NotePad. Save this file to a folder on your local computer. You'll save the other scripts to this folder as well.

```
ExchangeLocation,SharePointLocation,ContentMatchQuery,StartDate,EndDate
sarad@contoso.onmicrosoft.com,https://contoso-
my.sharepoint.com/personal/sarad_contoso_onmicrosoft_com,(lawsuit OR legal),1/1/2000,12/31/2005
sarad@contoso.onmicrosoft.com,https://contoso-
my.sharepoint.com/personal/sarad_contoso_onmicrosoft_com,(lawsuit OR legal),1/1/2006,12/31/2010
sarad@contoso.onmicrosoft.com,https://contoso-
my.sharepoint.com/personal/sarad_contoso_onmicrosoft_com,(lawsuit OR legal),1/1/2011,3/21/2016
,https://contoso.sharepoint.com/sites/contoso,,,3/21/2016
,https://contoso-my.sharepoint.com/personal/davidl_contoso_onmicrosoft_com,,1/1/2015,
,https://contoso-my.sharepoint.com/personal/janets_contoso_onmicrosoft_com,,1/1/2015,
```

The first row, or header row, of the file lists the parameters that will be used by **New-ComplianceSearch** cmdlet (in the script in Step 3) to create a new Content Searches. Each parameter name is separated by a comma. Make sure there aren't any spaces in the header row. Each row under the header row represents the parameter values for each search. Be sure to replace the placeholder data in the CSV file with your actual data.

2. Open the .txt file in Excel, and then use the information in the following table to edit the file with information for each search.

PARAMETER	DESCRIPTION
`ExchangeLocation`	The SMTP address of the user's mailbox.
`SharePointLocation`	The URL for the user's OneDrive for Business site or the URL for any site in your organization. For the URL for OneDrive for Business sites, use this format: `https://<your organization>-my.sharepoint.com/personal/<user alias>_<your organization>_onmicrosoft_com`. For example, `https://contoso-my.sharepoint.com/personal/sarad_contoso_onmicrosoft_com`.
`ContentMatchQuery`	The search query for the search. For more information about creating a search query, see Keyword queries and search conditions for Content Search.
`StartDate`	For email, the date on or after a message was received by a recipient or sent by the sender. For documents on SharePoint or OneDrive for Business sites, the date on or after a document was last modified.
`EndDate`	For email, the date on or before a message was sent by a sent by the user. For documents on SharePoint or OneDrive for Business sites, the date on or before a document was last modified.

3. Save the Excel file as a CSV file to a folder on your local computer. The script that you create in Step 3 will use the information in this CSV file to create the searches.

Step 2: Connect to Security & Compliance Center PowerShell

The next step is to connect to Security & Compliance Center PowerShell for your organization. For step-by-step instructions, see Connect to Security & Compliance Center PowerShell.

Step 3: Run the script to create and start the searches

The script in this step will create a separate Content Search for each row in the CSV file that you created in Step 1. When you run this script, you'll be prompted for two values:

- **Search Group ID** - This name provides an easy way to organize the searches that are created from the CSV file. Each search that's created is named with the Search Group ID, and then a number is appended to the search name. For example, if you enter **ContosoCase** for the Search Group ID, then the searches are named **ContosoCase_1**, **ContosoCase_2**, **ContosoCase_3**, and so on. Note that the name you type is case sensitive. When you use the Search Group ID in Step 4 and Step 5, you have to use the same case as you did when you created it.
- **CSV file** - The name of the CSV file that you created in Step 1. Be sure to include the use the full filename, include the .csv file extension; for example, `ContosoCase.csv`.

To run the script:

1. Save the following text to a Windows PowerShell script file by using a filename suffix of .ps1; for example, `CreateSearches.ps1`. Save the file to the same folder where you saved the other files.

```
# Get the Search Group ID and the location of the CSV input file
$searchGroup = Read-Host 'Search Group ID'
$csvFile = Read-Host 'Source CSV file'

# Do a quick check to make sure our group name will not collide with other searches
$searchCounter = 1
import-csv $csvFile |
  ForEach-Object{

  $searchName = $searchGroup +'_' + $searchCounter
  $search = Get-ComplianceSearch $searchName -EA SilentlyContinue
  if ($search)
  {
     Write-Error "The Search Group ID conflicts with existing searches.  Please choose a search group
name and restart the script."
     return
  }
  $searchCounter++
}

$searchCounter = 1
import-csv $csvFile |
  ForEach-Object{

  # Create the query
  $query = $_.ContentMatchQuery
  if(($_.StartDate -or $_.EndDate))
  {
       # Add the appropriate date restrictions.  NOTE: Using the Date condition property here because
it works across Exchange, SharePoint, and OneDrive for Business.
       # For Exchange, the Date condition property maps to the Sent and Received dates; for
SharePoint and OneDrive for Business, it maps to Created and Modified dates.
       if($query)
       {
           $query += " AND"
       }
       $query += " ("
       if($_.StartDate)
       {
           $query += "Date >= " + $_.StartDate
       }
       if($_.EndDate)
       {
           if($_.StartDate)
```

```
            {
                $query += " AND "
            }
            $query += "Date <= " + $_.EndDate
        }
        $query += ")"
    }

    # -ExchangeLocation can't be set to an empty string, set to null if there's no location.
    $exchangeLocation = $null
    if ( $_.ExchangeLocation)
    {
          $exchangeLocation = $_.ExchangeLocation
    }

    # Create and run the search
    $searchName = $searchGroup +'_' + $searchCounter
    Write-Host "Creating and running search: " $searchName -NoNewline
    $search = New-ComplianceSearch -Name $searchName -ExchangeLocation $exchangeLocation -
SharePointLocation $_.SharePointLocation -ContentMatchQuery $query

    # Start and wait for each search to complete
    Start-ComplianceSearch $search.Name
    while ((Get-ComplianceSearch $search.Name).Status -ne "Completed")
    {
       Write-Host " ." -NoNewline
       Start-Sleep -s 3
    }
    Write-Host ""

    $searchCounter++
}
```

2. In Windows PowerShell, go to the folder where you saved the script in the previous step, and then run the script; for example:

```
.\CreateSearches.ps1
```

3. At the **Search Group ID** prompt, type a search group name, and then press **Enter**; for example, `ContosoCase`. Remember that this name is case sensitive, so you'll have to type it the same way in the subsequent steps.

4. At the **Source CSV file** prompt, type the name of the CSV file, including the .csv file extension; for example, `ContosoCase.csv`.

5. Press **Enter** to continue running the script.

   The script displays the progress of creating and running the searches. When the script is complete, it returns to the prompt.

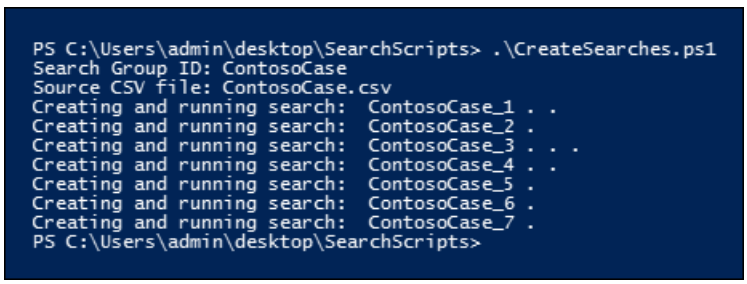

Step 4: Run the script to report the search estimates

After you create the searches, the next step is to run a script that displays a simple report of the number of

search hits for each search that was created in Step 3. The report also includes the size of results for each search, and the total number of hits and total size of all searches. When you run the reporting script, you'll be prompted for the Search Group ID, and a CSV filename if you want to save the report to a CSV file.

1. Save the following text to a Windows PowerShell script file by using a filename suffix of .ps1; for example, `SearchReport.ps1`. Save the file to the same folder where you saved the other files.

```
$searchGroup = Read-Host 'Search Group ID'
$outputFile = Read-Host 'Enter a file name or file path to save the report to a .csv file. Leave
blank to only display the report'
$searches = Get-ComplianceSearch | ?{$_.Name -clike $searchGroup + "_*"}
$allSearchStats = @()
foreach ($partialObj in $searches)
{
   $search = Get-ComplianceSearch $partialObj.Name
   $sizeMB = [System.Math]::Round($search.Size / 1MB, 2)
   $searchStatus = $search.Status
   if($search.Errors)
   {
       $searchStatus = "Failed"
   }elseif($search.NumFailedSources -gt 0)
   {
       $searchStatus = "Failed Sources"
   }
   $searchStats = New-Object PSObject
   Add-Member -InputObject $searchStats -MemberType NoteProperty -Name Name -Value $search.Name
   Add-Member -InputObject $searchStats -MemberType NoteProperty -Name ContentMatchQuery -Value
$search.ContentMatchQuery
   Add-Member -InputObject $searchStats -MemberType NoteProperty -Name Status -Value $searchStatus
   Add-Member -InputObject $searchStats -MemberType NoteProperty -Name Items -Value $search.Items
   Add-Member -InputObject $searchStats -MemberType NoteProperty -Name "Size" -Value $search.Size
   Add-Member -InputObject $searchStats -MemberType NoteProperty -Name "Size(MB)" -Value $sizeMB
   $allSearchStats += $searchStats
}
# Calculate the totals
$allItems = ($allSearchStats | Measure-Object Items -Sum).Sum
# Convert the total size to MB and round to the nearst 100th
$allSize = ($allSearchStats | Measure-Object 'Size' -Sum).Sum
$allSizeMB = [System.Math]::Round($allSize  / 1MB, 2)
# Get the total successful searches and total of all searches
$allSuccessCount = ($allSearchStats |?{$_.Status -eq "Completed"}).Count
$allCount = $allSearchStats.Count
$allStatus = [string]$allSuccessCount + " of " + [string]$allCount
# Totals Row
$totalSearchStats = New-Object PSObject
Add-Member -InputObject $totalSearchStats -MemberType NoteProperty -Name Name -Value "Total"
Add-Member -InputObject $totalSearchStats -MemberType NoteProperty -Name Status -Value $allStatus
Add-Member -InputObject $totalSearchStats -MemberType NoteProperty -Name Items -Value $allItems
Add-Member -InputObject $totalSearchStats -MemberType NoteProperty -Name "Size(MB)" -Value $allSizeMB
$allSearchStats += $totalSearchStats
# Just get the columns we're interested in showing
$allSearchStatsPrime = $allSearchStats | Select-Object Name, Status, Items, "Size(MB)",
ContentMatchQuery
# Print the results to the screen
$allSearchStatsPrime |ft -AutoSize -Wrap
# Save the results to a CSV file
if ($outputFile)
{
   $allSearchStatsPrime | Export-Csv -Path $outputFile -NoTypeInformation
}
```

2. In Windows PowerShell, go to the folder where you saved the script in the previous step, and then run the script; for example:

```
.\SearchReport.ps1
```

3. At the **Search Group ID** prompt, type a search group name, and then press **Enter**; for example `ContosoCase`. Remember that this name is case sensitive, so you'll have to type it the same way you did when you ran the script in Step 3.

4. At the **File path to save the report to a CSV file (leave blank to just display the report)** prompt, type a file name of complete filename path (including the .csv file extension) if you want to save the report to a CSV file. name of the CSV file, including the .csv file extension. For example, you could type `ContosoCaseReport.csv` to save it to the current directory or you could type `C:\Users\admin\OneDrive for Business\ContosoCase\ContosoCaseReport.csv` to save it to a different folder. You can also leave the prompt blank to display the report but not save it to a file.

5. Press **Enter**.

   The script displays the progress of creating and running the searches. When the script is complete, the report is displayed.

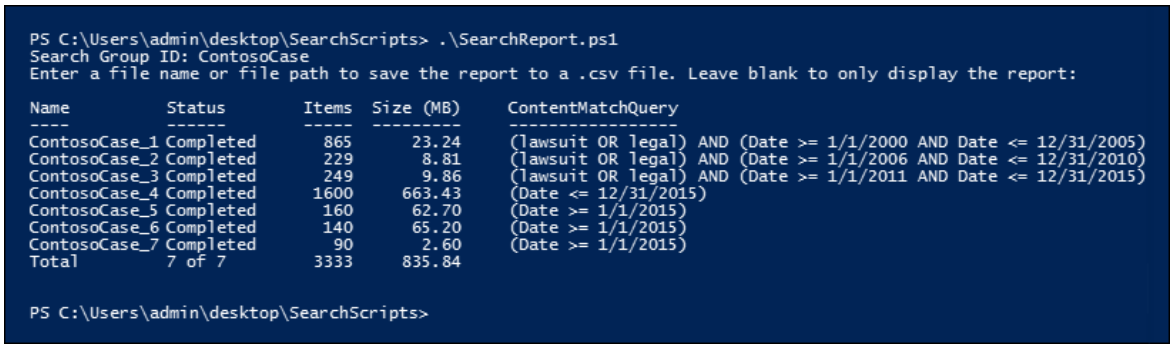

> **NOTE**
>
> If the same mailbox or site is specified as a content location in more than one search in a search group, the total results estimate in the report (for both the number of items and the total size) might include results for the same items. That's because the same email message or document will be counted more than once if it matches the query for different searches in the search group.

Step 5: Run the script to delete the searches

Because you might be creating a lot of searches, this last script just makes it easy to quickly delete the searches you created in Step 3. Like the other scripts, this one also prompts you for the Search Group ID. All searches with the Search Group ID in the search name will be deleted when you run this script.

1. Save the following text to a Windows PowerShell script file by using a filename suffix of .ps1; for example, `DeleteSearches.ps1`. Save the file to the same folder where you saved the other files.

```
# Delete all searches in a search group
$searchGroup = Read-Host 'Search Group ID'
Get-ComplianceSearch |
   ForEach-Object{
   # If the name matches the search group name pattern (case sensitive), delete the search
   if ($_.Name -cmatch $searchGroup + "_\d+")
   {
       Write-Host "Deleting search: " $_.Name
       Remove-ComplianceSearch $_.Name -Confirm:$false
   }
}
```

2. In Windows PowerShell, go to the folder where you saved the script in the previous step, and then run the script; for example:

```
.\DeleteSearches.ps1
```

3. At the **Search Group ID** prompt, type a search group name for the searches that you want to delete, and then press **Enter**; for example, `ContosoCase`. Remember that this name is case sensitive, so you'll have to type it the same way you did when you ran the script in Step 3.

   The script displays the name of each search that's deleted.

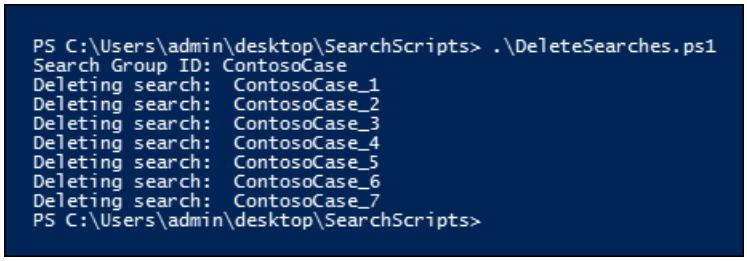

Clone a Content Search

11/2/2020 • 5 minutes to read • Edit Online

Creating a Content Search in the compliance center in Office 365 or Microsoft 365 that searches many mailboxes or SharePoint and OneDrive for Business sites can take a while. Specifying the sites to search can also be prone to errors if you mistype a URL. To avoid these issues, you can use the Windows PowerShell script in this article to quickly clone an existing Content Search. When you clone a search, a new search (with a different name) is created that contains the same properties (such as the content locations and the search query) as the original search. Then you can edit the new search by changing the keyword query or the date range, and run it.

Why clone Content Searches?

- To compare the results of different keyword search queries run on the same content locations.
- To save you from having to reenter a large number of content locations when you create a new search.
- To decrease the size of the search results. For example, if you have a search that returns too many results to export, you can clone the search and then add a search condition based on a date range to reduce the number of search results.

Script information

- You have to be a member of the eDiscovery Manager role group in the Security & Compliance Center to run the script described in this topic.
- The script includes minimal error handling. The primary purpose of the script is to quickly clone a content search.
- The script creates a new Content Search, but doesn't start it.
- This script takes into account whether the Content Search that you're cloning is associated with an eDiscovery case. If the search is associated with a case, the new search will also be associated with the same case. If the existing search isn't associated with a case, the new search will be listed on the **Content search** page in the compliance center.
- The sample script provided in this topic isn't supported under any Microsoft standard support program or service. The sample script is provided AS IS without warranty of any kind. Microsoft further disclaims all implied warranties including, without limitation, any implied warranties of merchantability or of fitness for a particular purpose. The entire risk arising out of the use or performance of the sample script and documentation remains with you. In no event shall Microsoft, its authors, or anyone else involved in the creation, production, or delivery of the scripts be liable for any damages whatsoever (including, without limitation, damages for loss of business profits, business interruption, loss of business information, or other pecuniary loss) arising out of the use of or inability to use the sample scripts or documentation, even if Microsoft has been advised of the possibility of such damages.

Step 1: Run the script to clone a search

The script in this step will create a new Content Search by cloning an existing one. When you run this script, you'll be prompted for the following information:

- **Your user credentials** - The script will use your credentials to connect to the Security & Compliance Center for your organization with Windows PowerShell. As previously stated, you have to be a member of the eDiscovery Manager role group in the Security & compCompliance Center to run the script.

- **The name of the existing search** - This is the Content Search that you want to clone.
- **The name of the new search that will be created** - If you leave this value blank, the script will create a name for the new search that is based on the name of the search that you're cloning.

To clone a search:

1. Save the following text to a Windows PowerShell script file by using a filename suffix of .ps1; for example, `CloneSearch.ps1`.

```
# This PowerShell script clones an existing content search in the Security &amp;amp; Compliance Center.
# Get login credentials from the user
if(!$UserCredential)
{
    $UserCredential = Get-Credential
    $Session = New-PSSession -ConfigurationName Microsoft.Exchange -ConnectionUri
https://ps.compliance.protection.outlook.com/powershell-liveid -Credential $UserCredential -Authentication
Basic -AllowRedirection
    if (!$Session)
    {
        Write-Error "Couldn't create a remote PowerShell session."
        return
    }
    Import-PSSession $Session -AllowClobber -DisableNameChecking
    $Host.UI.RawUI.WindowTitle = $UserCredential.UserName + " (Security & Compliance Center)"
}
# Ask for the name of the search you want to clone
$searchName = Read-Host 'Enter the name of the search that you want to clone'
# Ask for the name of the new search
$newSearchName = Read-Host 'Enter a name for the new search [leave blank to automatically generate a name]'
$originalSearch = Get-ComplianceSearch $searchName -EA SilentlyContinue
# Make sure we have a valid search before continuing
if(!$originalSearch)
{
    Write-Error "Couldn't find search: $searchName"
    return
}
$searchNameCounter = 1
# Find a suitable name for the new search
while(!$newSearchName)
{
    $newSearchName = $originalSearch.Name + "_" + $searchNameCounter
    $tempSearch = Get-ComplianceSearch $newSearchName -EA SilentlyContinue
    if ($tempSearch)
    {
        $newSearchName = $null
        $searchNameCounter++
    }
}
$caseName
# Determine if the search is part of a case; if so get the case name
if ($originalSearch.CaseId)
{
    $searchCase = Get-ComplianceCase $originalSearch.CaseId
    $caseName = $searchCase.Name
}
# Need to cast this value as a Boolean the old fashion way
$allowNotFoundExchangeLocationsEnabled = $false
if ($originalSearch.AllowNotFoundExchangeLocationsEnabled)
{
    $allowNotFoundExchangeLocationsEnabled = $true
}
$newSearch = New-ComplianceSearch -Name $newSearchName -AllowNotFoundExchangeLocationsEnabled
$allowNotFoundExchangeLocationsEnabled -Case $caseName -ContentMatchQuery $originalSearch.ContentMatchQuery
-Description $originalSearch.Description -ExchangeLocation $originalSearch.ExchangeLocation -
ExchangeLocationExclusion $originalSearch.ExchangeLocationExclusion -Language $originalSearch.Language -
SharePointLocation $originalSearch.SharePointLocation -SharePointLocationExclusion
$originalSearch.SharePointLocationExclusion -PublicFolderLocation $originalSearch.PublicFolderLocation
if ($newSearch)
{
    Write-Host $newSearch.Name "was successfully created" -ForegroundColor Yellow
}
```

2. Open Windows PowerShell and go to the folder where you saved the script.

3. Run the script; for example:

```
.\CloneSearch.ps1
```

4. When prompted for your credentials, enter your email address and password, and then click **OK**.
5. Enter following information when prompted by the script. Type each piece of information and then press **Enter**.
   - The name of the existing search.
   - The name of the new search.

   The script creates the new Content Search, but doesn't start it. This gives you a chance to edit and run the search in the next step. You can view the properties of the new search by running the **Get-ComplianceSearch** cmdlet or by going to the **Content search** or **eDiscovery** page in the compliance center, depending on whether the new search is associated with a case.

Step 2: Edit and run the cloned search in the compliance center

After you run the script to clone an existing Content Search, the next step is to go to the compliance center to edit and run the new search. As previously stated, you can edit a search by changing the keyword search query and adding or removing search conditions. For more information, see:

- Content Search in Office 365
- Keyword queries and search conditions for Content Search
- eDiscovery cases

Get started with Core eDiscovery

2/18/2021 • 7 minutes to read • Edit Online

Core eDiscovery in Microsoft 365 provides a basic eDiscovery tool that organizations can use to search and export content in Microsoft 365 and Office 365. You can also use Core eDiscovery to place an eDiscovery hold on content locations, such as Exchange mailboxes, SharePoint sites, OneDrive accounts, and Microsoft Teams. Nothing is needed to deploy Core eDiscovery, but there are some prerequisite tasks that an IT admin and eDiscovery manager have to complete before your organization can start using Core eDiscovery to search, export, and preserve content.

This article discusses the steps necessary to set up Core eDiscovery. This includes ensuring the proper licensing required to access Core eDiscovery and place an eDiscovery hold on content locations, as well as assigning permissions to your IT, legal, and investigation team so they can access and manage cases. This article also provides a high-level overview of using cases to search for and export content.

Step 1: Verify and assign appropriate licenses

Licensing for Core eDiscovery requires the appropriate organization subscription and per-user licensing.

- **Organization subscription:** To access Core eDiscovery in the Microsoft 365 compliance center or the Office 365 Security & Compliance Center and use the hold and export features, your organization must have a Microsoft 365 E3 or Office 365 E3 subscription or higher.
- **Per-user licensing:** To place an eDiscovery hold on mailboxes and sites, a user must be assigned one of the following licenses, depending on your organization subscription:

  - A Microsoft 365 E3 or Office 365 E3 license or higher

  OR

  - Office 365 E1 license with an Exchange Online Plan 2 or Exchange Online Archiving add-on license

  AND

  - Office 365 E1 license with an SharePoint Online Plan 2 or OneDrive for Business Plan 2 add-on license

  For information about how to assign licenses, see Assign licenses to users.

For information about licensing:

- Download and see the "Discover & Respond" solution in the Microsoft 365 Compliance Licensing Comparison.
- See the Security & Compliance Center service description.

Step 2: Assign eDiscovery permissions

To access Core eDiscovery or be added as a member of a Core eDiscovery case, a user must be assigned the appropriate permissions. Specifically, a user must be added as a member of the eDiscovery Manager role group in the Office 365 Security & Compliance Center. Members of this role group can create and manage Core eDiscovery cases. They can add and remove members, place an eDiscovery hold on users, create and edit searches, and export content from a Core eDiscovery case.

Complete the following steps to add users to the eDiscovery Manager role group:

1. Go to https://protection.office.com/permissions and sign in using the credentials for an admin account in

your Microsoft 365 or Office 365 organization.

2. On the **Permissions** page, select the **eDiscovery Manager** role group.
3. On the eDiscovery Manager flyout page, click **Edit** next to the **eDiscovery Manager** section.
4. On the **Choose eDiscovery Manager** page in the edit role group wizard, click **Choose Discovery Manager**.
5. Click **Add** then select the checkbox for all users you want to add to the role group.
6. Click **Add** to add the selected users, and then click **Done**.
7. Click **Save** to add the users to the role group, and then click **Close** to complete the step.

More information about the eDiscovery Manager role group

There are two subgroups in the eDiscovery Manager role group. The difference between these subgroups is based on scope.

- **eDiscovery Manager**: Can view and manage the Core eDiscovery cases they create or are a member of. If another eDiscovery Manager creates a case but doesn't add a second eDiscovery Manager as a member of that case, the second eDiscovery Manager won't be able to view or open the case on the Core eDiscovery page in the compliance center. In general, most people in your organization can be added to the eDiscovery Manager subgroup.
- **eDiscovery Administrator**: Can perform all case management tasks that an eDiscovery Manager can do. Additionally, an eDiscovery Administrator can:
  - View all cases that are listed on the Core eDiscovery page.
  - Manage any case in the organization after they add themselves as a member of the case.
  - Access and export case data for any case in the organization.

  Because of the broad scope of access, an organization should have only a few admins who are members of the eDiscovery Administrators subgroup.

For more information about eDiscovery permissions and a description of each role that's assigned to the eDiscovery Manager role group, see Assign eDiscovery permissions.

Step 3: Create a Core eDiscovery case

The next step is to create a case and start using Core eDiscovery. Complete the following steps to create a case and add members. The user who creates the case is automatically added as a member.

1. Go to https://compliance.microsoft.com and sign in using the credentials for a user account that has been assigned the appropriate eDiscovery permissions. Members of the Organization Management role group can also create Core eDiscovery cases.
2. In the left navigation pane of the Microsoft 365 compliance center, click **Show all**, and then click **eDiscovery > Core**.
3. On the **Core eDiscovery** page, click **Create a case**.
4. On the **New case** flyout page, give the case a name (required), and then type an optional case number and description. The case name must be unique in your organization.
5. Click **Save** to create the case.

   The new case is created and displayed on the Core eDiscovery page. You may have to click **Refresh** to display the new case.

Step 4 (optional): Add members to a Core eDiscovery case

If you create a case in Step 3 and you're the only person who will use the case, then you don't have to perform this step. You can start using the case to create eDiscovery holds, search for content, or export search results. Perform this step if you want to give other users (or roles group) access to the case.

1. On the **Core eDiscovery** page in the Microsoft 365 compliance center, click the name of the case that you want to add members to.
2. On the **Manage this case** flyout page, under **Manage members**, click **Add** to add members to the case.

   You can also choose to add role group as members of a case. Under **Manage role groups**, click **Add**. You can only assign the role groups that you are a member of to a case. That's because role groups control who can assign members to an eDiscovery case.
3. In the list of people or role groups that can be added as members of the case, click the check box next to the names of the people (or role groups) that you want to add. If you have a large list of people who can added as members, use the **Search** box to search for a specific person in the list.
4. After you select the people or role groups to add as members of the case, click **Add**.
5. Click **Save** to save the new list of case members.

Explore the Core eDiscovery workflow

To get you started using core eDiscovery, here's a simple workflow of creating eDiscovery holds for people of interest, searching for content that relevant to your investigation, and then exporting that data for further review. In each of these steps, we'll also highlight some extended Core eDiscovery functionality that you can explore.

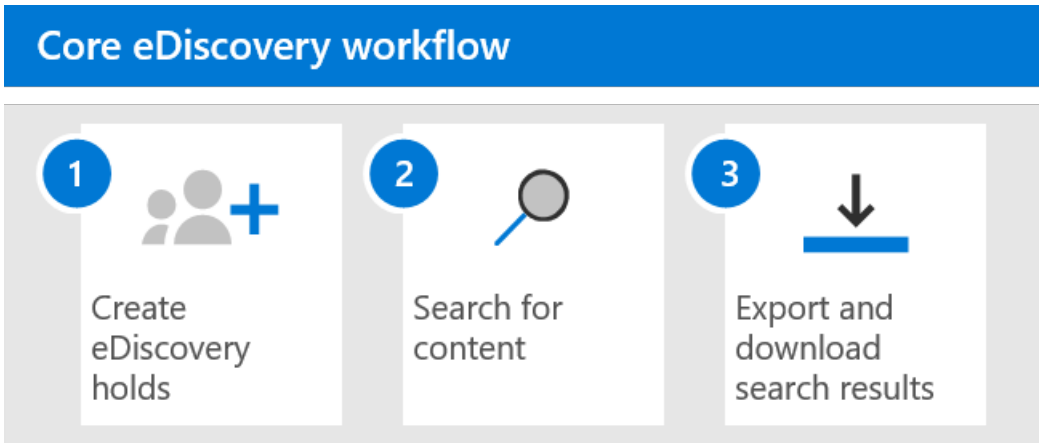

1. Create an eDiscovery hold. The first step after creating a case is placing a hold (also called an *eDiscovery hold*) on the content locations of the people of interest in your investigation. Content locations include Exchange mailboxes, SharePoint sites, OneDrive accounts, as well as the mailboxes and sites associated with Microsoft Teams and Office 365 Groups. While this step is optional, creating an eDiscovery hold preserves content that may be relevant to the case during the investigation. When you create an eDiscovery hold you can preserve all content in specific content locations or you can create a query-based hold to preserve only the content that matches a hold query. In addition to preserving content, another good reason to create eDiscovery holds is to quickly search the content locations on hold (instead of having to select each location to search) when you create and run searches in the next step. After you complete your investigation, you can release any hold that you created.
2. Search for content. After you create eDiscovery holds, use the built-in search tool to search the content locations on hold. You can also search other content locations for data that may be relevant to the case.

You can create and run different searches that are associated with the case. You use keywords, properties, and conditions to build search queries that return search results with the data that's most likely relevant to the case. You can also:

- View search statistics that may help you refine a search query to narrow the results.
- Preview the search results to quickly verify whether the relevant data is being found.
- Revise a query and rerun the search.

3. **Export and download search results**. After you search for and find data that's relevant to your investigation, you can export it out of Office 365 for review by people outside of the investigation team. Exporting data is a two-step process. The first step is to export the results of a search in the case out of Office 365. This is accomplished by copying the results of a search to a Microsoft-provided Azure Storage location. The next step is to use the eDiscovery Export tool to download the content to a local computer. In addition to the exported data files, the contains of the export package also contains an export report, a summary report, and an error report.

Create an eDiscovery hold

2/18/2021 • 18 minutes to read • Edit Online

You can use a Core eDiscovery case to create holds to preserve content that might be relevant to the case. You can place a hold on the Exchange mailboxes and OneDrive for Business accounts of people you're investigating in the case. You can also place a hold on the mailboxes and sites that are associated with Microsoft Teams, Office 365 Groups, and Yammer Groups. When you place content locations on hold, content is preserved until you remove the hold from the content location or until you delete the hold.

After you create an eDiscovery hold, it may take up to 24 hours for the hold to take effect.

When you create a hold, you have the following options to scope the content that is preserved in the specified content locations:

- You create an infinite hold where all content in the specified locations is placed on hold. Alternatively, you can create a query-based hold where only the content in the specified locations that matches a search query is placed on hold.
- You can specify a date range to preserve only the content that was sent, received, or created within that date range. Alternatively, you can hold all content in specified locations regardless of when it was sent, received, or created.

How to create an eDiscovery hold

To create an eDiscovery hold that's associated with a Core eDiscovery case:

1. Go to https://compliance.microsoft.com and sign in using the credentials for user account that has been assigned the appropriate eDiscovery permissions.
2. In the left navigation pane of the Microsoft 365 compliance center, click **Show all**, and then click **eDiscovery** > **Core**.
3. On the **Core eDiscovery** page, select the case that you want to create the hold in, and then click **Open case**.
4. On the **Home** page for the case, click the **Holds** tab.
5. On the **Holds** page, click **Create**.
6. On the **Name your hold** wizard page, give the hold a name and add an optional description, and then click **Next**. The name of the hold must be unique in your organization.
7. On the **Content locations** page, choose the content locations that you want to place on hold. You can place mailboxes, sites, and public folders on hold.

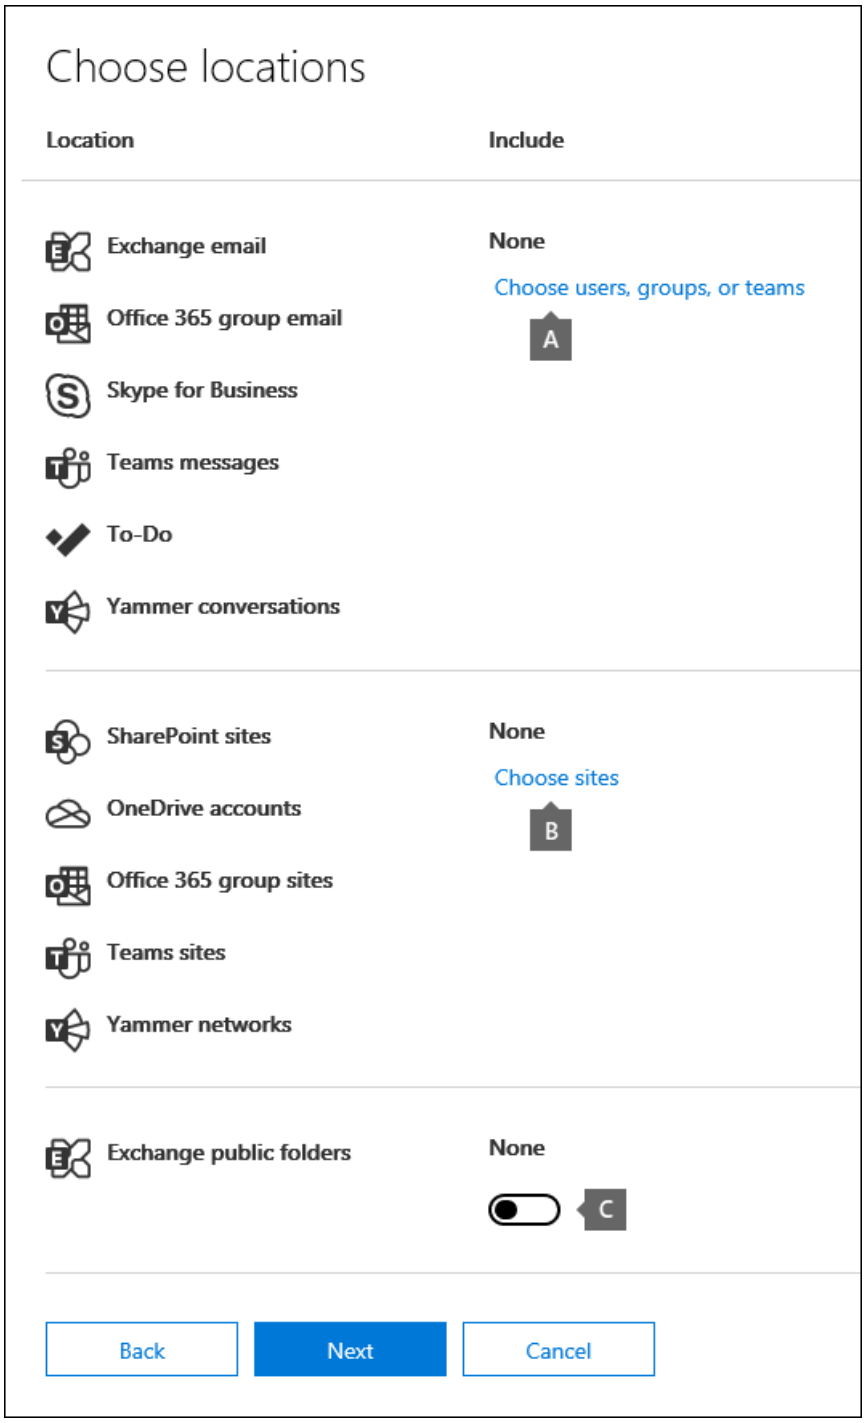

a. **Mailbox locations** - Click **Choose users, groups, or teams** and then click **Choose users, groups, or teams** again to specify the mailboxes to place on hold. Use the search box to find user mailboxes and distribution groups (to place a hold on the mailboxes of group members) to place on hold. You can also place a hold on the associated mailbox for a Microsoft Team, Office 365 Group, or Yammer Group. Select the user, group, team check box, click **Choose**, and then click **Done**.

b. **Site locations** - Click **Choose sites** and then click **Choose sites** again to specify SharePoint and OneDrive accounts to place on hold. Type the URL for each site that you want to place on hold. You can also add the URL for the SharePoint site for a Microsoft Team, Office 365 Group or a Yammer Group. Click **Choose**, and then click **Done**.

c. **Exchange public folders.** Move the toggle switch to the **All** position to put all public folders in your Exchange Online organization on hold. You can't choose specific public folders to put on hold. Leave the toggle switch set to **None** if you don't want to put a hold on public folders.

> **NOTE**
> You must add at least one content location to the hold. Otherwise, the eDiscovery hold statics will show that no items are on hold.

8. When you're done adding content locations to the hold, click **Next**.
9. To create a query-based hold with conditions, complete the following. Otherwise, to preserve all content in the specified content locations, click **Next**.

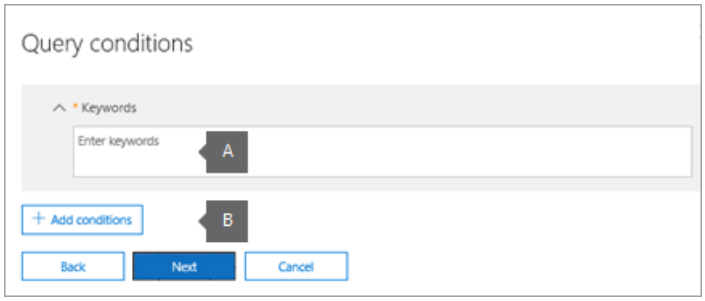

   a. In the box under **Keywords**, type a search query so that only the content that meets the search criteria is preserved. You can specify keywords, email message properties, or document properties, such as file names. You can also use more complex queries that use a Boolean operator, such as **AND**, **OR**, or **NOT**.

   b. Click **Add conditions** to add one or more conditions to narrow the search query for the hold. Each condition adds a clause to the KQL search query that is created and run when you create the hold. For example, you can specify a date range so that email or site documents that were created within the date ranged are placed on hold. A condition is logically connected to the keyword query (specified in the **Keywords** box) by the **AND** operator. That means that items have to satisfy both the keyword query and the condition to be preserved.

   For more information about creating a search query and using conditions, see Keyword queries and search conditions for Content Search.

10. After configuring a query-based hold, click **Next**.
11. Review your settings (and edit them if necessary), and then click **Create this hold**.

Query-based holds placed on site documents

Keep the following things in mind when you place a query-based eDiscovery hold on documents located in SharePoint sites:

- A query-based hold initially preserves all documents in a site for a short period of time after they are deleted. That means when a document is deleted, it will be moved to the Preservation Hold library even if it doesn't match the criteria of the query-based hold. However, deleted documents that don't match a query-based hold will be removed by a timer job that processes the Preservation Hold library. The timer job runs periodically and compares all documents in the Preservation Hold library to your query-based eDiscovery holds (and other types of holds and retention policies). The timer job deletes the documents that don't match a query-based hold and preserves the documents that do.
- Query-based holds should not be used to perform targeted preservation, like preserving documents in a specific folder or site or by using other location-based hold criteria. Doing so may have unintended results. We recommend using non-location based hold criteria such as keywords, date ranges, or other document properties to preserve site documents.

eDiscovery hold statistics

After you create an eDiscovery hold, information about the new hold is displayed on the flyout page for the selected hold. This information includes the number of mailboxes and sites on hold and statistics about the content that was placed on hold, such as the total number and size of items placed on hold and the last time the hold statistics were calculated. These hold statistics help you identify the amount of content related to the case is being preserved.

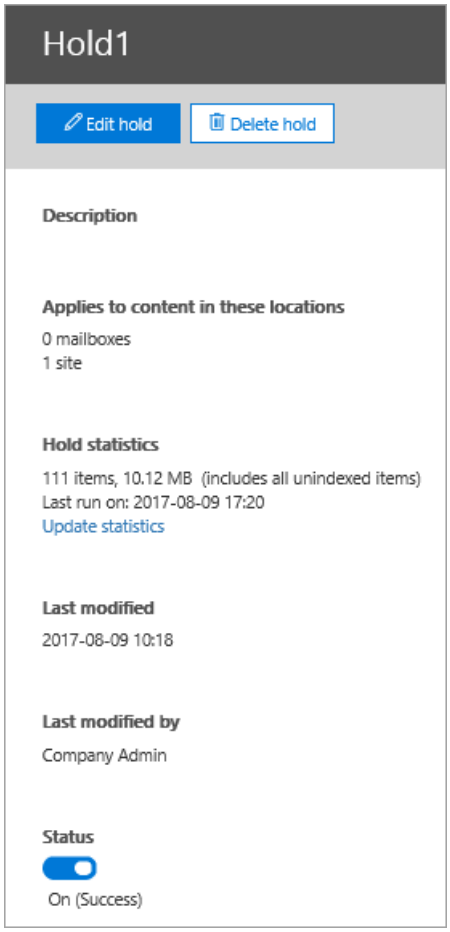

Keep the following things in mind about eDiscovery hold statistics:

- The total number of items on hold indicates the number of items from all content sources that are placed on hold. If you've created a query-based hold, this statistic indicates the number of items that match the query.
- The number of items on hold also includes unindexed items found in the content locations. If you create a query-based hold, all unindexed items in the content locations are placed on hold. This includes unindexed items that don't match the search criteria of a query-based hold and unindexed items that might fall outside of a date range condition. This is different than what happens when you run a search, in which unindexed items that don't match the search query or are excluded by a date range condition aren't included in the search results. For more information about unindexed items, see Partially indexed items.
- You can get the latest hold statistics by clicking **Update statistics** to rerun a search estimate that calculates the current number of items on hold.
- It's normal for the number of items on hold to increase over time because users whose mailbox or site is on hold are typically sending or receiving new email message and creating new documents in SharePoint and OneDrive.
- If an Exchange mailbox, SharePoint site, or OneDrive account is moved to a different region in a multi-geo environment, the statistics for that site won't be included in the hold statistics. But the content in those

locations will still be preserved. Also, if a mailbox or site is moved to a different region, the SMTP address or URL that's displayed in the hold will not automatically be updated. You'll have to edit the hold and update the URL or SMTP address so the content locations are once again included in the hold statistics

Search locations on eDiscovery hold

When you search for content in a Core eDiscovery case, you can quickly configure the search to only search the content locations that have been placed on a hold associated with the case.

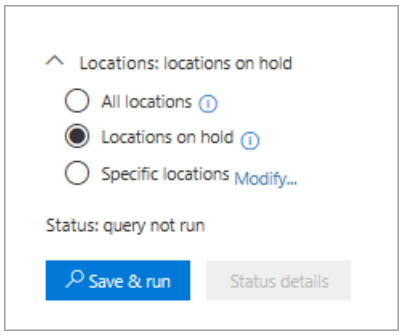

Select the **Locations on hold** option to search all the content locations that have been placed on hold. If the case contains multiple eDiscovery holds, the content locations from all holds will be searched when you select this option. Additionally, if a content location was placed on a query-based hold, only the items that match the hold query will be searched when you run the search. In other words, only the content that matches both the hold criteria and the search criteria is returned with the search results. For example, if a user was placed on query-based case hold that preserves items that were sent or created before a specific date, only those items would be searched. This is accomplished by connecting the case hold query and the search query by an **AND** operator.

Here are some other things to keep in mind when searching locations on eDiscovery hold:

- If a content location is part of multiple holds within the same case, the hold queries are combined by **OR** operators when you search that content location using the all case content option. Similarly, if a content location is part of two different holds, where one is query-based and the other is an infinite hold (where all content is placed on hold), then all content is search because of the infinite hold.
- If a search is configured it to search locations on hold and then you change an eDiscovery hold in the case (by adding or removing a location or changing a hold query), the search configuration is updated with those changes. However, you have to rerun the search after the hold is changed to update the search results.
- If multiple eDiscovery holds are placed on a single location in an eDiscovery case and you select to search locations on hold, the maximum number of keywords for that search query is 500. That's because the search combines all the query-based holds by using the **OR** operator. If there are more than 500 keywords in the combined hold queries and the search query, then all content in the mailbox is searched, not just that content that matches the query-based case holds.
- If an eDiscovery hold has a status of **Turning on**, you can still search the locations on hold while the hold is being turned on.

Preserve content in Microsoft Teams

Conversations that are part of a Microsoft Teams channel are stored in the mailbox that's associated with the Microsoft Team. Similarly, files that team members share in a channel are stored on the team's SharePoint site. Therefore, you have to place the Team mailbox and SharePoint site on eDiscovery hold to preserve conversations and files in a channel.

Alternatively, conversations that are part of the Chat list in Teams (called *1:1 chats* or *1:N group chats*) are stored

in the mailboxes of the users who participate in the chat. And files that users share in chat conversations are stored in the OneDrive account of the user who shares the file. Therefore, you have to add the individual user mailboxes and OneDrive accounts to an eDiscovery hold to preserve conversations and files in the chat list. It's a good idea to place a hold on the mailboxes of members of a Microsoft Team in addition to placing the team mailbox and site on hold.

> **IMPORTANT**
> In a cloud-based organization, users who participate in conversations that are part of the chat list in Teams must have an Exchange Online mailbox in order to retain chat conversations when the mailbox is placed on an eDiscovery hold. That's because conversations that are part of the chat list are stored in the cloud-based mailboxes of the chat participants. If a chat participant doesn't have an Exchange Online mailbox, you won't be able to preserve those chat conversations. For example, in an Exchange hybrid deployment, users with an on-premises mailbox may be able to participate in conversations that are part of the chat list in Teams. But in this case, content from these conversation can't be preserved because these users don't have a cloud-based mailboxes that can be placed on hold.

For more information about preserving Teams content, see Place a Microsoft Teams user or team on legal hold.

Preserve card content

Similarly, card content generated by apps in Teams channels, 1:1 chats, and 1:N group chats is stored in mailboxes and is preserved when a mailbox is placed on an eDiscovery hold. A *card* is a UI container for short pieces of content. Cards can have multiple properties and attachments, and can include buttons that trigger card actions. For more information, see Cards. Like other Teams content, where card content is stored is based on where the card was used. Content for cards used in a Teams channel is stored in the Teams group mailbox. Card content for 1:1 and 1xN chats are stored in the mailboxes of the chat participants.

Preserve meeting and call information

Summary information for meetings and calls in a Teams channel is also stored in the mailboxes of users who dialed into the meeting or call. This content is also preserved when an eDiscovery hold is placed on user mailboxes.

Preserve content in private channels

Starting in February 2020, we also turned on ability to preserve content in private channels. Because private channel chats are stored in the mailboxes of the chat participants, placing a user mailbox on eDiscovery hold will preserve private channel chats. Also, if a user mailbox was placed on an eDiscovery hold prior to February 2020, the hold will now automatically apply to private channel messages stored in that mailbox. Preserving files shared in private channels is also supported.

Preserve wiki content

Every Team or team channel also contains a Wiki for note taking and collaboration. The Wiki content is automatically saved to a file with a .mht format. This file is stored in the Teams Wiki Data document library on the team's SharePoint site. You can preserve the wiki content by adding the team's SharePoint site to an eDiscovery hold.

> **NOTE**
> The capability to preserve Wiki content for a Team or team channel (when you place the team's SharePoint site on hold) was released on June 22, 2017. If a team site is on hold, the Wiki content will be retained starting on that date. However, if a team site is on hold and the Wiki content was deleted before June 22, 2017, the Wiki content was not preserved.

Office 365 Groups

Teams is built on Office 365 Groups. Therefore, placing Office 365 Groups on eDiscovery hold is similar placing Teams content on hold.

Keep the following things in mind when placing both Teams and Office 365 Groups on an eDiscovery hold:

- As previously explained, to place content located in Teams and Office 365 Groups on hold, you have to specify the mailbox and SharePoint site that associated with a group or team.
- Run the **Get-UnifiedGroup** cmdlet in Exchange Online PowerShell to view properties for Teams and Office 365 Groups. This is a good way to get the URL for the site that's associated with a Team or Office 365 Group. For example, the following command displays selected properties for an Office 365 Group named Senior Leadership Team:

```
Get-UnifiedGroup "Senior Leadership Team" | FL DisplayName,Alias,PrimarySmtpAddress,SharePointSiteUrl

DisplayName            : Senior Leadership Team
Alias                  : seniorleadershipteam
PrimarySmtpAddress     : seniorleadershipteam@contoso.onmicrosoft.com
SharePointSiteUrl      : https://contoso.sharepoint.com/sites/seniorleadershipteam
```

> **NOTE**
> To run the **Get-UnifiedGroup** cmdlet, you have to be assigned the View-Only Recipients role in Exchange Online or be a member of a role group that's assigned the View-Only Recipients role.

- When a user's mailbox is searched, any Team or Office 365 Group that the user is a member of won't be searched. Similarly, when you place a Team or Office 365 Group on eDiscovery hold, only the group mailbox and group site are placed on hold. The mailboxes and OneDrive for Business sites of group members aren't placed on hold unless you explicitly add them to the eDiscovery hold. So if you have to place a Team or Office 365 Group on hold for a legal reason, consider adding the mailboxes and OneDrive accounts of team or group members on the same hold.
- To get a list of the members of a Team or Office 365 Group, you can view the properties on the **Groups** page in the Microsoft 365 admin center. Alternatively, you can run the following command in Exchange Online PowerShell:

```
Get-UnifiedGroupLinks <group or team name> -LinkType Members | FL DisplayName,PrimarySmtpAddress
```

> **NOTE**
> To run the **Get-UnifiedGroupLinks** cmdlet, you have to be assigned the View-Only Recipients role in Exchange Online or be a member of a role group that's assigned the View-Only Recipients role.

Preserve content in OneDrive accounts

To collect a list of the URLs for the OneDrive for Business sites in your organization so you can add them to a hold or search associated with an eDiscovery case, see Create a list of all OneDrive locations in your organization. The script in this article creates a text file that contains a list of all OneDrive sites in your organization. To run this script, you have to install and use the SharePoint Online Management Shell. Be sure to append the URL for your organization's MySite domain to each OneDrive site that you want to search. This is the domain that contains all your OneDrive; for example, `https://contoso-my.sharepoint.com`. Here's an example of a URL for a user's OneDrive site: `https://contoso-my.sharepoint.com/personal/sarad_contoso_onmicrosoft.com`.

IMPORTANT

The URL for a user's OneDrive account includes their user principal name (UPN) (for example, `https://alpinehouse-my.sharepoint.com/personal/sarad_alpinehouse_onmicrosoft_com`). In the rare case that a person's UPN is changed, their OneDrive URL will also change to incorporate the new UPN. If a user's OneDrive account is part of an eDiscovery hold, old and their UPN is changed, you need to update the hold and you'll have to update the hold and add the user's new OneDrive URL and remove the old one. For more information, see How UPN changes affect the OneDrive URL.

Removing content locations from an eDiscovery hold

After a mailbox, SharePoint site, or OneDrive account is removed from an eDiscovery hold, a *delay hold* is applied. This means that the actual removal of the hold is delayed for 30 days to prevent data from being permanently deleted (purged) from a content location. This gives admins an opportunity to search for or recover content that will be purged after an eDiscovery hold is removed. The details of how the delay hold works for mailboxes and sites are different.

- **Mailboxes:** A delay hold is placed on a mailbox the next time the Managed Folder Assistant processes the mailbox and detects that an eDiscovery hold was removed. Specifically, a delay hold is applied to a mailbox when the Managed Folder Assistant sets one of the following mailbox properties to **True**:
  - **DelayHoldApplied:** This property applies to email-related content (generated by people using Outlook and Outlook on the web) that's stored in a user's mailbox.
  - **DelayReleaseHoldApplied:** This property applies to cloud-based content (generated by non-Outlook apps such as Microsoft Teams, Microsoft Forms, and Microsoft Yammer) that's stored in a user's mailbox. Cloud data generated by a Microsoft app is typically stored in a hidden folder in a user's mailbox.

  When a delay hold is placed on the mailbox (when either of the previous properties is set to **True**), the mailbox is still considered to be on hold for an unlimited hold duration, as if the mailbox was on Litigation Hold. After 30 days, the delay hold expires, and Microsoft 365 will automatically attempt to remove the delay hold (by setting the DelayHoldApplied or DelayReleaseHoldApplied property to **False**) so that the hold is removed. After either of these properties are set to **False**, the corresponding items that are marked for removal are purged the next time the mailbox is processed by the Managed Folder Assistant.

  For more information, see Managing mailboxes on delay hold.

- **SharePoint and OneDrive sites:** Any SharePoint or OneDrive content that's being retained in the Preservation Hold library isn't deleted during the 30-day delay hold period after a site is removed from an eDiscovery hold. This is similar to what happens when a site is released from a retention policy. Additionally, you can't manually delete this content in the Preservation Hold library during the 30-day delay hold period.

  For more information, see Releasing a policy for retention.

A delay hold is also applied to content locations on hold when you close a Core eDiscovery case because holds are turned off when a case is closed. For more information about closing a case, see Close, reopen, and delete a Core eDiscovery case.

eDiscovery hold limits

The following table lists the limits for eDiscovery cases and case holds.

DESCRIPTION OF LIMIT	LIMIT
Maximum number of cases for an organization	No limit
Maximum number of eDiscovery holds for an organization	10,000
Maximum number of mailboxes in a single eDiscovery hold	1,000
Maximum number of SharePoint and OneDrive for Business sites in a single eDiscovery hold	100
Maximum number of cases displayed on the eDiscovery home page, and the maximum number of items displayed on the Holds, Searches, and Export tabs within a case. [1]	1,000

NOTE

[1] To view a list of more than 1,000 cases, holds, searches, or exports, you can use the corresponding Office 365 Security & Compliance PowerShell cmdlet:

- Get-ComplianceCase
- Get-CaseHoldPolicy
- Get-ComplianceSearch
- Get-ComplianceSearchAction

Search for content in a Core eDiscovery case

5/13/2020 • 5 minutes to read • Edit Online

After a Core eDiscovery case is created and people of interest in the case are placed on hold, you can create and run one or more searches for content relevant to the case. Searches associated with a Core eDiscovery case aren't listed on the **Content search** page in the Microsoft 365 compliance center. These searches are listed on the **Searches** page of the Core eDiscover case the searches are associated with. This also means that searches associated with a case can only be accessed by case members.

To create a Core eDiscovery search:

1. Go to https://compliance.microsoft.com and sign in using the credentials for user account that has been assigned the appropriate eDiscovery permissions.
2. In the left navigation pane of the Microsoft 365 compliance center, click **Show all**, and then click **eDiscovery > Core**.
3. On the **Core eDiscovery** page, select the case that you want to create an associated search, and then click **Open case**.
4. On the **Home** page for the case, click the **Searches** tab.
5. On the **Search** page, click **New search**.
6. On the **New search** page, you can add keywords and conditions to create the search query.

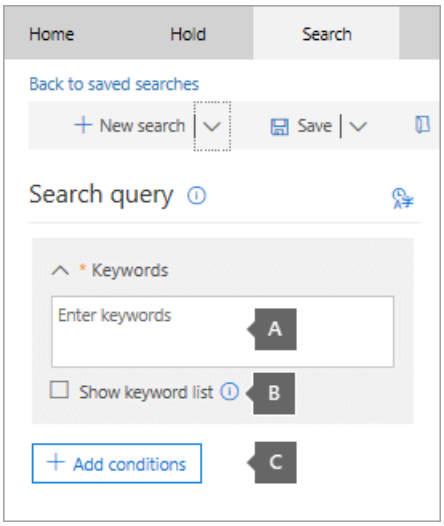

a. You can specify keywords, message properties, such as sent and received dates, or document properties, such as file names or the date that a document was last changed. You can use more complex queries that use a Boolean operator, such as **AND**, **OR**, **NOT**, or **NEAR**. You can also search for sensitive information (such as social security numbers) in documents, or search for documents that have been shared externally. If you leave the keyword box empty, all content located in the specified content locations will be included in the search results.

b. You can click the **Show keyword list** check box and the type a keyword in each row. If you do this, the keywords on each row are connected by the **OR** operator in the search query that's created. You can enter a maximum of 20 keywords to the list.

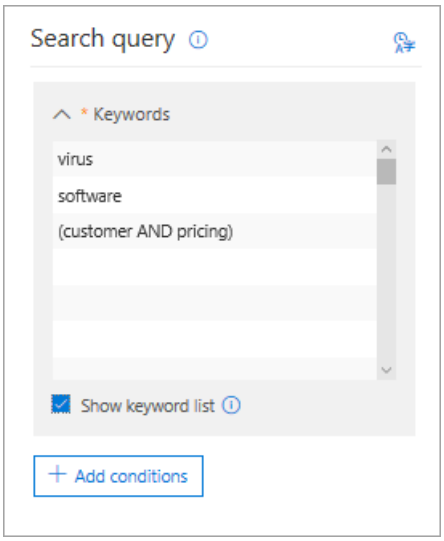

Why use the keyword list? You can get statistics that show how many items match each keyword. This can help you quickly identify which keywords are the most (and least) effective. You can also use a keyword phrase (surrounded by parentheses) in a row. For more information about search statistics, see View keyword statistics for Content Search results.

For more information about using the keywords list, see Building a search query.

c. You can click **Conditions** and add conditions to a search query to narrow a search and return a more refined set of results. Each condition adds a clause to the KQL search query that is created and run when you start the search. A condition is logically connected to the keyword query (specified in the keyword box) by the **AND** operator. That means that items have to satisfy both the keyword query and each condition to be included in the results. This is how conditions help to narrow your results.

For more information about creating a search query and using conditions, see Keyword queries for Content Search.

7. Under **Locations: locations on hold**, choose the content locations that you want to search. You can search mailboxes, sites, and public folders in the same search.

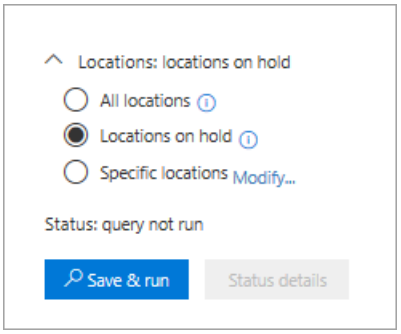

- **All locations**. Select this option to search all content locations in your organization. When you select this option, you can choose to search all Exchange mailboxes (which includes the mailboxes for all Microsoft Teams, Yammer Groups, and Office 365 Groups), all SharePoint and OneDrive for Business sites (which includes the sites for all Microsoft Teams, Yammer Groups, and Office 365 Groups), and all public folders.

- **All locations on hold**. Select this option to search all the content locations that have been placed on eDiscovery hold in the case. If the case contains multiple holds, the content locations from all holds will be searched. Additionally, if a content location was placed on a query-based hold, only the items that are on hold will be searched when you run the content search that you're creating in this step. For example, if a user was placed on query-based case hold that preserves items that were sent or created before a specific date, only those items would be searched. This is accomplished by connecting the case hold query and the content search query by an **AND**

operator. For more information, see Search locations on eDiscovery hold.

- **Specific locations**. Select this option to select the mailboxes and sites that you want to search. When you select this option and click **Modify**, a list of locations appears. You can choose to search any or all users, groups, teams, or site locations. You can also search the public folders in your organization.

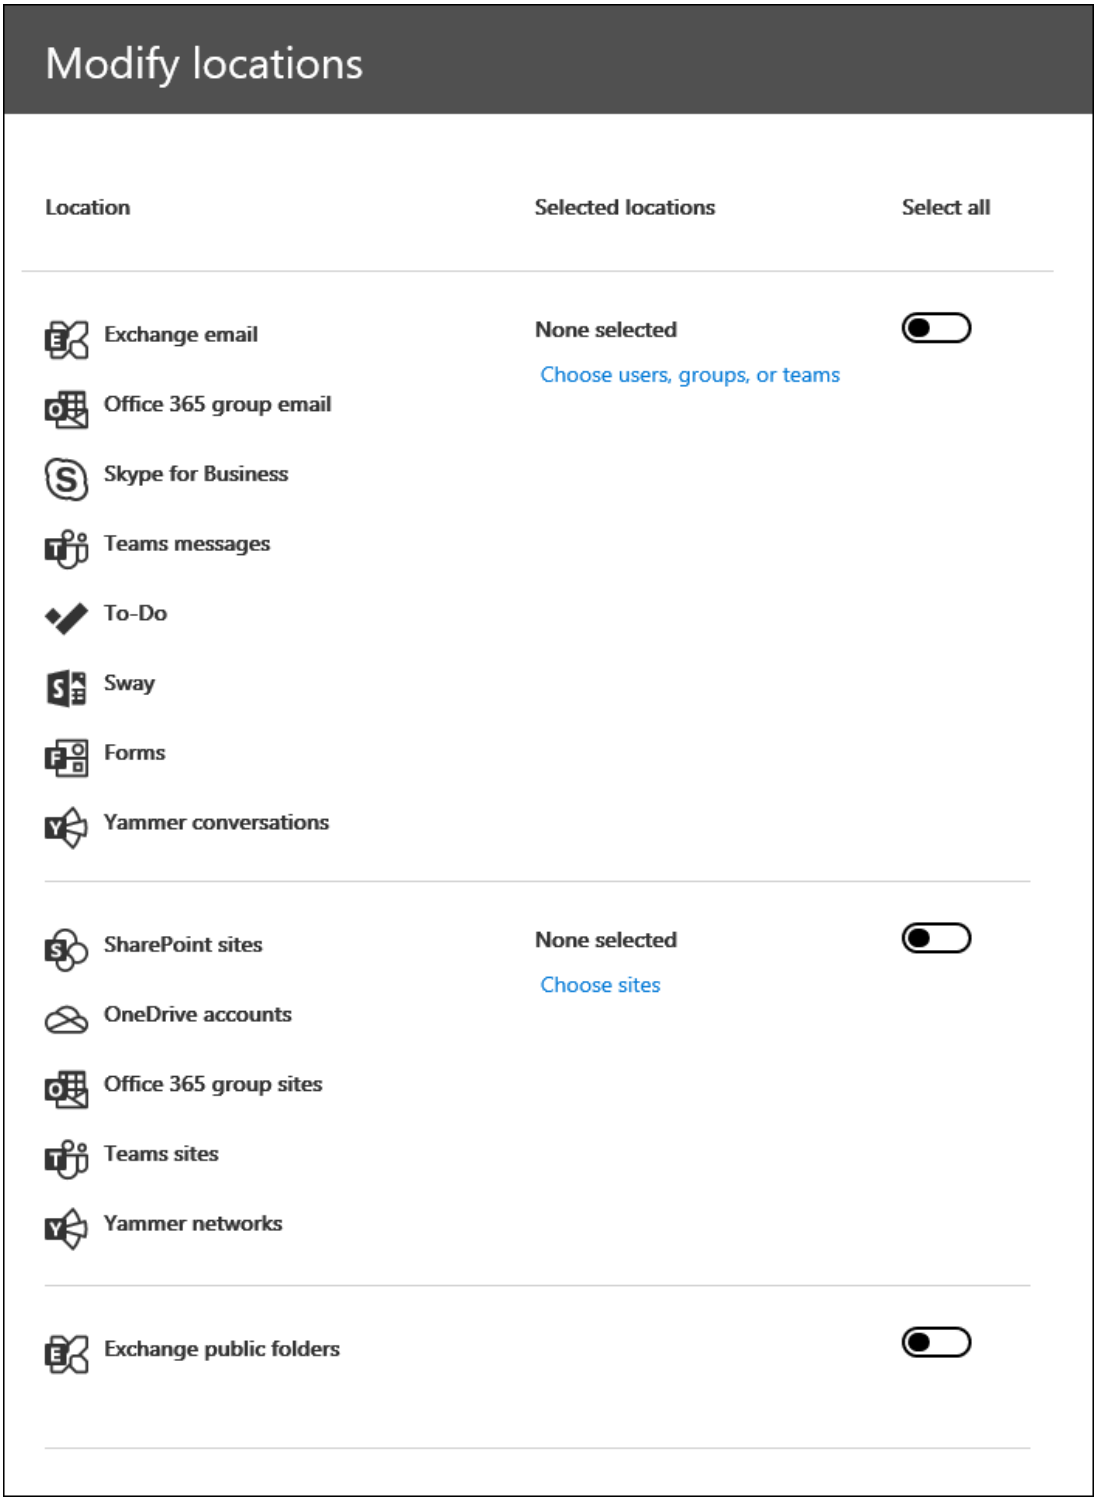

If you select this option and search any content location that's on hold, any query from a query-based case hold won't be applied to the search query. In other words, all content is searched, not just the content that's preserved by a query-based case hold.

8. After you select the content locations to search, click **Done** and then click **Save**.
9. On the **New search** page, click **Save & run** and then type a name for the search. Searches associated with a Core eDiscovery case must have names that are unique within your Office 365 organization.

10. Click **Save** to save the search settings and start the search.

    After the search is completed, you can preview the search results. If necessary, click **Refresh** on the **Searches** page to display the search you created in the list.

11. Click the search to display the flyout page, which contains statistics about the search and to perform other tasks such as viewing search statistics and exporting the search results.

More information about searching content locations

- When you click **Choose users, groups, or teams** to specify mailboxes to search, the mailbox picker that's displayed is empty. This is by design to enhance performance. To add recipients to this list, click **Choose users, groups, or teams**, type a name (a minimum of 3 characters) in the search box, select the check box next to the name, and then click **Choose**.
- You can add inactive mailboxes, Microsoft Teams, Yammer Groups, Office 365 Groups, and distribution groups to the list of mailboxes to search. Dynamic distribution groups aren't supported. If you add Microsoft Teams, Yammer Groups, or Office 365 Groups, the group or team mailbox is searched; the mailboxes of the group members aren't searched.
- To add sites click **Choose sites**, click **Choose sites** again, and then type the URL for each site that you want to search. You can also add the URL for the SharePoint site for a Microsoft Team, a Yammer Group, or an Office 365 Group.

Export content from a Core eDiscovery case

2/18/2021 • 5 minutes to read • Edit Online

After a search is successfully run, you can export the search results. When you export search results, mailbox items are downloaded in PST files or as individual messages. When you export content from SharePoint and OneDrive for Business sites, copies of native Office documents and other documents are exported. A Results.csv file that contains information about every item that's exported and a manifest file (in XML format) that contains information about every search result is also exported.

You can export the results of a single search associated with a case or you can export the results of multiple searches associated with a case.

Export the results of a single search

1. Go to https://compliance.microsoft.com and sign in using the credentials for user account that has been assigned the appropriate eDiscovery permissions.
2. In the left navigation pane of the Microsoft 365 compliance center, click **Show all**, and then click **eDiscovery > Core**.
3. On the **Core eDiscovery** page, select the case that you want to export search results from, and then click **Open case**.
4. On the **Home** page for the case, click the **Searches** tab.
5. In the list of searches for the case, click the search that you want to export search results from, and then click **Export results** on the flyout.

   The **Export results** page is displayed.

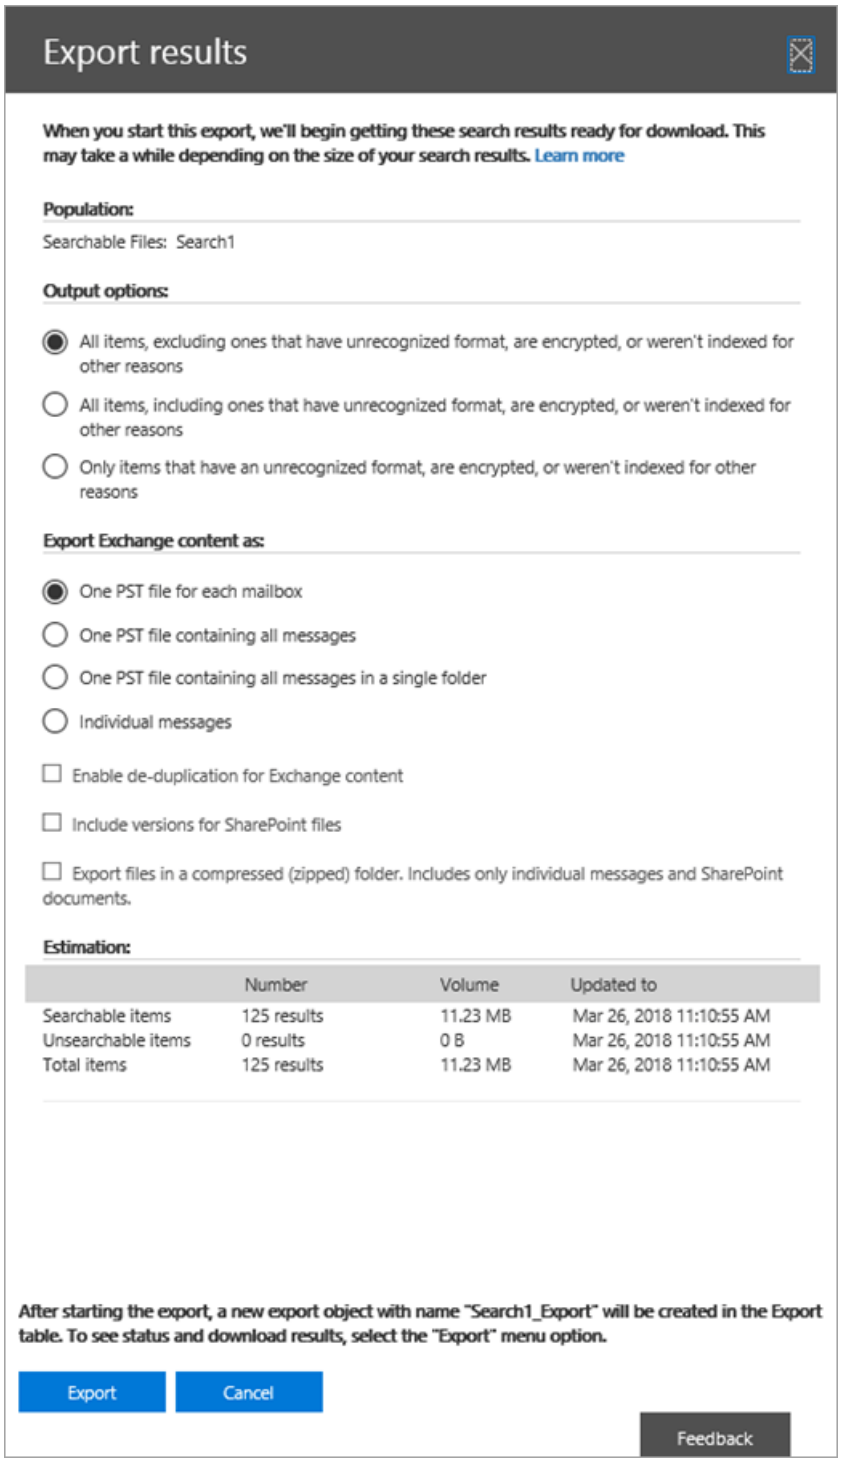

The workflow to export the results of a search associated with a Core eDiscovery case is that same as exporting the search results for a search on the **Content search** page. For step-by-step instructions, see Export content search results.

> **NOTE**
>
> When you export search results, you have the option to enable de-duplication so that only one copy of an email message is exported even though multiple instances of the same message might have been found in the mailboxes that were searched. For more information about de-duplication and how duplicate items are identified, see De-duplication in eDiscovery search results.

After you start the export, the search results are prepared for downloading, which means they are uploaded to a Microsoft-provided Azure Storage location in the Microsoft cloud.

6. Click the **Export** tab to display the list of export jobs for the case.

You may have to click **Refresh** to update the list of export jobs so that it shows the export job you created. Export jobs have the same name as the corresponding search with **_Export** appended to the search name.

7. Click the export job you created to display status information on the flyout page. This information includes the percentage of items that have been transferred to the Azure Storage location.

8. After all items have been transferred, click **Download results** to download the search results to your local computer. For more information downloading search results, see Step 2 in Export content search results

Export the results of multiple searches

As an alternative to exporting the results of a single search associated with a case, you can export the results of multiple searches from the same case in a single export job. Exporting the results of multiple searches is faster and easier than exporting the results one search at a time.

> **NOTE**
> You can't export the results of multiple searches if one of those searches was configured to search locations on hold.

1. Go to https://compliance.microsoft.com and sign in using the credentials for user account that has been assigned the appropriate eDiscovery permissions.

2. In the left navigation pane of the Microsoft 365 compliance center, click **Show all**, and then click **eDiscovery** > **Core**.

3. On the **Core eDiscovery** page, select the case that you want to export search results from, and then click **Open case**.

4. On the **Home** page for the case, click the **Searches** tab.

5. In the list of searches for the case, select the checkbox next to two or more searches you want to export search results from.

   The **Bulk actions** flyout page appears.

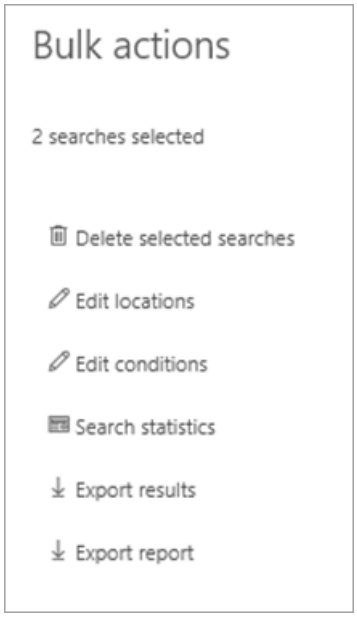

6. Click **Export results**.

   The **Export results** page is displayed.

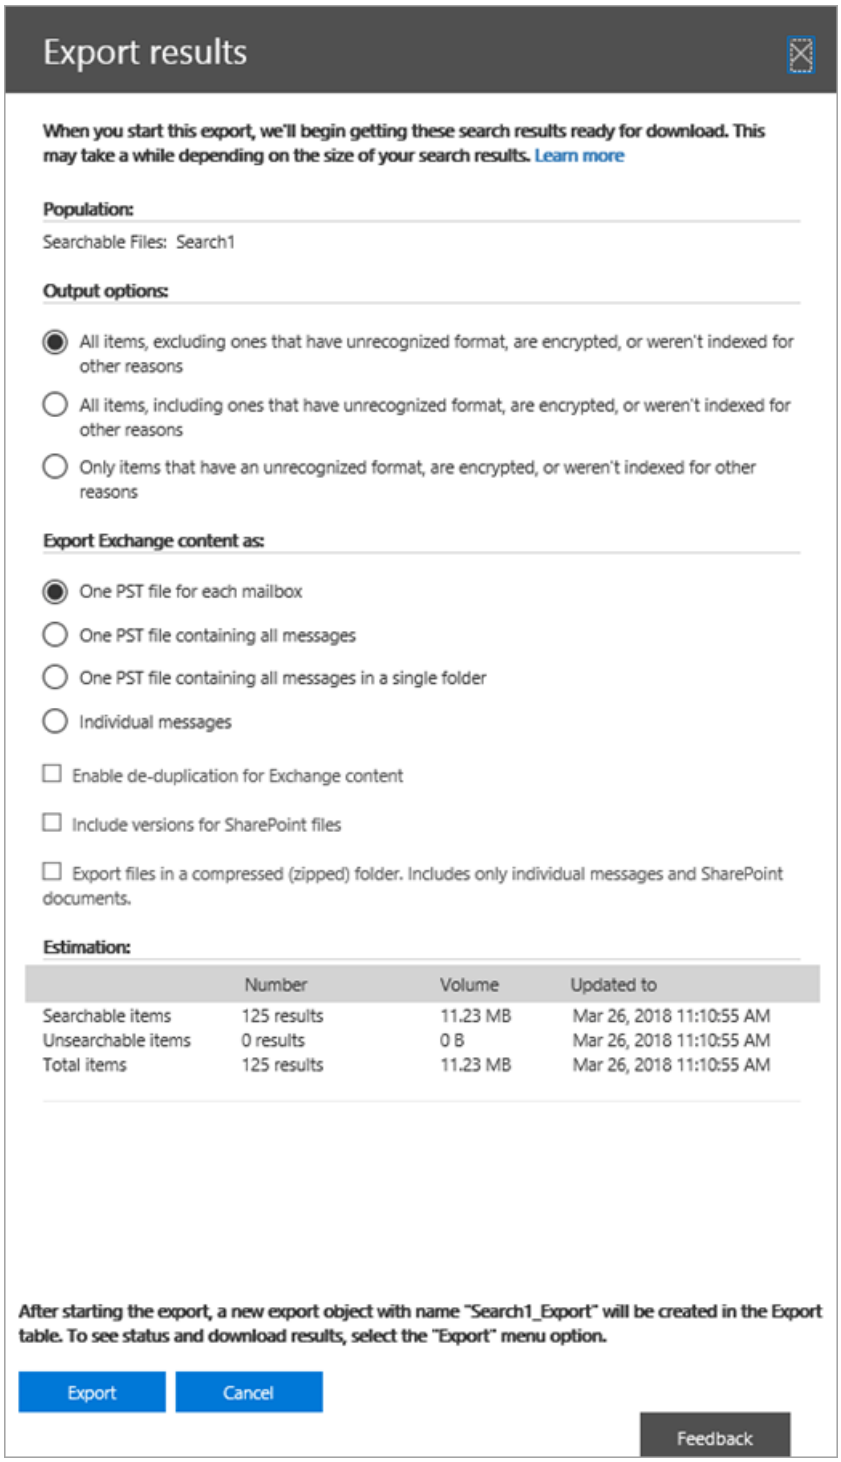

At this point, the workflow to export the results of multiple searches associated with a Core eDiscovery case is that same as exporting the search results for a single search. See step 5 in the previous section.

More information about exporting the results of multiple searches

- When you export the results of multiple searches, the search queries from all the searches are combined by using **OR** operators, and then the combined search is started. The estimated results of the combined search are displayed in the flyout page of the selected export job. The search results are then copied to the Azure Storage location in the Microsoft cloud. The status of the copy job is also displayed on the flyout page. As previously stated, after all the search results have been copied, you can download them to a local computer.

- The maximum number of keywords from queries for all searches that you want to export is 500. This is the same limit for a single search. That's because the export job combines all the search queries by using the **OR** operator. If you exceed this limit, an error will be returned. In this case, you have to export the results from fewer searches or simplify the search queries of the original searches that you want to export.

- The search results that are exported are organized by the content location the item was found in. That means a content location in the export results may have items returned by different searches. For example, if you choose to export email messages in one PST file for each mailbox, the PST file might have results from multiple searches.
- If the same email item or document from the same content location is returned by more than one of the searches that you export, only one copy of the item will be exported.
- You can't edit an export for multiple searches after you create it. For example, you can't add or remove searches from the export job. You have to create an export job to change which search results are exported. After an export job is created, you only can download the results to a computer, restart the export, or delete the export job.
- If you restart the export, any changes to the queries of the searches that make up the export job won't affect the search results that are retrieved. When you restart an export, the same combined search query job that was run when the export job was created will be run again.
- Also, if you restart an export, the search results that are copied to the Azure Storage location overwrites the previous results. The previous results that were copied won't be available to be downloaded.

Close, reopen, and delete a Core eDiscovery case

11/2/2020 • 4 minutes to read • Edit Online

This article describes how to close, reopen, and delete Core eDiscovery cases in Microsoft 365.

Close a case

When the legal case or investigation supported by a Core eDiscovery case is completed, you can close the case. Here's what happens when you close a case:

- If the case contains any content locations on eDiscovery hold, those holds will be turned off. After the hold is turned off, a 30-day grace period (called a *delay hold*) is applied to content locations that were on hold. This helps prevent content from being immediately deleted and provides admins the opportunity to search for and restore content before it may be permanently deleted after the delay hold period expires. For more information, see Removing content locations from an eDiscovery hold.
- Closing a case only turns off the holds that are associated with that case. If other holds are placed on a content location (such as a Litigation Hold, a retention policy, or a hold from a different Core eDiscovery case) those holds will still be maintained.
- The case is still listed on the Core eDiscovery page in the Microsoft 365 compliance center. The details, holds, searches, and members of a closed case are retained.
- You can edit a case after it's closed. For example, you can add or removing members, create searches, and export search results. The primary difference between active and closed cases is that eDiscovery holds are turned off when a case is closed.

To close a case:

1. In the Microsoft 365 compliance center, click **eDiscovery** > **Core** to display the list of Core eDiscovery cases in your organization.
2. Click the name of the case that you want to close.

   The **Manage this case** flyout page is displayed.
3. Under **Manage case status**, click **Close case**.

   A warning is displayed saying that the holds associated with the case will be turned off.
4. Click **Yes** to close the case.

   The status on the **Manage this case** flyout page is changed from **Active** to **Closing**.
5. Close the **Manage this case** page.
6. On the **Core eDiscovery** page, click **Refresh** to update the status of the closed case. It might take up to 60 minutes for the closing process to complete.

   When the process is complete, the status of the case is changed to **Closed** on the **Core eDiscovery** page. Click the name of the case again to display the **Manage this case** flyout page, which contains information about when the case was closed and who closed it.

Reopen a closed case

When you reopen a case, any eDiscovery holds that were in place when the case was closed won't be automatically reinstated. After the case is reopened, you'll have to go to the **Holds** page and turn on the previous holds. To turn on a hold, select it to display the flyout page, and then set the **Status** toggle to **On**.

1. In the Microsoft 365 compliance center, click **eDiscovery** > **Core** to display the list of Core eDiscovery cases in your organization.
2. Click the name of the case that you want to reopen.

   The **Manage this case** flyout page is displayed.
3. Under **Manage case status**, click **Reopen case**.

   A warning is displayed saying that the holds that were associated with the case when it was closed won't be turned on automatically.
4. Click **Yes** to reopen the case.

   The status on the **Manage this case** flyout page is changed from **Closed** to **Active**.
5. Close the **Manage this case** page.
6. On the **Core eDiscovery** page, click **Refresh** to update the status of the reopened case. It might take up to 60 minutes for the reopening process to complete.

   When the process is complete, the status of the case is changed to **Active** on the **Core eDiscovery** page.

Delete a case

You can also delete active and closed Core eDiscovery cases. When you delete a case, all searches and exports in the case are deleted, and the case is removed from the list of cases on the **Core eDiscovery** page in the Microsoft 365 compliance center. You can't reopen a deleted case.

Before you can delete a case (whether it's active or closed), you must first delete *all* eDiscovery holds associated with the case. That includes deleting holds with a status of **Off**.

To delete an eDiscovery hold:

1. Go the **Holds** tab in the case that you want to delete.
2. Click the hold that you want to delete.
3. On the flyout page, click **Delete hold**.

To delete a case:

1. In the Microsoft 365 compliance center, click **eDiscovery** > **Core** to display the list of Core eDiscovery cases in your organization.
2. Click the name of the case that you want to delete.
3. Under **Manage case status** on the flyout page, click **Delete case**.

If the case you're trying to delete still contains eDiscovery holds, you'll receive an error message. You'll have to delete all holds associated with the case and then try again to delete the case.

Assign eDiscovery permissions in the Security & Compliance Center

2/18/2021 • 10 minutes to read • Edit Online

If you want people to use any of the eDiscovery-related tools in the Security & Compliance Center in Office 365 or the Microsoft 365 compliance center, you have to assign them the appropriate permissions. The easiest way to do this is to add the person the appropriate role group on the **Permissions** page in the Security & Compliance Center. This topic describes the permissions required to perform eDiscovery- and Content Search-related tasks using the Security & Compliance Center.

The primary eDiscovery-related role group in Security & Compliance Center is called **eDiscovery Manager**. There are two subgroups within this role group.

- **eDiscovery Managers** - An eDiscovery Manager can use the Content Search tool in the Security & Compliance Center to search content locations in the organization, and perform various search-related actions such as preview and export search results. Members can also create and manage cases in Core eDiscovery and Advanced eDiscovery, add and remove members to a case, create case holds, run searches associated with a case, and access case data. eDiscovery Managers can only access and manage the cases they create. They can't access or manage cases created by other eDiscovery Managers.
- **eDiscovery Administrators** - An eDiscovery Administrator is a member of the eDiscovery Manager role group, and can perform the same content search and case management-related tasks that an eDiscovery Manager can perform. Additionally, an eDiscovery Administrator can:
  - Access all cases that are listed on the **eDiscovery** and **Advanced eDiscovery** pages in the Security & Compliance Center.
  - Access case data in Advanced eDiscovery for any case in the organization.
  - Manage any eDiscovery case after they add themselves as a member of the case.

  For reasons why you might want eDiscovery Administrators in your organization, see More information.

> **NOTE**
> To analyze a user's data using Advanced eDiscovery, the user (the custodian of the data) must be assigned an Office 365 E5 or Microsoft 365 E5 license. Alternatively, users with an Office 365 E1 or a Office 365 or Microsoft 365 E3 license can be assigned an Microsoft 365 E5 Compliance or Microsoft 365 eDiscovery and Audit add-on license. Administrators, compliance officers, or legal personnel who are assigned to cases as members and use Advanced eDiscovery to collect, view, and analyze data don't need an E5 license. For more information about Advanced eDiscovery licensing, see Get started with Advanced eDiscovery.

Confirm your roles

- You have to be a member of the Organization Management role group or be assigned the Role Management role to assign eDiscovery permissions in the Security & Compliance Center.
- You can use the Add-RoleGroupMember cmdlet in Security & Compliance Center PowerShell to add a mail-enabled security group as a member of the eDiscovery Managers subgroup in the eDiscovery Manager role group. However, you can't add a mail-enabled security group to the eDiscovery Administrators subgroup. For details, see More information.

Assign eDiscovery permissions in the Security & Compliance Center

1. Go to https://protection.office.com.
2. Sign in using your work or school account.
3. In the left pane of the security and compliance center, select **Permissions**, and then select the checkbox next to **eDiscovery Manager**.
4. On the **eDiscovery Manager** flyout page, do one of the following based on the eDiscovery permissions that you want to assign.

   To make a user an eDiscovery Manager: Next to **eDiscovery Manager**, select **Edit**. In the **Choose eDiscovery Manager** section, select the **Choose eDiscovery Manager** hyperlink, and then select **+ Add**. Select the user (or users) you want to add as an eDiscovery manager, and then select **Add**. When you're finished adding users, select **Done**. Then, on the **Editing Choose eDiscovery Manager** flyout page, select **Save** to save the changes to the eDiscovery Manager membership.

   To make a user an eDiscovery Administrator: Next to **eDiscovery Manager**, select **Edit**. In the **Choose eDiscovery Administrator** section, Under **eDiscovery Administrators**, select **Choose eDiscovery Administrator**, select **Edit**, and then select **+ Add**. Select the user (or users) you want to add as an **eDiscovery Administrator**, and then **Add**. When you're finished adding users, select **Done**. Then, on the **Editing Choose eDiscovery Administrator** flyout page, select **Save** to save the changes to the eDiscovery Administrator membership.

> **NOTE**
>
> You can also use the **Add-eDiscoveryCaseAdmin** cmdlet to make a user an eDiscovery Administrator. However, the user must be assigned the Case Management role before you can use this cmdlet to make them an eDiscovery Administrator. For more information, see Add-eDiscoveryCaseAdmin.

On the **Permissions** page in the Security & Compliance Center, you can also assign users eDiscovery-related permissions by adding them to the Compliance Administrator, Organization Management, and Reviewer role groups. For a description of the eDiscovery-related RBAC roles assigned to each of these role groups, see RBAC roles related to eDiscovery.

RBAC roles related to eDiscovery

The following table lists the eDiscovery-related RBAC roles in the Security & Compliance Center, and indicates the built-in role groups that each role is assigned to by default.

ROLE	COMPLIANCE ADMINISTRATOR	EDISCOVERY MANAGER & ADMINISTRATOR	ORGANIZATION MANAGEMENT	REVIEWER
Case Management	✓	✓	✓	
Communication		✓		
Compliance Search	✓	✓	✓	
Custodian		✓		
Export		✓		

ROLE	COMPLIANCE ADMINISTRATOR	EDISCOVERY MANAGER & ADMINISTRATOR	ORGANIZATION MANAGEMENT	REVIEWER
Hold	✓	✓	✓	
Preview		✓		
Review		✓		✓
RMS Decrypt		✓		
Search And Purge			✓	

The following sections describe each of the eDiscovery-related RBAC roles listed in the previous table.

Case Management

This role lets users create, edit, delete, and control access to Core eDiscovery and Advanced eDiscovery cases in the Security & Compliance Center. As previously explained, a user must be assigned the Case Management role before you can use the **Add-eDiscoveryCaseAdmin** cmdlet to make them an eDiscovery Administrator.

For more information, see:

- Get started with Core eDiscovery
- Get started with Advanced eDiscovery

Communication

This role lets users manage all communications with the custodians identified in an Advanced eDiscovery case. This includes creating hold notifications, hold reminders, and escalations to management. The user can also track custodian acknowledgment of hold notifications and manage access to the custodian portal that is used by each custodian to track communications for the cases where they were identified as a custodian.

For more information, see Work with communications in Advanced eDiscovery.

Compliance Search

This role lets users run the Content Search tool in the Security & Compliance Center to search mailboxes and public folders, SharePoint Online sites, OneDrive for Business sites, Skype for Business conversations, Microsoft 365 groups, and Microsoft Teams, and Yammer groups. This role allows a user to get an estimate of the search results and create export reports, but additional roles are needed to initiate content search actions such as previewing, exporting, or deleting search results.

Users who are assigned the Compliance Search role but don't have the Preview role can preview the results of a search in which the preview action has been initiated by a user who is assigned the Preview role. The user without the Preview role can preview results for up to two weeks after the initial preview action was created.

Similarly, users who are assigned the Compliance Search role but don't have the Export role can download the results of a search in which the export action was initiated by a user who is assigned the Export role. The user without the Export role can download the results of a search for up to two weeks after the initial export action was created. After that, they can't download the results unless someone with the Export role restarts the export.

For more information, see Content search in Office 365.

Custodian

This role lets users identify and manage custodians for Advanced eDiscovery cases and use the information

from Azure Active Directory and other sources to find data sources associated with custodians. The user can associate other data sources such as mailboxes, SharePoint sites, and Teams with custodians in a case. The user can also place a legal hold on the data sources associated with custodians to preserve content in the context of a case.

For more information, see Work with custodians in Advanced eDiscovery.

Export

The role lets users export the results of a Content Search to a local computer. It also lets them prepare search results for analysis in Advanced eDiscovery.

For more information about exporting search results, see Export search results from Security & Compliance Center.

Hold

This role lets users place content on hold in mailboxes, public folders, sites, Skype for Business conversations, and Microsoft 365 groups. When content is on hold, content owners can still modify or delete the original content, but the content will be preserved until the hold is removed or until the hold duration expires.

For more information about holds, see:

- Create a hold in Core eDiscovery
- Create a hold in Advanced eDiscovery

Preview

This role lets users view a list of items that were returned from a Content Search. They can also open and view each item from the list to view its contents.

Review

This role lets users access review sets in Advanced eDiscovery. Users who are assigned this role can see and open the list of cases on the **eDiscovery > Advanced** page in the Microsoft 365 compliance center that they're members of. After the user accesses an Advanced eDiscovery case, they can select **Review sets** to access case data. This role doesn't allow the user to preview the results of a collection search that's associated with the case or do other search or case management tasks. Users with this role can only access the data in a review set.

RMS Decrypt

This role lets users view rights-protected email messages when previewing search results and export decrypted rights-protected email messages. This role also lets users view (and export) a file that's encrypted with a Microsoft encryption technology when the encrypted file is attached to an email message that's included in the results of an eDiscovery search. Additionally, this role lets users review and query encrypted email attachments that are added to a review set in Advanced eDiscovery. For more information about decryption in eDiscovery, see Decryption in Microsoft 365 eDiscovery tools.

Search And Purge

This role lets users perform bulk removal of data matching the criteria of a content search. For more information, see Search for and delete email messages in your organization.

More information

- **Why create an eDiscovery Administrator?** As previously explained, an eDiscovery Administrator is member of the eDiscovery Manager role group who can view and access all eDiscovery cases in your organization. This ability to access all the eDiscovery cases has two important purposes:
    - If a person who is the only member of an eDiscovery case leaves your organization, no one (including members of the Organization Management role group or another member of the

eDiscovery Manager role group) can access that eDiscovery case because they aren't a member of a case. In this situation, there would be no way to access the data in the case. But because an eDiscovery Administrator can access all eDiscovery cases in the organization, they can view the case and add themselves or another eDiscovery manager as a member of the case.

  - Because an eDiscovery Administrator can view and access all Core eDiscovery and Advanced eDiscovery cases, they can audit and oversee all cases and associated compliance searches. This can help to prevent any misuse of compliance searches or eDiscovery cases. And because eDiscovery Administrators can access potentially sensitive information in the results of a compliance search, you should limit the number of people who are eDiscovery Administrators.

- **Can I add a group as a member of the eDiscovery Manager role group?** As previously explained, you can add a mail-enabled security group as a member of the eDiscovery Managers subgroup in the eDiscovery Manager role group by using the **Add-RoleGroupMember** cmdlet in Security & Compliance Center PowerShell. For example, you can run the following command to add a mail-enabled security group to the eDiscovery Manager role group.

  ```
  Add-RoleGroupMember "eDiscovery Manager" -Member <name of security group>
  ```

  Exchange distribution groups and Microsoft 365 Groups aren't supported. You must use a mail-enabled security group, which you can create in Exchange Online PowerShell by running `New-DistributionGroup -Type Security`. You can also create a mail-enabled security group (and add members) in the Exchange admin center or in the Microsoft 365 admin center. It might take up to 60 minutes after you create it for a new mail-enabled security to be available to add to the eDiscovery Managers role group.

  Also as previously stated, you can't make a mail-enabled security group an eDiscovery Administrator by using the **Add-eDiscoveryCaseAdmin** cmdlet in Security & Compliance Center PowerShell. You can only add individual users as eDiscovery Administrators.

  You also can't add a mail-enabled security group as a member of a case.

Keyword queries and search conditions for Content Search and eDiscovery

2/18/2021 • 34 minutes to read • Edit Online

This topic describes the email and document properties that you can search for in email items in Exchange Online and documents stored on SharePoint and OneDrive for Business sites by using the Content Search feature in the Microsoft 365 compliance center. You can also use the ***-ComplianceSearch** cmdlets in Security & Compliance Center PowerShell to search for these properties. The topic also describes:

- Using Boolean search operators, search conditions, and other search query techniques to refine your search results.
- Searching for sensitive data types and custom sensitive data types in SharePoint and OneDrive for Business.
- Searching for site content that's shared with users outside of your organization

For step-by-step instructions on how to create a Content Search, see Content Search.

> **NOTE**
> Content Search in the Microsoft 365 compliance center and the corresponding ***-ComplianceSearch** cmdlets in Security & Compliance Center PowerShell use the Keyword Query Language (KQL). For more detailed information, see Keyword Query Language syntax reference.

Searchable email properties

The following table lists email message properties that can be searched by using the Content Search feature in the Microsoft 365 compliance center or by using the **New-ComplianceSearch** or the **Set-ComplianceSearch** cmdlet. The table includes an example of the *property:value* syntax for each property and a description of the search results returned by the examples. You can type these `property:value` pairs in the keywords box for a Content Search.

> **NOTE**
> When searching email properties, it's not possible to search for items in which the specified property is empty or blank. For example, using the *property:value* pair of **subject:""** to search for email messages with an empty subject line will return zero results. This also applies when searching site and contact properties.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLES	SEARCH RESULTS RETURNED BY THE EXAMPLES

PROPERTY	PROPERTY DESCRIPTION	EXAMPLES	SEARCH RESULTS RETURNED BY THE EXAMPLES

AttachmentNames	The names of files attached to an email message.	`attachmentnames:annualreport.ppt`
		`attachmentnames:annual*`
		`attachmentnames:.pptx`
			Messages that have an attached file named annualreport.ppt. In the second example, using the wildcard returns messages with the word "annual" in the file name of an attachment. The third example returns all attachments with the pptx file extension.

Bcc	The Bcc field of an email message.[1]	`bcc:pilarp@contoso.com`
		`bcc:pilarp`
		`bcc:"Pilar Pinilla"`
			All examples return messages with Pilar Pinilla included in the Bcc field.

Category	The categories to search. Categories can be defined by users by using Outlook or Outlook on the web (formerly known as Outlook Web App). The possible values are: blue green orange purple red yellow	`category:"Red Category"`	Messages that have been assigned the red category in the source mailboxes.
Cc	The Cc field of an email message.[1]	`cc:pilarp@contoso.com` `cc:"Pilar Pinilla"`	In both examples, messages with Pilar Pinilla specified in the Cc field.
Folderid	The folder ID (GUID) of a specific mailbox folder. If you use this property, be sure to search the mailbox that the specified folder is located in. Only the specified folder will be searched. Any subfolders in the folder won't be searched. To search subfolders, you need to use the Folderid property for the subfolder you want to search. For more information about searching for the Folderid property and using a script to obtain the folder IDs for a specific mailbox, see [Use Content Search for targeted collections]().	`folderid:4D6DD7F943C29041A65787E30F02AD1F00000000013A0000` `folderid:2370FB455F82FC44BE31397F47B632A70000000001160000 AND participants:garthf@contoso.com`	The first example returns all items in the specified mailbox folder. The second example returns all items in the specified mailbox folder that were sent or received by garthf@contoso.com.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLES	SEARCH RESULTS RETURNED BY THE EXAMPLES
From	The sender of an email message.[1]	`from:pilarp@contoso.com` `from:contoso.com`	Messages sent by the specified user or sent from a specified domain.
HasAttachment	Indicates whether a message has an attachment. Use the values **true** or **false**.	`from:pilar@contoso.com AND hasattachment:true`	Messages sent by the specified user that have attachments.

Importance	The importance of an email message, which a sender can specify when sending a message. By default, messages are sent with normal importance, unless the sender sets the importance as **high** or **low**.	`importance:high`
		`importance:medium`
		`importance:low`
			Messages that are marked as high importance, medium importance, or low importance.

IsRead	Indicates whether messages have been read. Use the values **true** or **false**.	`isread:true` `isread:false`	The first example returns messages with the IsRead property set to **True**. The second example returns messages with the IsRead property set to **False**.
ItemClass	Use this property to search specific third-party data types that your organization imported to Office 365. Use the following syntax for this property: `itemclass:ipm.externaldata.<third-party data type>*`	`itemclass:ipm.externaldata.Facebook* AND subject:contoso` `itemclass:ipm.externaldata.Twitter* AND from:"Ann Beebe" AND "Northwind Traders"`	The first example returns Facebook items that contain the word "contoso" in the Subject property. The second example returns Twitter items that were posted by Ann Beebe and that contain the keyword phrase "Northwind Traders". For a complete list of values to use for third-party data types for the ItemClass property, see Use Content Search to search third-party data that was imported to Office 365.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLES	SEARCH RESULTS RETURNED BY THE EXAMPLES

Kind	The type of email message to search for. Possible values:
	contacts
	docs
	email
	externaldata
	faxes
	im
	journals
	meetings
	microsoftteams (returns items from chats, meetings, and calls in Microsoft Teams)
	notes
	posts
	rssfeeds
	tasks
	voicemail
		`kind:email`
		`kind:email OR kind:im OR kind:voicemail`
		`kind:externaldata`
			The first example returns email messages that meet the search criteria. The second example returns email messages, instant messaging conversations (including Skype for Business conversations and chats in Microsoft Teams), and voice messages that meet the search criteria. The third example returns items that were imported to mailboxes in Microsoft 365 from third-party data sources, such as Twitter, Facebook, and Cisco Jabber, that meet the search criteria. For more information, see Archiving third-party data in Office 365.

Participants	All the people fields in an email message. These fields are From, To, Cc, and Bcc.[1]	`participants:garthf@contoso.com` `participants:contoso.com`	Messages sent by or sent to garthf@contoso.com. The second example returns all messages sent by or sent to a user in the contoso.com domain.
Received	The date that an email message was received by a recipient.	`received:04/15/2016` `received>=01/01/2016 AND received<=03/31/2016`	Messages that were received on April 15, 2016. The second example returns all messages received between January 1, 2016 and March 31, 2016.
Recipients	All recipient fields in an email message. These fields are To, Cc, and Bcc.[1]	`recipients:garthf@contoso.com` `recipients:contoso.com`	Messages sent to garthf@contoso.com. The second example returns messages sent to any recipient in the contoso.com domain.
Sent	The date that an email message was sent by the sender.	`sent:07/01/2016` `sent>=06/01/2016 AND sent<=07/01/2016`	Messages that were sent on the specified date or sent within the specified date range.
Size	The size of an item, in bytes.	`size>26214400` `size:1..1048567`	Messages larger than 25?? MB. The second example returns messages from 1 through 1,048,567 bytes (1 MB) in size.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLES	SEARCH RESULTS RETURNED BY THE EXAMPLES
Subject	The text in the subject line of an email message. **Note:** When you use the Subject property in a query, the search returns all messages in which the subject line contains the text you're searching for. In other words, the query doesn't return only those messages that have an exact match. For example, if you search for `subject:"Quarterly Financials"`, your results will include messages with the subject "Quarterly Financials 2018".	`subject:"Quarterly Financials"` `subject:northwind`	Messages that contain the phrase "Quarterly Financials" anywhere in the text of the subject line. The second example returns all messages that contain the word northwind in the subject line.

To	The To field of an email message.[1]	`to:annb@contoso.com`
		`to:annb`
		`to:"Ann Beebe"`
			All examples return messages where Ann Beebe is specified in the To: line.

NOTE

[1] For the value of a recipient property, you can use email address (also called *user principal name* or UPN), display name, or alias to specify a user. For example, you can use annb@contoso.com, annb, or "Ann Beebe" to specify the user Ann Beebe.

Recipient expansion

When searching any of the recipient properties (From, To, Cc, Bcc, Participants, and Recipients), Microsoft 365 attempts to expand the identity of each user by looking them up in Azure Active Directory (Azure AD). If the user is found in Azure AD, the query is expanded to include the user's email address (or UPN), alias, display name, and LegacyExchangeDN. For example, a query such as `participants:ronnie@contoso.com` expands to `participants:ronnie@contoso.com OR participants:ronnie OR participants:"Ronald Nelson" OR participants:"<LegacyExchangeDN>"`.

To prevent recipient expansion, add a wild card character (asterisk) to the end of the email address and use a reduced domain name; for example, `participants:"ronnie@contoso*"` Be sure to surround the email address with double quotation marks.

However, be aware that preventing recipient expansion in the search query may result in relevant items not being returned in the search results. Email messages in Exchange can be saved with different text formats in the recipient fields. Recipient expansion is intended to help mitigate this fact by returning messages that may contain different text formats. So preventing recipient expansion may result in the search query not returning all items that may be relevant to your investigation.

NOTE

If you need to review or reduce the items returned by a search query due to recipient expansion, consider using Advanced eDiscovery. You can search for messages (taking advantage of recipient expansion), add them to a review set, and then use review set queries or filters to review or narrow the results. For more information, see Collect data for a case and Query the data in a review set.

Searchable site properties

The following table lists some of the SharePoint and OneDrive for Business properties that can be searched by using the Content Search feature in the Security & Compliance Center or by using the **New-ComplianceSearch** or the **Set-ComplianceSearch** cmdlet. The table includes an example of the *property:value* syntax for each property and a description of the search results returned by the examples.

For a complete list of SharePoint properties that can be searched, see Overview of crawled and managed properties in SharePoint. Properties marked with a **Yes** in the **Queryable** column can be searched.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLE	SEARCH RESULTS RETURNED BY THE EXAMPLES
Author	The author field from Office documents, which persists if a document is copied. For example, if a user creates a document and the emails it to someone else who then uploads it to SharePoint, the document will still retain the original author. Be sure to use the user's display name for this property.	`author:"Garth Fort"`	All documents that are authored by Garth Fort.
ContentType	The SharePoint content type of an item, such as Item, Document, or Video.	`contenttype:document`	All documents would be returned.
Created	The date that an item is created.	`created>=06/01/2016`	All items created on or after June 1, 2016.
CreatedBy	The person that created or uploaded an item. Be sure to use the user's display name for this property.	`createdby:"Garth Fort"`	All items created or uploaded by Garth Fort.
DetectedLanguage	The language of an item.	`detectedlanguage:english`	All items in English.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLE	SEARCH RESULTS RETURNED BY THE EXAMPLES
DocumentLink	The path (URL) of a specific folder on a SharePoint or OneDrive for Business site. If you use this property, be sure to search the site that the specified folder is located in. To return items located in subfolders of the folder that you specify for the documentlink property, you have to add /* to the URL of the specified folder; for example, `documentlink: "https://contoso.sharepoint.com/Shared Documents/*"` For more information about searching for the documentlink property and using a script to obtain the documentlink URLs for folders on a specific site, see Use Content Search for targeted collections.	`documentlink:"https://contoso-my.sharepoint.com/personal/garthf_contoso_com/Documents/Private` `documentlink:"https://contoso-my.sharepoint.com/personal/garthf_contoso_com/Documents/Shared with Everyone/*" AND filename:confidential`	The first example returns all items in the specified OneDrive for Business folder. The second example returns documents in the specified site folder (and all subfolders) that contain the word "confidential" in the file name.
FileExtension	The extension of a file; for example, docx, one, pptx, or xlsx.	`fileextension:xlsx`	All Excel files (Excel 2007 and later)
FileName	The name of a file.	`filename:"marketing plan"` `filename:estimate`	The first example returns files with the exact phrase "marketing plan" in the title. The second example returns files with the word "estimate" in the file name.
LastModifiedTime	The date that an item was last changed.	`lastmodifiedtime>=05/01/2016` `lastmodifiedtime>=05/10/2016 AND lastmodifiedtime<=06/1/2016`	The first example returns items that were changed on or after May 1, 2016. The second example returns items changed between May 1, 2016 and June 1, 2016.
ModifiedBy	The person who last changed an item. Be sure to use the user's display name for this property.	`modifiedby:"Garth Fort"`	All items that were last changed by Garth Fort.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLE	SEARCH RESULTS RETURNED BY THE EXAMPLES
Path	The path (URL) of a specific site in a SharePoint or OneDrive for Business site. To return items located in folders in the site that you specify for the path property, you have to add /* to the URL of the specified site; for example, `path: "https://contoso.sharepoint.com/Shared Documents/*"` **Note:** Using the `Path` property to search OneDrive locations won't return media files, such as .png, .tiff, or .wav files, in the search results. Use a different site property in your search query to search for media files in OneDrive folders.	`path:"https://contoso-my.sharepoint.com/personal/garthf_contoso_com/"` `path:"https://contoso-my.sharepoint.com/personal/garthf_contoso_com/*" AND filename:confidential`	The first example returns all items in the specified OneDrive for Business site. The second example returns documents in the specified site (and folders in the site) that contain the word "confidential" in the file name.
SharedWithUsersOWSUser	Documents that have been shared with the specified user and displayed on the **Shared with me** page in the user's OneDrive for Business site. These are documents that have been explicitly shared with the specified user by other people in your organization. When you export documents that match a search query that uses the SharedWithUsersOWSUser property, the documents are exported from the original content location of the person who shared the document with the specified user. For more information, see Searching for site content shared within your organization.	`sharedwithusersowsuser:garthf` `sharedwithusersowsuser:"garthf@contoso.com"`	Both examples return all internal documents that have been explicitly shared with Garth Fort and that appear on the **Shared with me** page in Garth Fort's OneDrive for Business account.
Site	The URL of a site or group of sites in your organization.	`site:"https://contoso-my.sharepoint.com"` `site:"https://contoso.sharepoint.com/sites/teamA"`	The first example returns items from the OneDrive for Business sites for all users in the organization. The second example returns items from all team sites.
Size	The size of an item, in bytes.	`size>=1` `size:1..10000`	The first example returns items larger than 1 byte. The second example returns items from 1 through 10,000 bytes in size.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLE	SEARCH RESULTS RETURNED BY THE EXAMPLES
Title	The title of the document. The Title property is metadata that's specified in Microsoft Office documents. It's different from the file name of the document.	`title:"communication plan"`	Any document that contains the phrase "communication plan" in the Title metadata property of an Office document.

Searchable contact properties

The following table lists the contact properties that are indexed and that you can search for using Content Search. These are the properties that are available for users to configure for the contacts (also called personal contacts) that are located in the personal address book of a user's mailbox. To search for contacts, you can select the mailboxes to search and then use one or more contact properties in the keyword query.

> **TIP**
> To search for values that contain spaces or special characters, use double quotation marks (" ") to contain the phrase; for example, `businessaddress:"123 Main Street"`.

PROPERTY	PROPERTY DESCRIPTION		
BusinessAddress	The address in the **Business Address** property. The property is also called the **Work** address on the contact properties page.		
BusinessPhone	The phone number in any of the **Business Phone** number properties.		
CompanyName	The name in the **Company** property.		
Department	The name in the **Department** property.		
DisplayName	The display name of the contact. This is the name in the **Full Name** property of the contact.		
EmailAddress	The address for any email address property for the contact. Users can add multiple email addresses for a contact. Using this property would return contacts that match any of the contact's email addresses.		

PROPERTY	PROPERTY DESCRIPTION		
FileAs	The **File as** property. This property is used to specify how the contact is listed in the user's contact list. For example, a contact could be listed as *FirstName,LastName* or *LastName,FirstName*.		
GivenName	The name in the **First Name** property.		
HomeAddress	The address in any of the **Home** address properties.		
HomePhone	The phone number in any of the **Home** phone number properties.		
IMAddress	The IM address property, which is typically an email address used for instant messaging.		
MiddleName	The name in the **Middle** name property.		
MobilePhone	The phone number in the **Mobile** phone number property.		
Nickname	The name in the **Nickname** property.		
OfficeLocation	The value in **Office** or **Office location** property.		
OtherAddress	The value for the **Other** address property.		
Surname	The name in the **Last** name property.		
Title	The title in the **Job title** property.		

Searchable sensitive data types

You can use eDiscovery search tools in the Microsoft 365 compliance center to search for sensitive data, such as credit card numbers or social security numbers, that is stored in documents on SharePoint and OneDrive for Business sites. You can do this by using the `SensitiveType` property and the name (or ID) of a sensitive information type in a keyword query. For example, the query `SensitiveType:"Credit Card Number"` returns documents that contain a credit card number. The query `SensitiveType:"U.S. Social Security Number (SSN)"` returns documents that contain a U.S. social security number.

To see a list of the sensitive information types that you can search for, go to **Data classifications** > **Sensitive info types** in the Microsoft 365 compliance center. Or you can use the **Get-DlpSensitiveInformationType** cmdlet in Security & Compliance Center PowerShell to display a list of sensitive information types.

For more information about creating queries using the `SensitiveType` property, see Form a query to find sensitive data stored on sites.

Limitations for searching sensitive data types

- To search for custom sensitive information types, you have to specify the ID of the sensitive information type in the `SensitiveType` property. Using the name of a custom sensitive information type (as shown in the example for built-in sensitive information types in the previous section) will return no results. Use the **Publisher** column on the **Sensitive info types** page in the compliance center (or the **Publisher** property in PowerShell) to differentiate between built-in and custom sensitive information types. Built-in sensitive data types have a value of `Microsoft Corporation` for the **Publisher** property.

  To display the name and ID for the custom sensitive data types in your organization, run the following command in Security & Compliance Center PowerShell:

  ```
  Get-DlpSensitiveInformationType | Where-Object {$_.Publisher -ne "Microsoft Corporation"} | FT
  Name,Id
  ```

  Then you can use the ID in the `SensitiveType` search property to return documents that contain the custom sensitive data type; for example, `SensitiveType:7e13277e-6b04-3b68-94ed-1aeb9d47de37`

- You can't use sensitive information types and the `SensitiveType` search property to search for sensitive data at-rest in Exchange Online mailboxes. However, you can use data loss prevention (DLP) policies to protect sensitive email data in transit. For more information, see Overview of data loss prevention policies and Search for and find personal data.

Search operators

Boolean search operators, such as **AND**, **OR**, and **NOT**, help you define more-precise searches by including or excluding specific words in the search query. Other techniques, such as using property operators (such as `>=` or `..`), quotation marks, parentheses, and wildcards, help you refine a search query. The following table lists the operators that you can use to narrow or broaden search results.

OPERATOR	USAGE	DESCRIPTION	
AND	keyword1 AND keyword2	Returns items that include all of the specified keywords or `property:value` expressions. For example, `from:"Ann Beebe" AND subject:northwind` would return all messages sent by Ann Beebe that contained the word northwind in the subject line. [2]	

OPERATOR	USAGE	DESCRIPTION	
+	keyword1 + keyword2 + keyword3	Returns items that contain *either* `keyword2` or `keyword3` *and* that also contain `keyword1`. Therefore, this example is equivalent to the query `(keyword2 OR keyword3) AND keyword1`. The query `keyword1 + keyword2` (with a space after the **+** symbol) isn't the same as using the **AND** operator. This query would be equivalent to `"keyword1 + keyword2"` and return items with the exact phase `"keyword1 + keyword2"`.	
OR	keyword1 OR keyword2	Returns items that include one or more of the specified keywords or `property:value` expressions. [2]	
NOT	keyword1 NOT keyword2 NOT from:"Ann Beebe" NOT kind:im	Excludes items specified by a keyword or a `property:value` expression. In the second example excludes messages sent by Ann Beebe. The third example excludes any instant messaging conversations, such as Skype for Business conversations that are saved to the Conversation History mailbox folder. [2]	
-	keyword1 -keyword2	The same as the **NOT** operator. So this query returns items that contain `keyword1` and would exclude items that contain `keyword2`.	
NEAR	keyword1 NEAR(n) keyword2	Returns items with words that are near each other, where n equals the number of words apart. For example, `best NEAR(5) worst` returns any item where the word "worst" is within five words of "best". If no number is specified, the default distance is eight words. [2]	

OPERATOR	USAGE	DESCRIPTION	
:	property:value	The colon (:) in the `property:value` syntax specifies that the value of the property being searched for contains the specified value. For example, `recipients:garthf@contoso.com` returns any message sent to garthf@contoso.com.	
=	property=value	The same as the : operator.	
<	property<value	Denotes that the property being searched is less than the specified value. [1]	
>	property>value	Denotes that the property being searched is greater than the specified value.[1]	
<=	property<=value	Denotes that the property being searched is less than or equal to a specific value.[1]	
>=	property>=value	Denotes that the property being searched is greater than or equal to a specific value.[1]	
..	property:value1..value2	Denotes that the property being searched is greater than or equal to value1 and less than or equal to value2.[1]	
" "	"fair value" subject:"Quarterly Financials"	Use double quotation marks (" ") to search for an exact phrase or term in keyword and `property:value` search queries.	

OPERATOR	USAGE	DESCRIPTION	
*	cat* subject:set*	Prefix wildcard searches (where the asterisk is placed at the end of a word) match for zero or more characters in keywords or `property:value` queries. For example, `title:set*` returns documents that contain the word set, setup, and setting (and other words that start with "set") in the document title. **Note:** You can use only prefix wildcard searches; for example, **cat*** or **set***. Suffix searches (***cat**), infix searches (**c*t**), and substring searches (***cat***) are not supported.	
( )	(fair OR free) AND (from:contoso.com) (IPO OR initial) AND (stock OR shares) (quarterly financials)	Parentheses group together Boolean phrases, `property:value` items, and keywords. For example, `(quarterly financials)` returns items that contain the words quarterly and financials.	

NOTE

[1] Use this operator for properties that have date or numeric values.

[2] Boolean search operators must be uppercase; for example, **AND**. If you use a lowercase operator, such as **and**, it will be treated as a keyword in the search query.

Search conditions

You can add conditions to a search query to narrow a search and return a more refined set of results. Each condition adds a clause to the KQL search query that is created and run when you start the search.

Conditions for common properties

Conditions for mail properties

Conditions for document properties

Operators used with conditions

Guidelines for using conditions

Examples of using conditions in search queries

Conditions for common properties

Create a condition using common properties when searching mailboxes and sites in the same search. The following table lists the available properties to use when adding a condition.

CONDITION	DESCRIPTION
Date	For email, the date a message was received by a recipient or sent by the sender. For documents, the date a document was last modified.
Sender/Author	For email, the person who sent a message. For documents, the person cited in the author field from Office documents. You can type more than one name, separated by commas. Two or more values are logically connected by the **OR** operator.
Size (in bytes)	For both email and documents, the size of the item (in bytes).
Subject/Title	For email, the text in the subject line of a message. For documents, the title of the document. As previously explained, the Title property is metadata specified in Microsoft Office documents. You can type the name of more than one subject/title, separated by commas. Two or more values are logically connected by the **OR** operator.
Compliance label	For both email and documents, retention labels that have been assigned to messages and documents automatically by autolabel policies or retention labels that have been manually assigned by users. Retention labels are used to classify email and documents for information governance and enforce retention rules based on the settings defined by the label. You can type part of the retention label name and use a wildcard or type the complete label name. For more information about retention labels, see Learn about retention policies and retention labels.

Conditions for mail properties

Create a condition using mail properties when searching mailboxes or public folders. The following table lists the email properties that you can use for a condition. These properties are a subset of the email properties that were previously described. These descriptions are repeated for your convenience.

CONDITION	DESCRIPTION
Message kind	The message type to search. This is the same property as the Kind email property. Possible values: contacts docs email externaldata faxes im journals meetings microsoftteams notes posts rssfeeds tasks voicemail

CONDITION	DESCRIPTION
Participants	All the people fields in an email message. These fields are From, To, Cc, and Bcc.
Type	The message class property for an email item. This is the same property as the ItemClass email property. It's also a multi-value condition. So to select multiple message classes, hold the **CTRL** key and then click two or more message classes in the drop-down list that you want to add to the condition. Each message class that you select in the list will be logically connected by the **OR** operator in the corresponding search query. For a list of the message classes (and their corresponding message class ID) that are used by Exchange and that you can select in the **Message class** list, see Item Types and Message Classes.
Received	The date that an email message was received by a recipient. This is the same property as the Received email property.
Recipients	All recipient fields in an email message. These fields are To, Cc, and Bcc.
Sender	The sender of an email message.
Sent	The date that an email message was sent by the sender. This is the same property as the Sent email property.
Subject	The text in the subject line of an email message.
To	The recipient of an email message in the To field.

Conditions for document properties

Create a condition using document properties when searching for documents on SharePoint and OneDrive for Business sites. The following table lists the document properties that you can use for a condition. These properties are a subset of the site properties that were previously described. These descriptions are repeated for your convenience.

CONDITION	DESCRIPTION
Author	The author field from Office documents, which persists if a document is copied. For example, if a user creates a document and the emails it to someone else who then uploads it to SharePoint, the document will still retain the original author.
Title	The title of the document. The Title property is metadata that's specified in Office documents. It's different than the file name of the document.
Created	The date that a document is created.
Last modified	The date that a document was last changed.

CONDITION	DESCRIPTION
File type	The extension of a file; for example, docx, one, pptx, or xlsx. This is the same property as the FileExtension site property.

Operators used with conditions

When you add a condition, you can select an operator that is relevant to type of property for the condition. The following table describes the operators that are used with conditions and lists the equivalent that is used in the search query.

OPERATOR	QUERY EQUIVALENT	DESCRIPTION
After	`property>date`	Used with date conditions. Returns items that were sent, received, or modified after the specified date.
Before	`property<date`	Used with date conditions. Returns items that were sent, received, or modified before the specified date.
Between	`date..date`	Use with date and size conditions. When used with a date condition, returns items there were sent, received, or modified within the specified date range. When used with a size condition, returns items whose size is within the specified range.
Contains any of	`(property:value) OR (property:value)`	Used with conditions for properties that specify a string value. Returns items that contain any part of one or more specified string values.
Doesn't contain any of	`-property:value` `NOT property:value`	Used with conditions for properties that specify a string value. Returns items that don't contain any part of the specified string value.
Doesn't equal any of	`-property=value` `NOT property=value`	Used with conditions for properties that specify a string value. Returns items that don't contain the specific string.
Equals	`size=value`	Returns items that are equal to the specified size.[1]
Equals any of	`(property=value) OR (property=value)`	Used with conditions for properties that specify a string value. Returns items that are an exact match of one or more specified string values.
Greater	`size>value`	Returns items where the specified property is greater than the specified value.[1]

OPERATOR	QUERY EQUIVALENT	DESCRIPTION
Greater or equal	`size>=value`	Returns items where the specified property is greater than or equal to the specified value.[1]
Less	`size<value`	Returns items that are greater than or equal to the specific value.[1]
Less or equal	`size<=value`	Returns items that are greater than or equal to the specific value.[1]
Not equal	`size<>value`	Returns items that don't equal the specified size.[1]

NOTE

[1] This operator is available only for conditions that use the Size property.

Guidelines for using conditions

Keep the following in mind when using search conditions.

- A condition is logically connected to the keyword query (specified in the keyword box) by the **AND** operator. That means that items have to satisfy both the keyword query and the condition to be included in the results. This is how conditions help to narrow your results.
- If you add two or more unique conditions to a search query (conditions that specify different properties), those conditions are logically connected by the **AND** operator. That means only items that satisfy all the conditions (in addition to any keyword query) are returned.
- If you add more than one condition for the same property, those conditions are logically connected by the **OR** operator. That means items that satisfy the keyword query and any one of the conditions are returned. So, groups of the same conditions are connected to each other by the **OR** operator and then sets of unique conditions are connected by the **AND** operator.
- If you add multiple values (separated by commas or semi-colons) to a single condition, those values are connected by the **OR** operator. That means items are returned if they contain any of the specified values for the property in the condition.
- The search query that is created by using the keywords box and conditions is displayed on the **Search** page, in the details pane for the selected search. In a query, everything to the right of the notation `(c:c)` indicates conditions that are added to the query.
- Conditions only add properties to the search query; the don't add operators. This is why the query displayed in the detail pane doesn't show operators to the right of the `(c:c)` notation. KQL adds the logical operators (according to the previously explained rules) when the executing the query.
- You can use the drag and drop control to resequence the order of conditions. Click on the control for a condition and move it up or down.
- As previously explained, some condition properties allow you to type multiple values. Each value is logically connected by the **OR** operator. This results in the same logic as having multiple instances of the same condition, where each has a single value. The following illustrations show an example of a single condition with multiple values and an example of multiple conditions (for the same property) with a single value. Both examples result in the same query:
  `(filetype:docx) OR (filetype:pptx) OR (filetype:xlsx)`

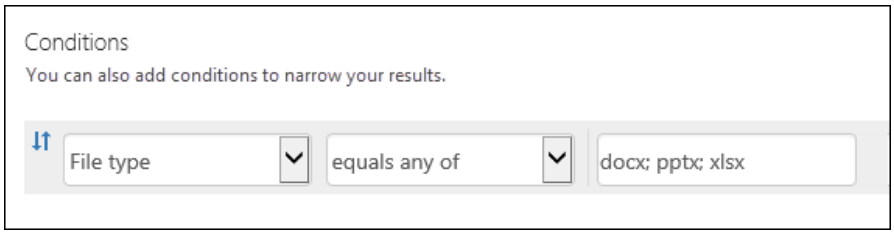

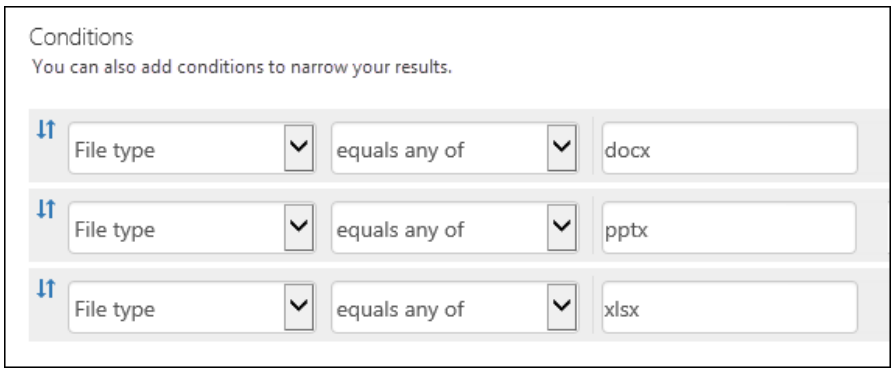

TIP

If a condition accepts multiple values, we recommend that you use a single condition and specify multiple values (separated by commas or semi-colons). This helps ensure the query logic that's applied is what you intend.

Examples of using conditions in search queries

The following examples show the GUI-based version of a search query with conditions, the search query syntax that is displayed in the details pane of the selected search (which is also returned by the **Get-ComplianceSearch** cmdlet), and the logic of the corresponding KQL query.

Example 1

This example returns documents on SharePoint and OneDrive for Business sites that contain a credit card number and were last modified before January 1, 2016.

GUI

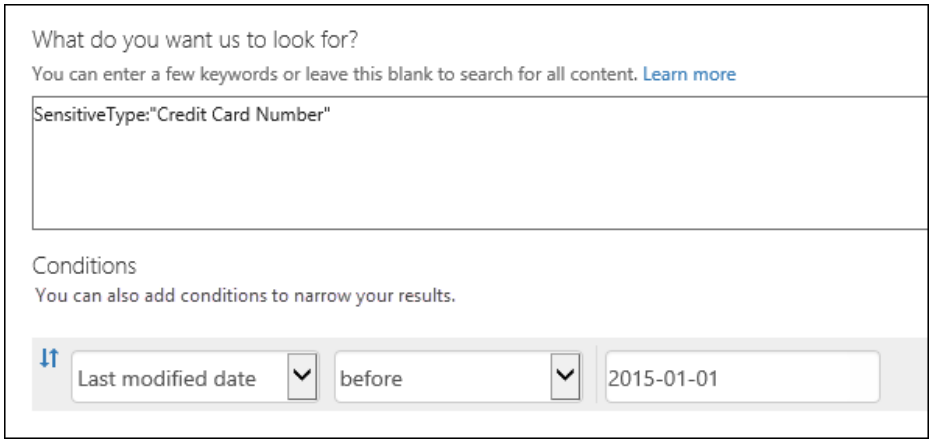

Search query syntax

```
SensitiveType:"Credit Card Number"(c:c)(lastmodifiedtime<2016-01-01)
```

Search query logic

```
SensitiveType:"Credit Card Number" AND (lastmodifiedtime<2016-01-01)
```

Example 2

This example returns email items or documents that contain the keyword "report", that were sent or created before April 1, 2105, and that contain the word "northwind" in the subject field of email messages or in the title

property of documents. The query excludes Web pages that meet the other search criteria.

GUI

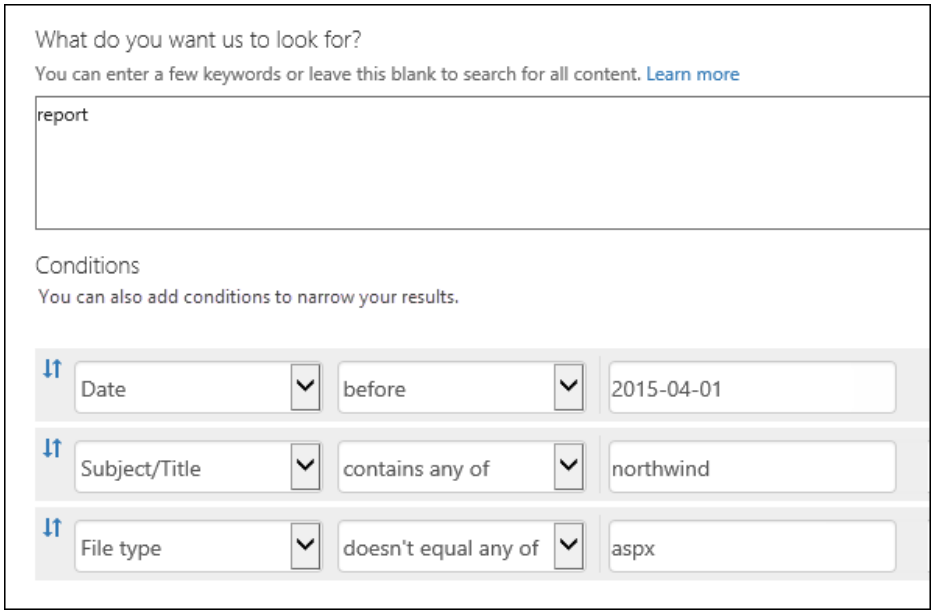

Search query syntax

```
report(c:c)(date<2016-04-01)(subjecttitle:"northwind")(-filetype:aspx)
```

Search query logic

```
report AND (date<2016-04-01) AND (subjecttitle:"northwind") NOT (filetype:aspx)
```

Example 3

This example returns email messages or calendar meetings that were sent between 12/1/2016 and 11/30/2016 and that contain words that start with "phone" or "smartphone".

GUI

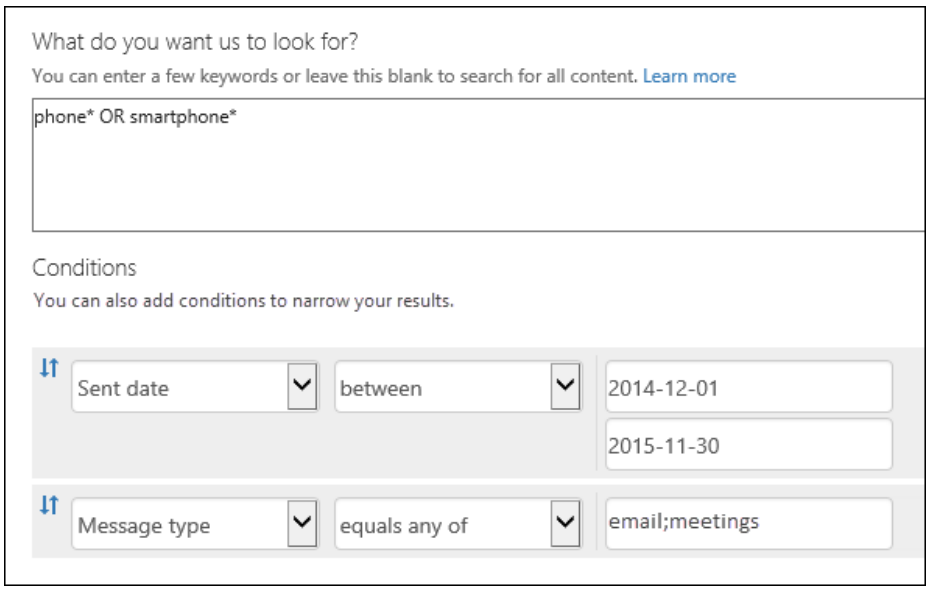

Search query syntax

```
phone* OR smartphone*(c:c)(sent=2016-12-01..2016-11-30)(kind="email")(kind="meetings")
```

Search query logic

```
phone* OR smartphone* AND (sent=2016-12-01..2016-11-30) AND ((kind="email") OR (kind="meetings"))
```

Special characters

Some special characters are not included in the search index and therefore are not searchable. This also includes the special characters that represent search operators in the search query. Here's a list of special characters that are either replaced by a blank space in the actual search query or cause a search error.

```
+ - = : ! @ # % ^ & ; _ / ? ( ) [ ] { }
```

Searching for site content shared with external users

You can also use the Content Search feature in the Security & Compliance Center to search for documents stored on SharePoint and OneDrive for Business sites that have been shared with people outside of your organization. This can help you identify sensitive or proprietary information that's being shared outside your organization. You can do this by using the `ViewableByExternalUsers` property in a keyword query. This property returns documents or sites that have been shared with external users by using one of the following sharing methods:

- A sharing invitation that requires users to sign in to your organization as an authenticated user.
- An anonymous guest link, which allows anyone with this link to access the resource without having to be authenticated.

Here are some examples:

- The query `ViewableByExternalUsers:true AND SensitiveType:"Credit Card Number"` returns all items that have been shared with people outside your organization and contain a credit card number.
- The query
  ```
  ViewableByExternalUsers:true AND ContentType:document AND
  site:"https://contoso.sharepoint.com/Sites/Teams"
  ```
  returns a list of documents on all team sites in the organization that have been shared with external users.

> **TIP**
> A search query such as `ViewableByExternalUsers:true AND ContentType:document` might return a lot of .aspx files in the search results. To eliminate these (or other types of files), you can use the `FileExtension` property to exclude specific file types; for example `ViewableByExternalUsers:true AND ContentType:document NOT FileExtension:aspx`.

What is considered content that is shared with people outside your organization? Documents in your organization's SharePoint and OneDrive for Business sites that are shared by sending a sharing invitation or that are shared in public locations. For example, the following user activities result in content that is viewable by external users:

- A user shares a file or folder with a person outside your organization.
- A user creates and sends a link to a shared file to a person outside your organization. This link allows the external user to view (or edit) the file.
- A user sends a sharing invitation or a guest link to a person outside your organization to view (or edit) a shared file.

Issues using the ViewableByExternalUsers property

While the `ViewableByExternalUsers` property represents the status of whether a document or site is shared with external users, there are some caveats to what this property does and doesn't reflect. In the following scenarios, the value of the `ViewableByExternalUsers` property won't be updated, and the results of a Content Search query that uses this property may be inaccurate.

- Changes to sharing policy, such as turning off external sharing for a site or for the organization. The property will still show previously shared documents as being externally accessible even though external access might have been revoked.
- Changes to group membership, such as adding or removing external users to Microsoft 365 Groups or Microsoft 365 security groups. The property won't automatically be updated for items the group has access to.
- Sending sharing invitations to external users where the recipient hasn't accepted the invitation, and therefore doesn't yet have access to the content.

In these scenarios, the `ViewableByExternalUsers` property won't reflect the current sharing status until the site or document library is recrawled and reindexed.

Searching for site content shared within your organization

As previously explained, you can use the `SharedWithUsersOWSUser` property so search for documents that have been shared between people in your organization. When a person shares a file (or folder) with another user inside your organization, a link to the shared file appears on the **Shared with me** page in the OneDrive for Business account of the person who the file was shared with. For example, to search for the documents that have been shared with Sara Davis, you can use the query `SharedWithUsersOWSUser:"sarad@contoso.com"`. If you export the results of this search, the original documents (located in the content location of the person who shared the documents with Sara) will be downloaded.

Documents must be explicitly shared with a specific user to be returned in search results when using the `SharedWithUsersOWSUser` property. For example, when a person shares a document in their OneDrive account, they have the option to share it with anyone (inside or outside the organization), share it only with people inside the organization, or share it with a specific person. Here's a screenshot of the **Share** window in OneDrive, that shows the three sharing options.

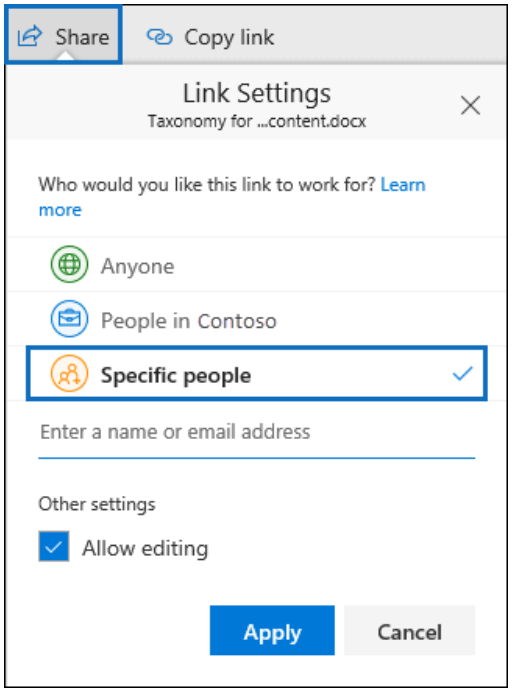

Only documents that are shared by using the third option (shared with **Specific people**) will be returned by a search query that uses the `SharedWithUsersOWSUser` property.

Searching for Skype for Business conversations

You can use the following keyword query to specifically search for content in Skype for Business conversations:

```
kind:im
```

The previous search query also returns chats from Microsoft Teams. To prevent this, you can narrow the search results to include only Skype for Business conversations by using the following keyword query:

```
kind:im AND subject:conversation
```

The previous keyword query excludes chats in Microsoft Teams because Skype for Business conversations are saved as email messages with a Subject line that starts with the word "Conversation".

To search for Skype for Business conversations that occurred within a specific date range, use the following keyword query:

```
kind:im AND subject:conversation AND (received=startdate..enddate)
```

Search tips and tricks

- Keyword searches are not case-sensitive. For example, **cat** and **CAT** return the same results.
- The Boolean operators **AND**, **OR**, **NOT**, and **NEAR** must be uppercase.
- A space between two keywords or two `property:value` expressions is the same as using **AND**. For example, `from:"Sara Davis" subject:reorganization` returns all messages sent by Sara Davis that contain the word reorganization in the subject line.
- Use syntax that matches the `property:value` format. Values are not case-sensitive, and they can't have a space after the operator. If there is a space, your intended value will be a full-text search. For example `to: pilarp` searches for "pilarp" as a keyword, rather than for messages that were sent to pilarp.
- When searching a recipient property, such as To, From, Cc, or Recipients, you can use an SMTP address, alias, or display name to denote a recipient. For example, you can use pilarp@contoso.com, pilarp, or "Pilar Pinilla".
- You can use only prefix wildcard searches; for example, **cat*** or **set***. Suffix searches (***cat**), infix searches (**c*t**), and substring searches (***cat***) are not supported.
- When searching a property, use double quotation marks (" ") if the search value consists of multiple words. For example `subject:budget Q1` returns messages that contain **budget** in the subject line and that contain **Q1** anywhere in the message or in any of the message properties. Using `subject:"budget Q1"` returns all messages that contain **budget Q1** anywhere in the subject line.
- To exclude content marked with a certain property value from your search results, place a minus sign (-) before the name of the property. For example, `-from:"Sara Davis"` excludes any messages sent by Sara Davis.
- You can export items based on message type. For example, to export Skype conversations and chats in Microsoft Teams, use the syntax `kind:im`. To return only email messages, you would use `kind:email`. To return chats, meetings, and calls in Microsoft Teams, use `kind:microsoftteams`.

Configure permissions filtering for Content Search

2/18/2021 • 18 minutes to read • Edit Online

You can use search permissions filtering to let an eDiscovery manager search only a subset of mailboxes and sites in your organization. You can also use permissions filtering to let that same eDiscovery manager search only for mailbox or site content that meets a specific search criteria. For example, you might let an eDiscovery manager search only the mailboxes of users in a specific location or department. You do this by creating a filter that uses a supported recipient filter to limit which mailboxes a specific user or group of users can search. You can also create a filter that specifies what mailbox content a user can search for. This is done by creating a filter that uses a searchable message property. Similarly, you can let an eDiscovery manager search only specific SharePoint sites in your organization. You do this by creating a filter that limits which site can be searched. You can also create a filter that specifies what site content can be searched. This is done by creating a filter that uses a searchable site property.

You can also use search permissions filtering to create logical boundaries (called *compliance boundaries*) within an organization that control the user content locations (such as mailboxes, SharePoint sites, and OneDrive accounts) that specific eDiscovery managers can search. For more information, see Set up compliance boundaries for eDiscovery investigations in Office 365.

Search permissions filtering is supported by the Content Search feature in the Security & Compliance Center. These four cmdlets let you configure and manage search permissions filters:

New-ComplianceSecurityFilter

Get-ComplianceSecurityFilter

Set-ComplianceSecurityFilter

Remove-ComplianceSecurityFilter

Requirements to configure permissions filtering

- To run the compliance security filter cmdlets, you have to be a member of the Organization Management role group in the Security & Compliance Center. For more information, see Permissions in the Security & Compliance Center.
- You have to connect to both Exchange Online and Security & Compliance Center PowerShell to use the compliance security filter cmdlets. This is necessary because these cmdlets require access to mailbox properties, which is why you have to connect to Exchange Online PowerShell. See the steps in the next section.
- See the More information section for additional information about search permissions filters.
- Search permissions filtering is applicable to inactive mailboxes, which means you can use mailbox and mailbox content filtering to limit who can search an inactive mailbox. See the More information section for additional information about permissions filtering and inactive mailboxes.
- Search permissions filtering can't be used to limit who can search public folders in Exchange.
- There is no limit to the number of search permissions filters that can be created in an organization. But search performance will be impacted when there are more than 100 search permissions filters. To keep the number of search permissions filters in your organization as small as possible, create filters that combine rules for Exchange, SharePoint, and OneDrive in a single filter whenever possible.

Connect to Exchange Online and Security & Compliance Center PowerShell in a single session

Before you can successfully run the script in this section, you have to download and install the Exchange Online PowerShell V2 module. For information, see About the Exchange Online PowerShell V2 module.

1. Save the following text to a Windows PowerShell script file by using a filename suffix of **.ps1**. For example, you could save it to a file named **ConnectEXO-SCC.ps1**.

```
Import-Module ExchangeOnlineManagement
$UserCredential = Get-Credential
Connect-ExchangeOnline -Credential $UserCredential -ShowBanner:$false
Connect-IPPSSession -Credential $UserCredential
$Host.UI.RawUI.WindowTitle = $UserCredential.UserName + " (Exchange Online + Compliance Center)"
```

2. On your local computer, open Windows PowerShell, go to the folder where the script that you created in the previous step is located, and then run the script; for example:

```
.\ConnectEXO-SCC.ps1
```

How do you know if this worked? After you run the script, cmdlets from Exchange Online and Security & Compliance PowerShell are imported to your local Windows PowerShell session. If you don't receive any errors, you connected successfully. A quick test is to run an Exchange Online and Security & Compliance Center cmdlet. For example, you can run and **Get-Mailbox** and **Get-ComplianceSearch**.

For troubleshooting PowerShell connection errors, see:

- Connect to Exchange Online PowerShell
- Connect to Security & Compliance Center PowerShell

New-ComplianceSecurityFilter

The **New-ComplianceSecurityFilter** is used to create a search permissions filter. The following table describes the parameters for this cmdlet. All parameters are required to create a compliance security filter.

PARAMETER	DESCRIPTION
Action	The *Action* parameter specifies that type of search action that the filter is applied to. The possible Content Search actions are: **Export:** The filter is applied when exporting search results. **Preview:** The filter is applied when previewing search results. **Purge:** The filter is applied when purging search results. **Search:** The filter is applied when running a search. **All:** The filter is applied to all search actions.
FilterName	The *FilterName* parameter specifies the name of the permissions filter. This name is used to identity a filter when using the **Get-ComplianceSecurityFilter**, **Set-ComplianceSecurityFilter,** and **Remove-ComplianceSecurityFilter** cmdlets.
Filters	The *Filters* parameter specifies the search criteria for the compliance security filter. You can create three different types

PARAMETER	DESCRIPTION

	of filters:
	Mailbox filtering: This type of filter specifies the mailboxes the assigned users (specified by the *Users* parameter) can search. The syntax for this type of filter is **Mailbox_** *MailboxPropertyName*, where *MailboxPropertyName* specifies a mailbox property used to scope the mailboxes that can be searched. For example, the mailbox filter `"Mailbox_CustomAttribute10 -eq 'OttawaUsers'"` would allow the user assigned this filter to search only the mailboxes that have the value "OttawaUsers" in the CustomAttribute10 property.
	Any supported filterable recipient property can be used for the *MailboxPropertyName* property. For a list of supported properties, see Filterable properties for the -RecipientFilter parameter.
	
	Mailbox content filtering: This type of filter is applied on the content that can be searched. It specifies the mailbox content the assigned users can search for. The syntax for this type of filter is **MailboxContent_** *SearchablePropertyName:value*, where *SearchablePropertyName* specifies a Keyword Query Language (KQL) property that can be specified in a Content Search. For example, the mailbox content filter `MailboxContent_recipients:contoso.com` would allow the user assigned this filter to only search for messages sent to recipients in the contoso.com domain.
	For a list of searchable message properties, see Keyword queries and search conditions for Content Search.
	
	Important: A single search filter can't contain a mailbox filter and a mailbox content filter. To combine these in a single filter, you have to use a filters list. But a filter can contain a more complex query of the same type. For example,
	`"Mailbox_CustomAttribute10 -eq 'FTE' -and Mailbox_MemberOfGroup -eq '$($DG.DistinguishedName)'"`
	
	Site and site content filtering: There are two SharePoint and OneDrive for Business site-related filters that you can use to specify what site or site content the assigned users can search:
	
	- **Site_** *SearchableSiteProperty*
	- **SiteContent_** *SearchableSiteProperty*
	
	These two filters are interchangeable. For example,
	`"Site_Path -like 'https://contoso.sharepoint.com/sites/doctors*'"`
	and
	`"SiteContent_Path -like 'https://contoso.sharepoint.com/sites/doctors*'"`
	return the same results. But to help you identify what a filter does, you can use `Site_` to specify site-related properties (such as a site URL) and `SiteContent_` to specify content-related properties (such as document types. For example, the filter
	`"Site_Path -like 'https://contoso.sharepoint.com/sites/doctors*'"`
	would allow the user assigned this filter to only search for content in the https://contoso.sharepoint.com/sites/doctors site collection. The filter
	`"SiteContent_FileExtension -eq 'docx'"` would allow the user assigned this filter to only search for Word documents (Word 2007 and later).

PARAMETER	DESCRIPTION
	For a list of searchable site properties, see Overview of crawled and managed properties in SharePoint. Properties marked with a **Yes** in the **Queryable** column can be used to create a site or site content filter. **Important:** You have to create a search permissions filter to explicitly prevent users from searching content locations in a specific service (such as preventing a user from searching any Exchange mailbox or any SharePoint site). In other words, creating a search permissions filter that allows a user to search all SharePoint sites in the organization doesn't prevent that user from searching mailboxes. For example, to allow SharePoint admins to only search SharePoint sites, you have to create a filter that prevents them from searching mailboxes. Similarly, to allow Exchange admins to only search mailboxes, you have to create a filter that prevents them from searching sites.
Users	The *Users* parameter specifies the users who get this filter applied to their Content Searches. Identify users by their alias or primary SMTP address. You can specify multiple values separated by commas, or you can assign the filter to all users by using the value **All**. You can also use the *Users* parameter to specify a Security & Compliance Center role group. This lets you create a custom role group and then assign that role group a search permissions filter. For example, let's say you have a custom role group for eDiscovery managers for the U.S. subsidiary of a multi-national corporation. You can use the *Users* parameter to specify this role group (by using the Name property of the role group) and then use the *Filter* parameter to allow only mailboxes in the U.S. to be searched. You can't specify distribution groups with this parameter.

Using a filters list to combine filter types

A *filters list* is a filter that includes a mailbox filter and a site filter separated by a comma. Using a filters list is the only supported method for combining different types of filters. In the following example, notice that a comma separates the **Mailbox** and **Site** filters:

```
-Filters "Mailbox_CustomAttribute10 -eq 'OttawaUsers'", "Site_Path -like
'https://contoso.sharepoint.com/sites/doctors*'"
```

When a filter that contains a filters list is processed during the running of a content search, two search permissions filters are created from the filters list: One for each filter that's separated by a comma. So in the previous example, one mailbox search permissions filter and one site search permissions filter would be created.

An alternative to using a filters list would be to create two separate search permissions filters. So in the previous example, you'd create one filter for the mailbox attribute and one filter for the site attribute. In either case, the results are the same. Using a filters list or creating separate search permissions filters is a matter of preference.

Keep the following things in mind about using a filters list:

- You have to use a filters list to create a filter that includes a **Mailbox** filter and a **MailboxContent** filter.
- As previously suggested, you don't have to use a filters list to include a **Site** and a **SiteContent** filter in a single search permissions filter. For example, you can combine **Site** and a **SiteContent** filters using an **-or** operator.

```
-Filters "Site_ComplianceAttribute -eq 'FourthCoffee' -or Site_Path -like
'https://contoso.sharepoint.com/sites/FourthCoffee*'"
```

- Each component of a filters list can contain a complex filter syntax. For example, the mailbox and site filters can contain multiple filters separated by an **-or** operator:

```
-Filters "Mailbox_Department -eq 'CohoWinery' -or Mailbox_CustomAttribute10 -eq 'CohoUsers'",
"Site_ComplianceAttribute -eq 'CohoWinery' -or Site_Path -like
'https://contoso.sharepoint.com/sites/CohoWinery*'"
```

Examples of creating search permissions filters

Here are examples of using the **New-ComplianceSecurityFilter** cmdlet to create a search permissions filter.

This example allows the user annb@contoso.com to perform all Content Search actions only for mailboxes in Canada. This filter contains the three-digit numeric country code for Canada from ISO 3166-1.

```
New-ComplianceSecurityFilter -FilterName CountryFilter  -Users annb@contoso.com -Filters
"Mailbox_CountryCode  -eq '124'" -Action All
```

This example allows the users donh and suzanf to search only the mailboxes that have the value 'Marketing' for the CustomAttribute1 mailbox property.

```
New-ComplianceSecurityFilter -FilterName MarketingFilter  -Users donh,suzanf -Filters
"Mailbox_CustomAttribute1  -eq 'Marketing'" -Action Search
```

This example allows members of the "US Discovery Managers" role group to perform all Content Search actions only on mailboxes in the United States. This filter contains the three-digit numeric country code for the United States from ISO 3166-1.

```
New-ComplianceSecurityFilter -FilterName USDiscoveryManagers  -Users "US Discovery Managers" -Filters
"Mailbox_CountryCode  -eq '840'" -Action All
```

This example allows members of the eDiscovery Manager role group to search only the mailboxes of members of the Ottawa Users distribution group. The Get-DistributionGroup cmdlet in Exchange Online PowerShell is used to find the members of the Ottawa Users group.

```
$DG = Get-DistributionGroup "Ottawa Users"
```

```
New-ComplianceSecurityFilter -FilterName DGFilter  -Users eDiscoveryManager -Filters "Mailbox_MemberOfGroup
-eq '$($DG.DistinguishedName)'" -Action Search
```

This example prevents any user from deleting content from the mailboxes of members of the Executive Team distribution group. The Get-DistributionGroup cmdlet in Exchange Online PowerShell is used to find the members of the Executive Team group.

```
$DG = Get-DistributionGroup "Executive Team"
```

```
New-ComplianceSecurityFilter -FilterName NoExecutivesPreview  -Users All -Filters "Mailbox_MemberOfGroup -ne
'$($DG.DistinguishedName)'" -Action Purge
```

This example allows members of the OneDrive eDiscovery Managers custom role group to only search for content in OneDrive for Business locations in the organization.

```
New-ComplianceSecurityFilter -FilterName OneDriveOnly  -Users "OneDrive eDiscovery Managers" -Filters
"Site_Path -like 'https://contoso-my.sharepoint.com/personal*'" -Action Search
```

NOTE

To restrict users to searching specific sites, use the filter `Site_Path`, as shown in the previous example. Using `Site_Site` will not work.

This example restricts the user to performing all Content Search actions only on email messages sent during the calendar year 2015.

```
New-ComplianceSecurityFilter -FilterName EmailDateRestrictionFilter -Users donh@contoso.com -Filters
"MailboxContent_Received -ge '01-01-2015' -and MailboxContent_Received -le '12-31-2015'" -Action All
```

Similar to the previous example, this example restricts the user to performing all Content Search actions on documents that were last changed sometime in the calendar year 2015.

```
New-ComplianceSecurityFilter -FilterName DocumentDateRestrictionFilter -Users donh@contoso.com -Filters
"SiteContent_LastModifiedTime -ge '01-01-2015' -and SiteContent_LastModifiedTime -le '12-31-2015'" -Action
All
```

This example prevents members of the "OneDrive Discovery Managers" role group from performing content search actions on any mailbox in the organization.

```
New-ComplianceSecurityFilter -FilterName NoEXO -Users "OneDrive Discovery Managers" -Filters "Mailbox_Alias
-notlike '*'"  -Action All
```

This example prevents anyone in the organization from searching for email messages that were sent or received by janets or sarad.

```
New-ComplianceSecurityFilter -FilterName NoSaraJanet -Users All -Filters "MailboxContent_Participants -
notlike 'janets@contoso.onmicrosoft.com' -and MailboxContent_Participants -notlike
'sarad@contoso.onmicrosoft.com'" -Action Search
```

This example uses a filters list to combine mailbox and site filters.

```
New-ComplianceSecurityFilter -FilterName "Coho Winery Security Filter" -Users "Coho Winery eDiscovery
Managers", "Coho Winery Investigators" -Filters "Mailbox_Department -eq 'CohoWinery'",
"Site_ComplianceAttribute -eq 'CohoWinery' -or Site_Path -like
'https://contoso.sharepoint.com/sites/CohoWinery*'" -Action ALL
```

Get-ComplianceSecurityFilter

The **Get-ComplianceSecurityFilter** is used to return a list of search permissions filters. Use the *FilterName*

parameter to return information for a specific search filter.

Set-ComplianceSecurityFilter

The **Set-ComplianceSecurityFilter** is used to modify an existing search permissions filter. The only required parameter is *FilterName*.

PARAMETER	DESCRIPTION
Action	The *Action* parameter specifies that type of search action that the filter is applied to. The possible Content Search actions are: **Export:** The filter is applied when exporting search results. **Preview:** The filter is applied when previewing search results. **Purge:** The filter is applied when purging search results. **Search:** The filter is applied when running a search. **All:** The filter is applied to all search actions.
FilterName	The *FilterName* parameter specifies the name of the permissions filter.
Filters	The *Filters* parameter specifies the search criteria for the compliance security filter. You can create two different types of filters: **Mailbox filtering:** This type of filter specifies the mailboxes the assigned users (specified by the *Users* parameter) can search. The syntax for this type of filter is **Mailbox_** *MailboxPropertyName*, where *MailboxPropertyName* specifies a mailbox property used to scope the mailboxes that can be searched. For example, the mailbox filter `"Mailbox_CustomAttribute10 -eq 'OttawaUsers'"` would allow the user assigned this filter to search only the mailboxes that have the value "OttawaUsers" in the CustomAttribute10 property. Any supported filterable recipient property can be used for the *MailboxPropertyName* property. For a list of supported properties, see Filterable properties for the -RecipientFilter parameter. **Mailbox content filtering:** This type of filter is applied on the content that can be searched. It specifies the mailbox content the assigned users can search for. The syntax for this type of filter is **MailboxContent_** *SearchablePropertyName:value*, where *SearchablePropertyName* specifies a Keyword Query Language (KQL) property that can be specified in a Content Search. For example, the mailbox content filter `MailboxContent_recipients:contoso.com` would allow the user assigned this filter to only search for messages sent to recipients in the contoso.com domain. For a list of searchable message properties, see Keyword queries for Content Search. **Site and site content filtering:** There are two SharePoint and OneDrive for Business site-related filters that you can use to specify what site or site content the assigned users can search: - **Site_** *SearchableSiteProperty* - **SiteContent_** *SearchableSiteProperty*

PARAMETER	DESCRIPTION
	- SiteContent_*SearchableSiteProperty* These two filters are interchangeable. For example, `"Site_Path -like 'https://contoso.spoppe.com/sites/doctors*'"` and `"SiteContent_Path -like 'https://contoso.spoppe.com/sites/doctors*'"` returns the same results. But to help you identify what a filter does, you can use `Site_` to specify site-related properties (such as a site URL) and `SiteContent_` to specify content-related properties (such as document types. For example, the filter `"Site_Path -like 'https://contoso.spoppe.com/sites/doctors*'"` would allow the user assigned this filter to only search for content in the https://contoso.spoppe.com/sites/doctors site collection. The filter `"SiteContent_FileExtension -eq 'docx'"` would allow the user assigned this filter to only search for Word documents (Word 2007 and later). For a list of searchable site properties, see Overview of crawled and managed properties in SharePoint. Properties marked with a **Yes** in the **Queryable** column can be used to create a site or site content filter.
Users	The *Users* parameter specifies the users who get this filter applied to their Content Searches. Because this is a multi-value property, specifying a user or group of users with this parameter overwrite the existing list of users. See the following examples for the syntax to add and remove selected users. You can also use the *Users* parameter to specify a Security & Compliance Center role group. This lets you create a custom role group and then assign that role group a search permissions filter. For example, let's say you have a custom role group for eDiscovery managers for the U.S. subsidiary of a multi-national corporation. You can use the *Users* parameter to specify this role group (by using the Name property of the role group) and then use the *Filter* parameter to allow only mailboxes in the U.S. to be searched. You can't specify distribution groups with this parameter.

Examples of changing search permissions filters

These examples show how to use the **Get-ComplianceSecurityFilter** and **Set-ComplianceSecurityFilter** cmdlets to add or remove a user to the existing list of users that the filter is assigned to. When you add or remove users from a filter, specify the user by using their SMTP address.

This example adds a user to the filter.

```
$filterusers = Get-ComplianceSecurityFilter -FilterName OttawaUsersFilter
```

```
$filterusers.users.add("pilarp@contoso.com")
```

```
Set-ComplianceSecurityFilter -FilterName OttawaUsersFilter -Users $filterusers.users
```

This example removes a user from the filter.

```
$filterusers = Get-ComplianceSecurityFilter -FilterName OttawaUsersFilter
```

```
$filterusers.users.remove("annb@contoso.com")
```

```
Set-ComplianceSecurityFilter -FilterName OttawaUsersFilter -Users $filterusers.users
```

Remove-ComplianceSecurityFilter

The **Remove-ComplianceSecurityFilter** is used to delete a search filter. Use the *FilterName* parameter to specify the filter you want to delete.

More information

- **How does search permissions filtering work?** The permissions filter is added to the search query when a Content Search is run. The permissions filter is joined to the search query by the **AND** Boolean operator. For example, you have a permissions filter that allows Bob to perform all search actions on the mailboxes of members of the Workers distribution group. Then Bob runs a Content Search on all mailboxes in the organization with the search query `sender:jerry@adatum.com`. Because the permissions filter and the search query are logically combined by an **AND** operator, the search returns any message sent by jerry@adatum.com to any member of the Workers distribution group.

- **What happens if you have multiple search permissions filters?** 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. Let's take the previous example, where a search filter allows Bob to search only the mailboxes of the members of the Workers distribution group. Then we create another filter that prevents Bob from searching Phil's mailbox ("Mailbox_Alias -ne 'Phil'"). And let's also assume that Phil is a member of the Workers group. When Bob runs a Content Search (from the previous example) on all mailboxes in the organization, search results are returned for Phil's mailbox even though you applied filter to prevent Bob from searching Phil's mailbox. This is because the first filter, which allows Bob to search the Workers group, is true. And because Phil is a member of the Workers group, Bob can search Phil's mailbox.

- **Does search permissions filtering work for inactive mailboxes?** Yes, you can use mailbox and mailbox content filters to limit who can search inactive mailboxes in your organization. Like a regular mailbox, an inactive mailbox has to be configured with the recipient property that's used to create a permissions filter. If necessary, you can use the **Get-Mailbox -InactiveMailboxOnly** command to display the properties of inactive mailboxes. For more information, see Create and manage inactive mailboxes in Office 365.

- **Does search permissions filtering work for public folders?** No. As previously explained, search permissions filtering can't be used to limit who can search public folders in Exchange. For example, items in public folder locations can't be excluded from the search results by a permissions filter.

- **Does allowing a user to search all content locations in a specific service also prevent them from searching content locations in a different service?** No. As previously explained, you have to create a search permissions filter to explicitly prevent users from searching content locations in a specific

service (such as preventing a user from searching any Exchange mailbox or any SharePoint site). In other words, creating a search permissions filter that allows a user to search all SharePoint sites in the organization doesn't prevent that user from searching mailboxes. For example, to allow SharePoint admins to only search SharePoint sites, you have to create a filter that prevents them from searching mailboxes. Similarly, to allow Exchange admins to only search mailboxes, you have to create a filter that prevents them from searching sites.

Set up compliance boundaries for eDiscovery investigations

2/18/2021 • 20 minutes to read • Edit Online

The guidance in this article can be applied when using either Core eDiscovery or Advanced eDiscovery to manage investigations.

Compliance boundaries create logical boundaries within an organization that control the user content locations (such as mailboxes, OneDrive accounts, and SharePoint sites) that eDiscovery managers can search. Also, compliance boundaries control who can access eDiscovery cases used to manage the legal, human resources, or other investigations within your organization. The need for compliance boundaries is often necessary for multi-national corporations that have to respect geographical boarders and regulations and for governments, which are often divided into different agencies. In Microsoft 365, compliance boundaries help you meet these requirements when performing content searches and managing investigations with eDiscovery cases.

We use the example in the following illustration to explain how compliance boundaries work.

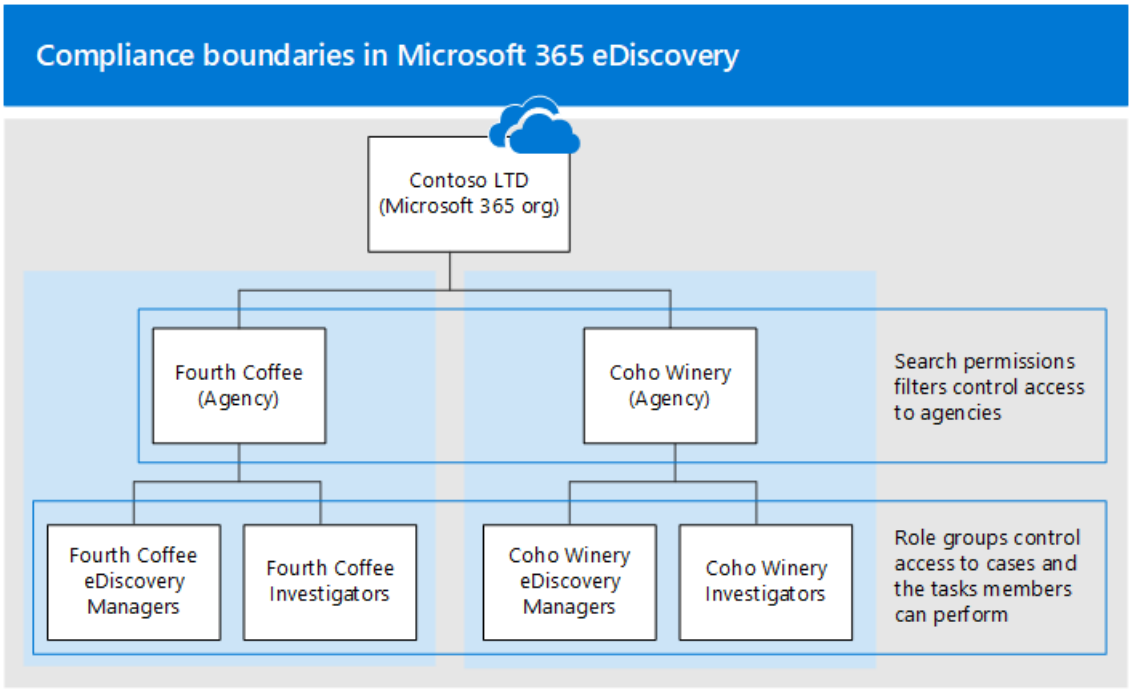

In this example, Contoso LTD is an organization that consists of two subsidiaries, Fourth Coffee and Coho Winery. The business requires that eDiscovery mangers and investigators can only search the Exchange mailboxes, OneDrive accounts, and SharePoint sites in their agency. Also, eDiscovery managers and investigators can only see eDiscovery cases in their agency, and they can only access the cases that they're a member of. Additionally in this scenario, investigators cannot place content locations on hold or export content from a case. Here's how compliance boundaries meet these requirements.

- The search permissions filtering functionality in Content search controls the content locations that eDiscovery managers and investigators can search. This means eDiscovery managers and investigators in the Fourth Coffee agency can only search content locations in the Fourth Coffee subsidiary. The same restriction applies to the Coho Winery subsidiary.
- Role groups provide the following functions for compliance boundaries:
    - Control who can see the eDiscovery cases in the Security & Compliance Center. This means that

eDiscovery managers and investigators can only see the eDiscovery cases in their agency.

- Control who can assign members to an eDiscovery case. This means eDiscovery managers and investigators can only assign members to cases that they themselves are a member of.
- Control the eDiscovery-related tasks that members can perform by adding or removing roles that assign specific permissions.

Here's the process for setting up compliance boundaries:

Step 1: Identify a user attribute to define your agencies

Step 2: File a request with Microsoft Support to synchronize the user attribute to OneDrive accounts

Step 3: Create a role group for each agency

Step 4: Create a search permissions filter to enforce the compliance boundary

Step 5: Create an eDiscovery case for an intra-agency investigations

Before you set up compliance boundaries

You have to meet the following prerequisites before the Azure Active Directory (Azure AD) attribute that you identity (in Step 1) can be successfully synched to a user's OneDrive account (in Step 2):

- Users must be assigned an Exchange Online license and a SharePoint Online license.
- User mailboxes must be at least 10 MB in size. If a user's mailbox is less than 10 MB, the attribute used to define your agencies won't be synched to the user's OneDrive account.
- Compliance boundaries and the attributes used to create search permissions filters require that Azure Active Directory (Azure AD) attributes are synchronized to user mailboxes. To verify that the attributes that you want to use have been synchronized, run the Get-User cmdlet in Exchange Online PowerShell. The output of this cmdlet displays the Azure AD attributes synchronized to Exchange Online.

Step 1: Identify a user attribute to define your agencies

The first step is to choose an Azure AD attribute to use that will define your agencies. This attribute is used to create the search permissions filter that limits an eDiscovery manager to search only the content locations of users who are assigned a specific value for this attribute. For example, let's say Contoso decides to use the **Department** attribute. The value for this attribute for users in the Fourth Coffee subsidiary would be `FourthCoffee` and the value for users in Coho Winery subsidiary would be `CohoWinery`. In Step 4, you use this `attribute:value` pair (for example, *Department:FourthCoffee*) to limit the user content locations that eDiscovery managers can search.

Here's a list of Azure AD user attributes that you can use for compliance boundaries:

- Company
- CustomAttribute1 - CustomAttribute15
- Department
- Office
- C (Two-letter country code) *

> **NOTE**
>
> * This attribute maps to the CountryOrRegion property that is returned by running the **Get-User** cmdlet in Exchange Online PowerShell. The cmdlet returns the localized country name, which is translated from the two-letter country code. For more information, see the CountryOrRegion parameter description in the Set-User cmdlet reference article.

Although more user attributes are available, particularly for Exchange mailboxes, the attributes listed above are the only ones currently supported by OneDrive.

Step 2: File a request with Microsoft Support to synchronize the user attribute to OneDrive accounts

The next step is to file a request with Microsoft Support to synchronize the Azure AD attribute that you chose in Step 1 to all OneDrive accounts in your organization. After this synchronization occurs, the attribute (and its value) that you chose in Step 1 will be mapped to a hidden managed property named `ComplianceAttribute`. You use this attribute to create the search permissions filter for OneDrive in Step 4.

Include the following information when you submit the request to Microsoft support:

- The default domain name of your organization
- The name of the Azure AD attribute (from Step 1)
- The following title or description of the purpose of the support request: "Enable OneDrive for Business Synchronization with Azure AD for Compliance Security Filters". This helps route the request to the eDiscovery engineering team who implements the request.

After the engineering change is made and the attribute is synchronized to OneDrive, Microsoft Support will send you the build number that the change was made in and an estimated deployment date. The deployment process usually takes 4–6 weeks after you submit the support request.

> **IMPORTANT**
>
> You can complete Step 3 through Step 5 before this attribute change is deployed. But running content searches won't return documents from OneDrive accounts that are specified in a search permissions filter until after the attribute synch is deployed.

Step 3: Create a role group for each agency

The next step is to create the role groups in the Security & Compliance Center that will align with your agencies. We recommend that you create a role group by copying the built-in eDiscovery Managers group, adding the appropriate members, and removing roles that may not be applicable to your needs. For more information about eDiscovery-related roles, see Assign eDiscovery permissions in the Office 365 Security & Compliance Center.

To create the role groups, go to the **Permissions** page in the Security & Compliance Center and create a role group for each team in each agency that will use compliance boundaries and eDiscovery cases to manage investigations.

Using the Contoso compliance boundaries scenario, four role groups need to be created and the appropriate members added to each one.

- Fourth Coffee eDiscovery Managers

- Fourth Coffee Investigators
- Coho Winery eDiscovery Managers
- Coho Winery Investigators

To meet the requirements of the Contoso compliance boundaries scenario, you would also remove the **Hold** and **Export** roles from the investigators role groups to prevent investigators from placing holds on content locations and exporting content from a case.

Step 4: Create a search permissions filter to enforce the compliance boundary

After you've created role groups for each agency, the next step is to create the search permissions filters that associate each role group to its specific agency and defines the compliance boundary itself. You need to create one search permissions filter for each agency. For more information about creating security permissions filters, see Configure permissions filtering for Content Search.

Here's the syntax that's used to create a search permissions filter used for compliance boundaries.

```
New-ComplianceSecurityFilter -FilterName <name of filter> -Users <role groups> -Filters
"Mailbox_<ComplianceAttribute>  -eq '<AttributeVale> '", "Site_<ComplianceAttribute>  -eq '<AttributeValue>'
-or Site_Path -like '<SharePointURL>*'" -Action <Action >
```

Here's a description of each parameter in the command:

- `FilterName`: Specifies the name of the filter. Use a name that describes or identifies the agency that the filter is used in.
- `Users`: Specifies the users or groups who get this filter applied to the Content Search actions they perform. For compliance boundaries, this parameter specifies the role groups (that you created in Step 3) in the agency that you're creating the filter for. Note this is a multi-value parameter so you can include one or more role groups, separated by commas.
- `Filters`: Specifies the search criteria for the filter. For the compliance boundaries, you define the following filters. Each one applies to a content location.
  - `Mailbox`: Specifies the mailboxes that the role groups defined in the `Users` parameter can search. For compliance boundaries, *ComplianceAttribute* is the same attribute that you identified in Step 1 and *AttributeValue* specifies the agency. This filter allows members of the role group to search only the mailboxes in a specific agency; for example, `"Mailbox_Department -eq 'FourthCoffee'"`.
  - `Site`: Specifies the OneDrive accounts that the role groups defined in the `Users` parameter can search. For the OneDrive filter, use the actual string `ComplianceAttribute`. This maps to the same attribute that you identified in Step 1 and that's synchronized to OneDrive accounts as a result of the support request that you submitted in Step 2; *AttributeValue* specifies the agency. This filter allows members of the role group to search only the OneDrive accounts in a specific agency; for example, `"Site_ComplianceAttribute -eq 'FourthCoffee'"`.
  - `Site_Path`: Specifies the SharePoint sites that the role groups defined in the `Users` parameter can search. The *SharePointURL* specifies the sites in the agency that members of the role group can search. For example, `"Site_Path -like 'https://contoso.sharepoint.com/sites/FourthCoffee*'"`. Notice the `Site` and `Site_Path` filters are connected by an **-or** operator.

NOTE

The syntax for the `Filters` parameter includes a *filters list*. A filters list is a filter that includes a mailbox filter and a site filter separated by a comma. In the previous example, notice that a comma separates **Mailbox_ComplianceAttribute** and **Site_ComplianceAttribute**:

```
-Filters "Mailbox_<ComplianceAttribute> -eq '<AttributeVale> '", "Site_ComplianceAttribute -eq '<AttributeValue>' -or Site_Path -like '<SharePointURL>*'"
```

. When this filter is processed during the running of a content search, two search permissions filters are created from the filters list: one mailbox filter and one site filter. An alternative to using a filters list would be to create two separate search permissions filters for each agency: one search permissions filter for the mailbox attribute and one filter for the site attributes. In either case, the results will be the same. Using a filters list or creating separate search permissions filters is a matter of preference.

- `Action`: Specifies the type of Compliance Search action that the filter is applied to. For example, `-Action Search` would only apply the filter when members of the role group defined in the `Users` parameter run a content search. In this case, the filter wouldn't be applied when exporting search results. For compliance boundaries, use `-Action All` so the filter applies to all search actions.

  For a list of the Content Search actions, see the "New-ComplianceSecurityFilter" section in Configure permissions filtering for Content Search.

Here are examples of the two search permissions filters that would be created to support the Contoso compliance boundaries scenario. Both of these examples include a comma-separated filters list, in which the mailbox and site filters are included in the same search permissions filter and are separated by a comma.

Fourth Coffee

```
New-ComplianceSecurityFilter -FilterName "Fourth Coffee Security Filter" -Users "Fourth Coffee eDiscovery Managers", "Fourth Coffee Investigators" -Filters "Mailbox_Department -eq 'FourthCoffee'", "Site_ComplianceAttribute -eq 'FourthCoffee' -or Site_Path -like 'https://contoso.sharepoint.com/sites/FourthCoffee*'" -Action ALL
```

Coho Winery

```
New-ComplianceSecurityFilter -FilterName "Coho Winery Security Filter" -Users "Coho Winery eDiscovery Managers", "Coho Winery Investigators" -Filters "Mailbox_Department -eq 'CohoWinery'", "Site_ComplianceAttribute -eq 'CohoWinery' -or Site_Path -like 'https://contoso.sharepoint.com/sites/CohoWinery*'" -Action ALL
```

Step 5: Create an eDiscovery case for intra-agency investigations

The final step is to create a Core eDiscovery case or Advanced eDiscovery case in the Microsoft 365 compliance center and then add the role group that you created in Step 3 as a member of the case. This results in two important characteristics of using compliance boundaries:

- Only members of the role group added to the case will be able to see and access the case in the Security & Compliance Center. For example, if the Fourth Coffee Investigators role group is the only member of a case, then members of the Fourth Coffee eDiscovery Managers role group (or members of any other role group) won't be able to see or access the case.
- When a member of the role group assigned to a case runs a search associated with the case, they will only be able to search the content locations within their agency (which is defined by the search permissions filter that you created in Step 4.)

To create a case and assign members:

1. Go to the **Core eDiscovery** or **Advanced eDiscovery** page in the Microsoft 365 compliance center and create a case.
2. In the list of cases, click the name of the case you created.
3. In the **Manage this case** flyout page, under **Manage role groups**, click **+ Add**.

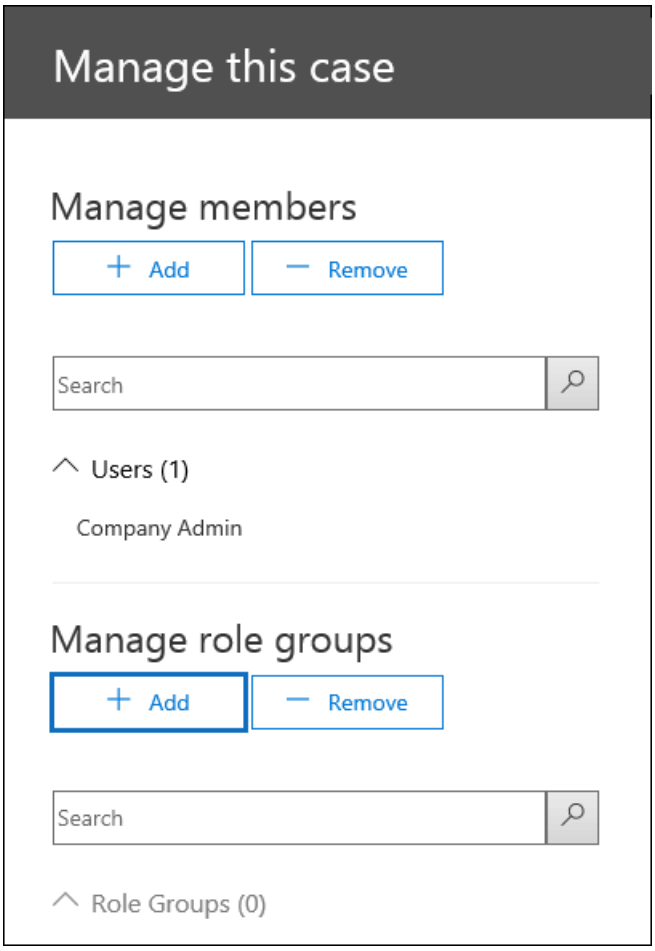

4. In the list of role groups, select one of the role groups that you created in Step 3, and click **Add**.
5. Click **Save** on the **Manage this case** flyout to save the change.

> **NOTE**

When adding a role group to a case, you can only add the role groups that you are a member of.

Searching and exporting content in Multi-Geo environments

Search permissions filters also let you control where content is routed for export and which datacenter can be searched when searching content locations in a SharePoint Multi-Geo environment.

- **Export search results:** You can export the search results from Exchange mailboxes, SharePoint sites, and OneDrive accounts from a specific datacenter. This means that you can specify the datacenter location that search results will be exported from.

  Use the **Region** parameter for **New-ComplianceSecurityFilter** or **Set-ComplianceSecurityFilter** cmdlets to create or change which datacenter the export will be routed through.

PARAMETER VALUE	DATACENTER LOCATION
NAM	North American (datacenters are in the US)
EUR	Europe
APC	Asia Pacific
CAN	Canada

- **Route content searches:** You can route the content searches of SharePoint sites and OneDrive accounts to a satellite datacenter. This means you can specify the datacenter location where searches will be run.

  Use one of the following values for the **Region** parameter to control the datacenter location that searches will run in when searching SharePoint sites and OneDrive accounts.

PARAMETER VALUE	DATACENTER ROUTING LOCATIONS FOR SHAREPOINT
NAM	US
EUR	Europe
APC	Asia Pacific
CAN	US
AUS	Asia Pacific
KOR	The organization's default datacenter
GBR	Europe
JPN	Asia Pacific
IND	Asia Pacific
LAM	US
NOR	Europe
BRA	North American datacenters

  If you don't specify the **Region** parameter for a search permissions filter, the organization's primary SharePoint region will be searched. Search results are exported to the closest datacenter.

  To simplify the concept, the **Region** parameter controls the datacenter that is used to search for content in SharePoint and OneDrive. This doesn't apply to searching for content in Exchange because Exchange content searches aren't bound by the geographic location of datacenters. Also, the same **Region** parameter value may also dictate the datacenter that exports are routed through. This is often necessary to control the movement of data across geographic boarders.

> **NOTE**
>
> If you're using Advanced eDiscovery, the **Region** parameter doesn't control the region that data is exported from. Data is exported from the organization's primary datacenter. Also, searching for content in SharePoint and OneDrive isn't bound by the geographic location of datacenters. All datacenters are searched. For more information about Advanced eDiscovery, see Overview of the Advanced eDiscovery solution in Microsoft 365.

Here are examples of using the **Region** parameter when creating search permission filters for compliance boundaries. This assumes that the Fourth Coffee subsidiary is located in North America and that Coho Winery is in Europe.

```
New-ComplianceSecurityFilter -FilterName "Fourth Coffee Security Filter" -Users "Fourth Coffee eDiscovery Managers", "Fourth Coffee Investigators" -Filters "Mailbox_Department -eq 'FourthCoffee'", "Site_Department -eq 'FourthCoffee' -or Site_Path -like 'https://contoso.sharepoint.com/sites/FourthCoffee*'" -Action ALL -Region NAM
```

```
New-ComplianceSecurityFilter -FilterName "Coho Winery Security Filter" -Users "Coho Winery eDiscovery Managers", "Coho Winery Investigators" -Filters "Mailbox_Department -eq 'CohoWinery'", "Site_Department -eq 'CohoWinery' -or Site_Path -like 'https://contoso.sharepoint.com/sites/CohoWinery*'" -Action ALL -Region EUR
```

Keep the following things in mind when searching and exporting content in multi-geo environments.

- The **Region** parameter doesn't control searches of Exchange mailboxes. All datacenters will be searched when you search mailboxes. To limit the scope of which Exchange mailboxes are searched, use the **Filters** parameter when creating or changing a search permissions filter.
- If it's necessary for an eDiscovery Manager to search across multiple SharePoint regions, you need to create a different user account for that eDiscovery manager to use in the search permissions filter to specify the region where the SharePoint sites or OneDrive accounts are located. For more information about setting this up, see the "Searching for content in a SharePoint Multi-Geo environment" section in Content Search.
- When searching for content in SharePoint and OneDrive, the **Region** parameter directs searches to either the primary or satellite location where the eDiscovery manager will conduct eDiscovery investigations. If an eDiscovery manager searches SharePoint and OneDrive sites outside of the region that's specified in the search permissions filter, no search results are returned.
- When exporting search results, content from all content locations (including Exchange, Skype for Business, SharePoint, OneDrive, and other services that you can search by using the Content Search tool) are uploaded to the Azure Storage location in the datacenter that's specified by the **Region** parameter. This helps organizations stay within compliance by not allowing content to be exported across controlled borders. If no region is specified in the search permissions filter, content is uploaded to the organization's primary datacenter.
- You can edit an existing search permissions filter to add or change the region by running the following command:

  ```
  Set-ComplianceSecurityFilter -FilterName <Filter name>  -Region <Region>
  ```

Using compliance boundaries for SharePoint hub sites

SharePoint hub sites often align with the same geographical or agency boundaries that eDiscovery compliance boundaries follow. That means you can use the site ID property of the hub site to create a compliance boundary.

To do this, use the Get-SPOHubSite cmdlet in SharePoint Online PowerShell to obtain the SiteId for the hub site and then use this value for the department ID property to create a search permissions filter.

Use the following syntax to create a search permissions filter for a SharePoint hub site:

```
New-ComplianceSecurityFilter -FilterName <Filter Name> -Users <User or Group> -Filters "Site_Departmentid -eq '{SiteId of hub site}'" -Action ALL
```

Here's an example of creating a search permissions filter for a hub site for the Coho Winery agency:

```
New-ComplianceSecurityFilter -FilterName "Coho Winery Hub Site Security Filter" -Users "Coho Winery eDiscovery Managers", "Coho Winery Investigators" -Filters "Site_Departmentid -eq '44252d09-62c4-4913-9eb0-a2a8b8d7f863'" -Action ALL
```

Compliance boundary limitations

Keep the following limitations in mind when managing eDiscovery cases and investigations that use of compliance boundaries.

- When creating and running a search, you can select content locations that are outside of your agency. However, because of the search permissions filter, content from those locations isn't included in the search results.
- Compliance boundaries don't apply to holds in eDiscovery cases. That means an eDiscovery manager in one agency can place a user in a different agency on hold. However, the compliance boundary will be enforced if the eDiscovery manager searches the content locations of the user who was placed on hold. That means the eDiscovery manager won't be able search the user's content locations, even though they were able to place the user on hold.

  Also, hold statistics will only apply to content locations in the agency.
- Search permissions filters aren't applied to Exchange public folders.

More information

- If a mailbox is de-licensed or soft-deleted, Azure AD attributes are no longer synchronized to the mailbox. If a hold was placed on the mailbox when it was deleted, the content preserved in the mailbox is still subject to a compliance boundary or search permissions filter based on the last time the Azure AD attributes were synchronized before the mailbox was deleted.

  Additionally, the synchronization between the user's mailbox and OneDrive account will cease if the mailbox is de-licensed or soft-deleted. The last stamped value of the compliance attribute for the OneDrive account will remain in effect.
- The compliance attribute is synchronized from a user's Exchange mailbox to their OneDrive account every seven days. As previously stated, this synchronization only occurs when the user is assigned both an Exchange Online and SharePoint Online license and the user's mailbox is at least 10 MB.
- If compliance boundaries and search permissions filters are implemented for both a user's mailbox and OneDrive account, then we recommend that you don't delete a user's mailbox and not their OneDrive account. In other words, if you delete a user's mailbox, you should also remove the user's OneDrive account.
- There are situations (such as a returning employee) where a user might have two or more OneDrive accounts. In these cases, only the primary OneDrive account associated with the user in Azure AD will be synchronized.

- Compliance boundaries and search permissions filters depend on attributes being stamped on content in Exchange, OneDrive, and SharePoint and the subsequent indexing of this stamped content.
- We don't recommend using exclusion filters (such as using `-not()` in a search permissions filter) for a content-based compliance boundary. Using an exclusion filter can have unexpected results if content with recently updated attributes hasn't been indexed.

Frequently asked questions

Who can create and manage search permissions filters (using New-ComplianceSecurityFilter and Set-ComplianceSecurityFilter cmdlets)?

To create, view, and modify search permissions filters, you have to be a member of the Organization Management role group in the Security & Compliance Center.

If an eDiscovery manager is assigned to more than one role group that spans multiple agencies, how do they search for content in one agency or the other?

The eDiscovery manager can add parameters to their search query that restrict the search to a specific agency. For example, if an organization has specified the **CustomAttribute10** property to differentiate agencies, they can append the following to their search query to search mailboxes and OneDrive accounts in a specific agency: `CustomAttribute10:<value> AND Site_ComplianceAttribute:<value>`.

What happens if the value of the attribute that's used as the compliance attribute in a search permissions filter is changed?

It takes up to three days for a search permissions filter to enforce the compliance boundary if the value of the attribute that's used in the filter is changed. For example, in the Contoso scenario let's say that a user in the Fourth Coffee agency is transferred to the Coho Winery agency. As a result, the value of the **Department** attribute on the user object is changed from *FourthCoffee* to *CohoWinery*. In this situation, Fourth Coffee eDiscovery and investors will get search results for that user for up three days after the attribute is changed. Similarly, it takes up to three days before Coho Winery eDiscovery managers and investigators get search results for the user.

Can an eDiscovery manager see content from two separate compliance boundaries?

Yes, this can be done when searching Exchange mailboxes by adding the eDiscovery manager to role groups that have visibility to both agencies. However when searching SharePoint sites and OneDrive accounts, an eDiscovery manager can search for content in different compliance boundaries only if the agencies are in the same region or geo location. **Note:** This limitation for sites doesn't apply in Advanced eDiscovery because searching for content in SharePoint and OneDrive isn't bound by geographic location.

Do search permissions filters work for eDiscovery case holds, Microsoft 365 retention policies, or DLP?

No, not at this time.

If I specify a region to control where content is exported, but I don't have a SharePoint organization in that region, can I still search SharePoint?

If the region specified in the search permissions filter doesn't exist in your organization, the default region will be searched.

What is the maximum number of search permissions filters that can be created in an organization?

There is no limit to the number of search permissions filters that can be created in an organization. However, search performance will be impacted when there are more than 100 search permissions filters. To keep the

number of search permissions filters in your organization as small as possible, create filters that combine rules for Exchange, SharePoint, and OneDrive into a single search permissions filter whenever possible.

Use a script to add users to a hold in a Core eDiscovery case

2/18/2021 • 12 minutes to read • Edit Online

Security & Compliance Center PowerShell provides cmdlets that let you automate time-consuming tasks related to creating and managing eDiscovery cases. Currently, using the Core eDiscovery case in the Security & Compliance Center to place a large number of custodian content locations on hold takes time and preparation. For example, before you create a hold, you have to collect the URL for each OneDrive for Business site that you want to place on hold. Then for each user you want to place on hold, you have to add their mailbox and their OneDrive for Business site to the hold. You can use the script in this article to automate this process.

The script prompts you for the name of your organization's My Site domain (for example, `contoso` in the URL https://contoso-my.sharepoint.com), the name of an existing eDiscovery case, the name of the new hold that associated with the case, a list of email addresses of the users you want to put on hold, and a search query to use if you want to create a query-based hold. The script then gets the URL for the OneDrive for Business site for each user in the list, creates the new hold, and then adds the mailbox and OneDrive for Business site for each user in the list to the hold. The script also generates log files that contain information about the new hold.

Here are the steps to make this happen:

Before you add users to a hold

- You have to be a member of the eDiscovery Manager role group in the Security & Compliance Center and a SharePoint Online administrator to run the script in Step 3. For more information, see Assign eDiscovery permissions in the Office 365 Security & Compliance Center.
- A maximum of 1,000 mailboxes and 100 sites can be added to a hold that's associated with an eDiscovery case in the Security & Compliance Center. Assuming that every user that you want to place on hold has a OneDrive for Business site, you can add a maximum of 100 users to a hold using the script in this article.
- Be sure to save the list of users that you create in Step 2 and the script in Step 3 to the same folder. That will make it easier to run the script.
- The script adds the list of users to a new hold that is associated with an existing case. Be sure the case that you want to associate the hold with is created before you run the script.
- The script in this article supports modern authentication when connecting to Security & Compliance Center PowerShell and SharePoint Online Management Shell. You can use the script as-is if you are a Microsoft 365 or a Microsoft 365 GCC organization. If you are an Office 365 Germany organization, a Microsoft 365 GCC High organization, or a Microsoft 365 DoD organization, you will have to edit the script to successfully run it. Specifically, you have to edit the line `Connect-IPPSSession` and use the *ConnectionUri* and *AzureADAuthorizationEndpointUri* parameters (and the appropriate values for your organization type) to connect to Security & Compliance Center PowerShell. For more information, see the examples in Connect to Security & Compliance Center PowerShell.

- The script automatically disconnects from Security & Compliance Center PowerShell and SharePoint Online Management Shell.
- The script includes minimal error handling. Its primary purpose is to quickly and easily place the mailbox and OneDrive for Business site of each user on hold.
- The sample scripts provided in this topic aren't supported under any Microsoft standard support program or service. The sample scripts are provided AS IS without warranty of any kind. Microsoft further disclaims all implied warranties including, without limitation, any implied warranties of merchantability or of fitness for a particular purpose. The entire risk arising out of the use or performance of the sample scripts and documentation remains with you. In no event shall Microsoft, its authors, or anyone else involved in the creation, production, or delivery of the scripts be liable for any damages whatsoever (including, without limitation, damages for loss of business profits, business interruption, loss of business information, or other pecuniary loss) arising out of the use of or inability to use the sample scripts or documentation, even if Microsoft has been advised of the possibility of such damages.

Step 1: Install the SharePoint Online Management Shell

The first step is to install the SharePoint Online Management Shell if it's not already installed on your local computer. You don't have to use the shell in this procedure, but you have to install it because it contains pre-requisites required by the script that you run in Step 3. These prerequisites allow the script to communicate with SharePoint Online to get the URLs for the OneDrive for Business sites.

Go to Set up the SharePoint Online Management Shell Windows PowerShell environment and perform Step 1 and Step 2 to install the SharePoint Online Management Shell on your local computer.

Step 2: Generate a list of users

The script in Step 3 will create a hold that's associated with an eDiscovery case, and the add the mailboxes and OneDrive for Business sites of a list of users to the hold. You can just type the email addresses in a text file, or you can run a command in Windows PowerShell to get a list of email addresses and save them to a file (located in same folder that you'll save the script to in Step 3).

Here's a PowerShell command (that you run by using remote PowerShell connected to your Exchange Online organization) to get a list of email addresses for all users in your organization and save it to a text file named HoldUsers.txt.

```
Get-Mailbox -ResultSize unlimited -Filter { RecipientTypeDetails -eq 'UserMailbox'} | Select-Object
PrimarySmtpAddress > HoldUsers.txt
```

After you run this command, open the text file and remove the header that contains the property name, `PrimarySmtpAddress`. Then remove all email addresses except the ones for the users that you want to add to the hold that you'll create in Step 3. Make sure there are no blank rows before or after the list of email addresses.

Step 3: Run the script to create a hold and add users

When you run the script in this step, it will prompt you for the following information. Be sure to have this information ready before you run the script.

- **Your user credentials:** The script will use your credentials to connect to Security & Compliance Center with PowerShell. It will also use these credentials to access SharePoint Online to get the OneDrive for Business URLs for the list of users.
- **Name of your SharePoint domain:** The script prompts you to enter this name so it can connect to the

SharePoint admin center. It also uses the domain name for the OneDrive URLs in your organization. For example, if the URL for your admin center is `https://contoso-admin.sharepoint.com` and the URL for OneDrive is `https://contoso-my.sharepoint.com`, then you would enter `contoso` when the script prompts you for your domain name.

- **Name of the case:** The name of an existing case. The script will create a new hold that is associated with this case.
- **Name of the hold:** The name of the hold the script will create and associate with the specified case.
- **Search query for a query-based hold:** You can create a query-based hold so that only the content that meets the specified search criteria is placed on hold. To place all content on hold, just press **Enter** when you're prompted for a search query.
- **Turning on the hold or not:** You can have the script turn on the hold after it's created or you can have the script create the hold without enabling it. If you don't have the script turn on the hold, you can turn it on later in the Security & Compliance Center or by running the following PowerShell commands:

  ```
  Set-CaseHoldPolicy -Identity <name of the hold> -Enabled $true
  ```

  ```
  Set-CaseHoldRule -Identity <name of the hold> -Disabled $false
  ```

- **Name of the text file with the list of users** - The name of the text file from Step 2 that contains the list of users to add to the hold. If this file is located in the same folder as the script, just type the name of the file (for example, HoldUsers.txt). If the text file is in another folder, type the full pathname of the file.

After you've collected the information that the script will prompt you for, the final step is to run the script to create the new hold and add users to it.

1. Save the following text to a Windows PowerShell script file by using a filename suffix of `.ps1`. For example, `AddUsersToHold.ps1`.

```
#script begin
" "
write-host "************************************************"
write-host "   Security & Compliance Center PowerShell  " -foregroundColor yellow -backgroundcolor darkgreen
write-host "   Core eDiscovery cases - Add users to a hold   " -foregroundColor yellow -backgroundcolor
darkgreen
write-host "************************************************"
" "
# Connect to SCC PowerShell using modern authentication
if (!$SccSession)
{
  Import-Module ExchangeOnlineManagement
  Connect-IPPSSession
}

# Get the organization's domain name. We use this to create the SharePoint admin URL and root URL for
OneDrive for Business.
""
$mySiteDomain = Read-Host "Enter the domain name for your SharePoint organization. We use this name to
connect to SharePoint admin center and for the OneDrive URLs in your organization. For example, 'contoso' in
'https://contoso-admin.sharepoint.com' and 'https://contoso-my.sharepoint.com'"
""

# Connect to PnP Online using modern authentication
Import-Module PnP.PowerShell
Connect-PnPOnline -Url https://$mySiteDomain-admin.sharepoint.com -UseWebLogin

# Load the SharePoint assemblies from the SharePoint Online Management Shell
```

```
# To install, go to https://go.microsoft.com/fwlink/p/?LinkId=255251
if (!$SharePointClient -or !$SPRuntime -or !$SPUserProfile)
{
    $SharePointClient = [System.Reflection.Assembly]::LoadWithPartialName("Microsoft.SharePoint.Client")
    $SPRuntime = [System.Reflection.Assembly]::LoadWithPartialName("Microsoft.SharePoint.Client.Runtime")
    $SPUserProfile =
[System.Reflection.Assembly]::LoadWithPartialName("Microsoft.SharePoint.Client.UserProfiles")
    if (!$SharePointClient)
    {
        Write-Error "The SharePoint Online Management Shell isn't installed. Please install it from:
https://go.microsoft.com/fwlink/p/?LinkId=255251 and then re-run this script."
        return;
    }
}

# Get other required information
do{
$casename = Read-Host "Enter the name of the case"
$caseexists = (get-compliancecase -identity "$casename" -erroraction SilentlyContinue).isvalid
if($caseexists -ne 'True')
{""
write-host "A case named '$casename' doesn't exist. Please specify the name of an existing case, or create a
new case and then re-run the script." -foregroundColor Yellow
""}
}While($caseexists -ne 'True')
""
do{
$holdName = Read-Host "Enter the name of the new hold"
$holdexists=(get-caseholdpolicy -identity "$holdname" -case "$casename" -erroraction
SilentlyContinue).isvalid
if($holdexists -eq 'True')
{""
write-host "A hold named '$holdname' already exists. Please specify a new hold name." -foregroundColor
Yellow
""}
}While($holdexists -eq 'True')
""
$holdQuery = Read-Host "Enter a search query to create a query-based hold, or press Enter to hold all
content"
""
$holdstatus = read-host "Do you want the hold enabled after it's created? (Yes/No)"
do{
""
$inputfile = read-host "Enter the name of the text file that contains the email addresses of the users to
add to the hold"
""
$fileexists = test-path -path $inputfile
if($fileexists -ne 'True'){write-host "$inputfile doesn't exist. Please enter a valid file name." -
foregroundcolor Yellow}
}while($fileexists -ne 'True')
#Import the list of addresses from the txt file.  Trim any excess spaces and make sure all addresses
    #in the list are unique.
  [array]$emailAddresses = Get-Content $inputfile -ErrorAction SilentlyContinue | where {$_.trim() -ne ""}
| foreach{ $_.Trim() }
  [int]$dupl = $emailAddresses.count
  [array]$emailAddresses = $emailAddresses | select-object -unique
  $dupl -= $emailAddresses.count
#Validate email addresses so the hold creation does not run in to an error.
if($emailaddresses.count -gt 0){
write-host ($emailAddresses).count "addresses were found in the text file. There were $dupl duplicate
entries in the file." -foregroundColor Yellow
""
Write-host "Validating the email addresses. Please wait..." -foregroundColor Yellow
""
$finallist =@()
foreach($emailAddress in $emailAddresses)
{
if((get-recipient $emailaddress -erroraction SilentlyContinue).isvalid -eq 'True')
{$finallist += $emailaddress}
```

```
{$finallist += $emailaddress}
else {"Unable to find the user $emailaddress"
[array]$excludedlist += $emailaddress}
}
""
#Find user's OneDrive account URL using email address
Write-Host "Getting the URL for each user's OneDrive for Business site." -foregroundColor Yellow
""
$AdminUrl = "https://$mySiteDomain-admin.sharepoint.com"
$mySiteUrlRoot = "https://$mySiteDomain-my.sharepoint.com"
$urls = @()
foreach($emailAddress in $emailAddresses)
{
try
{
$url=Get-PnPUserProfileProperty -Account $emailAddress | Select PersonalUrl
$urls += $url.PersonalUrl
       Write-Host "- $emailAddress => $url"
       [array]$ODadded += $url.PersonalUrl
       }catch {
 Write-Warning "Could not locate OneDrive for $emailAddress"
 [array]$ODExluded += $emailAddress
 Continue }
}
$urls | FL
if(($finallist.count -gt 0) -or ($urls.count -gt 0)){
""
Write-Host "Creating the hold named $holdname. Please wait..." -foregroundColor Yellow
if(($holdstatus -eq "Y") -or ($holdstatus -eq  "y") -or ($holdstatus -eq "yes") -or ($holdstatus -eq "YES"))
{
New-CaseHoldPolicy -Name "$holdName" -Case "$casename" -ExchangeLocation $finallist -SharePointLocation
$urls -Enabled $True | out-null
New-CaseHoldRule -Name "$holdName" -Policy "$holdname" -ContentMatchQuery $holdQuery | out-null
}
else{
New-CaseHoldPolicy -Name "$holdName" -Case "$casename" -ExchangeLocation $finallist -SharePointLocation
$urls -Enabled $false | out-null
New-CaseHoldRule -Name "$holdName" -Policy "$holdname" -ContentMatchQuery $holdQuery -disabled $true | out-
null
}
""
}
else {"No valid locations were identified. Therefore, the hold wasn't created."}
#write log files (if needed)
$newhold=Get-CaseHoldPolicy -Identity "$holdname" -Case "$casename" -erroraction SilentlyContinue
$newholdrule=Get-CaseHoldRule -Identity "$holdName" -erroraction SilentlyContinue
if(($ODAdded.count -gt 0) -or ($ODExluded.count -gt 0) -or ($finallist.count -gt 0) -or ($excludedlist.count
-gt 0) -or ($newhold.isvalid -eq 'True') -or ($newholdrule.isvalid -eq 'True'))
{
Write-Host "Generating output files..." -foregroundColor Yellow
if($ODAdded.count -gt 0){
"OneDrive Locations" | add-content .\LocationsOnHold.txt
"==================" | add-content .\LocationsOnHold.txt
$newhold.SharePointLocation.name | add-content .\LocationsOnHold.txt}
if($ODExluded.count -gt 0){
"Users without OneDrive locations" | add-content .\LocationsNotOnHold.txt
"================================" | add-content .\LocationsNotOnHold.txt
$ODExluded | add-content .\LocationsNotOnHold.txt}
if($finallist.count -gt 0){
" " | add-content .\LocationsOnHold.txt
"Exchange Locations" | add-content .\LocationsOnHold.txt
"==================" | add-content .\LocationsOnHold.txt
$newhold.ExchangeLocation.name | add-content .\LocationsOnHold.txt}
if($excludedlist.count -gt 0){
" "| add-content .\LocationsNotOnHold.txt
"Mailboxes not added to the hold" | add-content .\LocationsNotOnHold.txt
"================================" | add-content .\LocationsNotOnHold.txt
$excludedlist | add-content .\LocationsNotOnHold.txt}
$FormatEnumerationLimit=-1
if($newhold.isvalid -eq 'True'){$newhold|fl > .\GetCaseHoldPolicy.txt}
```

```
if($newhold.isvalid -eq 'True'){$newhold|Fl >.\GetCaseHoldPolicy.txt}
if($newholdrule.isvalid -eq 'True'){$newholdrule|Fl >.\GetCaseHoldRule.txt}
}
}
else {"The hold wasn't created because no valid entries were found in the text file."}
""
#Disconnect from SCC PowerShell and PnPOnline

Write-host "Disconnecting from SCC PowerShell and PnP Online" -foregroundColor Yellow
Get-PSSession | Remove-PSSession
Disconnect-PnPOnline

Write-host "Script complete!" -foregroundColor Yellow
""
#script end
```

2. On your local computer, open Windows PowerShell and go to the folder where you saved the script.
3. Run the script; for example:

   ```
   .\AddUsersToHold.ps1
   ```

4. Enter the information that the script prompts you for.

   The script connects to Security & Compliance Center PowerShell, and then creates the new hold in the eDiscovery case and adds the mailboxes and OneDrive for Business for the users in the list. You can go to the case on the **eDiscovery** page in the Security & Compliance Center to view the new hold.

After the script is finished running, it creates the following log files, and saves them to the folder where the script is located.

- **LocationsOnHold.txt:** Contains a list of mailboxes and OneDrive for Business sites that the script successfully placed on hold.
- **LocationsNotOnHold.txt:** Contains a list of mailboxes and OneDrive for Business sites that the script did not place on hold. If a user has a mailbox, but not a OneDrive for Business site, the user would be included in the list of OneDrive for Business sites that weren't placed on hold.
- **GetCaseHoldPolicy.txt:** Contains the output of the **Get-CaseHoldPolicy** cmdlet for the new hold, which the script ran after creating the new hold. The information returned by this cmdlet includes a list of users whose mailboxes and OneDrive for Business sites were placed on hold and whether the hold is enabled or disabled.
- **GetCaseHoldRule.txt:** Contains the output of the **Get-CaseHoldRule** cmdlet for the new hold, which the script ran after creating the new hold. The information returned by this cmdlet includes the search query if you used the script to create a query-based hold.

Create a report on holds in eDiscovery cases

11/2/2020 • 6 minutes to read • Edit Online

The script in this article lets eDiscovery administrators and eDiscovery managers generate a report that contains information about all holds that are associated with eDiscovery cases in the the compliance center in Office 365 or Microsoft 365. The report contains information such as the name of the case a hold is associated with, the content locations that are placed on hold, and whether the hold is query-based. If there are cases that don't have any holds, the script will create an additional report with a list of cases without holds.

See the More information section for a detailed description of the information included in the report.

Admin requirements and script information

- To generate a report on all eDiscovery cases in your organization, you have to be an eDiscovery Administrator in your organization. If you are an eDiscovery Manager, the report will only include information about the cases that you can access. For more information about eDiscovery permissions, see Assign eDiscovery permissions.
- The script in this article has minimal error handling. The primary purpose is to quickly create report about the holds that are associated with the eDiscovery cases in your organization.
- The sample scripts provided in this topic aren't supported under any Microsoft standard support program or service. The sample scripts are provided AS IS without warranty of any kind. Microsoft further disclaims all implied warranties including, without limitation, any implied warranties of merchantability or of fitness for a particular purpose. The entire risk arising out of the use or performance of the sample scripts and documentation remains with you. In no event shall Microsoft, its authors, or anyone else involved in the creation, production, or delivery of the scripts be liable for any damages whatsoever (including, without limitation, damages for loss of business profits, business interruption, loss of business information, or other pecuniary loss) arising out of the use of or inability to use the sample scripts or documentation, even if Microsoft has been advised of the possibility of such damages.

Step 1: Connect to the Security & Compliance Center PowerShell

The first step is to connect to Security & Compliance Center PowerShell for your organization. For step-by-step instructions, see Connect to Security & Compliance Center PowerShell.

Step 2: Run the script to report on holds associated with eDiscovery cases

After you've connected to Security & Compliance Center PowerShell, the next step is to create and run the script that collects information about the eDiscovery cases in your organization.

1. Save the following text to a Windows PowerShell script file by using a filename suffix of .ps1; for example, CaseHoldsReport.ps1.

```
#script begin
" "
write-host "*********************************************"
write-host "   Security & Compliance Center   " -foregroundColor yellow -backgroundcolor darkgreen
write-host "        eDiscovery cases - Holds report         " -foregroundColor yellow -
backgroundcolor darkgreen
```

```
write-host "*********************************************"
" "
#prompt users to specify a path to store the output files
$time=get-date
$Path = Read-Host 'Enter a file path to save the report to a .csv file'
$outputpath=$Path+'\'+'CaseHoldsReport'+' '+$time.day+'-'+$time.month+'-'+$time.year+'
'+$time.hour+'.'+$time.minute+'.csv'
$noholdsfilepath=$Path+'\'+'CaseswithNoHolds'+' '+$time.day+'-'+$time.month+'-'+$time.year+'
'+$time.hour+'.'+$time.minute+'.csv'
#add case details to the csv file
function add-tocasereport{
Param([string]$casename,
[String]$casestatus,
[datetime]$casecreatedtime,
[string]$casemembers,
[datetime]$caseClosedDateTime,
[string]$caseclosedby,
[string]$holdname,
[String]$Holdenabled,
[string]$holdcreatedby,
[string]$holdlastmodifiedby,
[string]$ExchangeLocation,
[string]$sharePointlocation,
[string]$ContentMatchQuery,
[datetime]$holdcreatedtime,
[datetime]$holdchangedtime
)
$addRow = New-Object PSObject
Add-Member -InputObject $addRow -MemberType NoteProperty -Name "Case name" -Value $casename
Add-Member -InputObject $addRow -MemberType NoteProperty -Name "Case status" -Value $casestatus
Add-Member -InputObject $addRow -MemberType NoteProperty -Name "Case members" -Value $casemembers
Add-Member -InputObject $addRow -MemberType NoteProperty -Name "Case created time" -Value
$casecreatedtime
Add-Member -InputObject $addRow -MemberType NoteProperty -Name "Case closed time" -Value
$caseClosedDateTime
Add-Member -InputObject $addRow -MemberType NoteProperty -Name "Case closed by" -Value $caseclosedby
Add-Member -InputObject $addRow -MemberType NoteProperty -Name "Hold name" -Value $holdname
Add-Member -InputObject $addRow -MemberType NoteProperty -Name "Hold enabled" -Value $Holdenabled
Add-Member -InputObject $addRow -MemberType NoteProperty -Name "Hold created by" -Value
$holdcreatedby
Add-Member -InputObject $addRow -MemberType NoteProperty -Name "Hold last changed by" -Value
$holdlastmodifiedby
Add-Member -InputObject $addRow -MemberType NoteProperty -Name "Exchange locations" -Value
$ExchangeLocation
Add-Member -InputObject $addRow -MemberType NoteProperty -Name "SharePoint locations" -Value
$sharePointlocation
Add-Member -InputObject $addRow -MemberType NoteProperty -Name "Hold query" -Value $ContentMatchQuery
Add-Member -InputObject $addRow -MemberType NoteProperty -Name "Hold created time (UTC)" -Value
$holdcreatedtime
Add-Member -InputObject $addRow -MemberType NoteProperty -Name "Hold changed time (UTC)" -Value
$holdchangedtime
$allholdreport = $addRow | Select-Object "Case name","Case status","Hold name","Hold enabled","Case
members", "Case created time","Case closed time","Case closed by","Exchange locations","SharePoint
locations","Hold query","Hold created by","Hold created time (UTC)","Hold last changed by","Hold
changed time (UTC)"
$allholdreport | export-csv -path $outputPath -notypeinfo -append -Encoding ascii
}
#get information on the cases and pass values to the case report function
" "
write-host "Gathering a list of cases and holds..."
" "
$edc =Get-ComplianceCase -ErrorAction SilentlyContinue
foreach($cc in $edc)
{
write-host "Working on case :" $cc.name
if($cc.status -eq 'Closed')
{
$cmembers = ((Get-ComplianceCaseMember -Case $cc.name).windowsLiveID)-join ';'
add-tocasereport -casename $cc.name -casestatus $cc.Status -caseclosedby $cc.closedby -
```

```
caseClosedDateTime $cc.ClosedDateTime -casemembers $cmembers
}
else{
$cmembers = ((Get-ComplianceCaseMember -Case $cc.name).windowsLiveID)-join ';'
$policies = Get-CaseHoldPolicy -Case $cc.Name | %{ Get-CaseHoldPolicy $_.Name -Case $_.CaseId -
DistributionDetail}
if ($policies -ne $NULL)
{
foreach ($policy in $policies)
{
$rule=Get-CaseHoldRule -Policy $policy.name
add-tocasereport -casename $cc.name -casemembers $cmembers -casestatus $cc.Status -casecreatedtime
$cc.CreatedDateTime -holdname $policy.name -holdenabled $policy.enabled -holdcreatedby
$policy.CreatedBy -holdlastmodifiedby $policy.LastModifiedBy -ExchangeLocation
(($policy.exchangelocation.name)-join ';') -SharePointLocation (($policy.sharePointlocation.name)-
join ';') -ContentMatchQuery $rule.ContentMatchQuery -holdcreatedtime $policy.WhenCreatedUTC -
holdchangedtime $policy.WhenChangedUTC
}
}
else{
write-host "No hold policies found in case:" $cc.name -foregroundColor 'Yellow'
" "
[string]$cc.name | out-file -filepath $noholdsfilepath -append
}
}
}

" "
Write-host "Script complete! Report files saved to this folder: '$Path'"
" "
#script end
```

2. In the Windows PowerShell session that opened in Step 1, go to the folder where you saved the script.
3. Run the script; for example:

   ```
   .\CaseHoldsReport.ps1
   ```

   The script will prompt for a target folder to save the report to.
4. Type the full path name of the folder to save the report to, and then press **Enter**.

   > **TIP**
   > To save the report in the same folder that the script is located in, type a period (".") when prompted for a target folder. To save the report in a subfolder in the folder where the script is located, just type the name of the subfolder.

   The script starts to collect information about all the eDiscovery cases in your organization. Don't access the report file while the script is running. After the script is complete, a confirmation message is displayed in the Windows PowerShell session. After this message is displayed, you can access the report in the folder that you specified in Step 4. The file name for the report is `CaseHoldsReport<DateTimeStamp>.csv`.

   Addtionally, the script also creates a report with a list of cases that don't have any holds. The file name for this report is `CaseswithNoHolds<DateTimeStamp>.csv`.

   Here's an example of running the CaseHoldsReport.ps1 script.

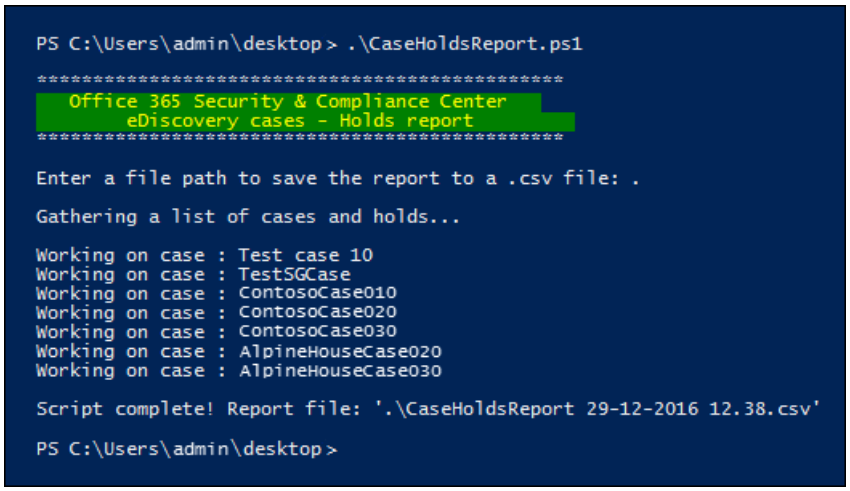

More information

The case holds report that's created when you run the script in this article contains the following information about each hold. As previously explained, you have to be an eDiscovery Administrator to return information for all holds in your organization. For more information about case holds, see eDiscovery cases.

- The name of the hold and the name of the eDiscovery case that the hold is associated with.
- Whether or not the eDiscovery case is active or closed.
- Whether or not the hold is enabled or disabled.
- The members of the eDiscovery case that the hold is associated with. Case members can view or manage a case, depending on the eDiscovery permissions they've been assigned.
- The time and date the case was created.
- If a case is closed, the person who closed it and the time and date it was closed.
- The Exchange mailboxes and SharePoint sites locations that are on hold.
- If the hold is query-based, the query syntax.
- The time and date the hold was created and the person who created it.
- The time and date the hold was last changed and the person who changed it.

Export a Content Search report

11/2/2020 • 7 minutes to read • Edit Online

Instead of exporting the full set of search results from a Content Search in the Security & Compliance Center (and from a Content Search that's associated with an eDiscovery case), you can export the same reports that are generated when you export search results.

When you export a report, it's downloaded to a folder that has the same name as the Content Search, but that's appended with *_ReportsOnly*. For example, if the Content Search is named *ContosoCase0815*, then the report is downloaded to a folder named *ContosoCase0815_ReportsOnly*. For a list of documents that are included in the report, see What's included in the report.

Assign roles and check system requirements

- To export a Content Search report, you have to be assigned the Compliance Search management role in the Security & Compliance Center. This role is assigned by default to the built-in eDiscovery Manager and Organization Management role groups. For more information, see Assign eDiscovery permissions.
- When you export a report, the data is temporarily stored in a unique Azure Storage area in the Microsoft cloud before it's downloaded to your local computer. Be sure that your organization can connect to the endpoint in Azure, which is ***.blob.core.windows.net** (the wildcard represents a unique identifier for your export). The search results data is deleted from the Azure Storage area two weeks after it's created.
- The computer you use to export the search results has to meet the following system requirements:
  - 32-bit or 64-bit versions of Windows 7 and later versions
  - Microsoft .NET Framework 4.7
- You have to use one of the following supported browsers to run the eDiscovery Export Tool[1]:
  - Microsoft Edge [2]

    OR
  - Microsoft Internet Explorer 10 and later versions

> **NOTE**
>
> [1] Microsoft doesn't manufacture third-party extensions or add-ons for ClickOnce applications. Exporting search results using an unsupported browser with third-party extensions or add-ons isn't supported.
>
> [2] As a result of recent changes to Microsoft Edge, ClickOnce support is no longer enabled by default. For instructions on enabling ClickOnce support in Edge, see Use the eDiscovery Export Tool in Microsoft Edge.

- If the estimated total size of the results returned by a Content Search exceeds 2 TB, exporting the report fails. To successfully export the report, try to narrow the scope and rerun the search so the estimated size of the results is less than 2 TB.
- Exporting Content Search reports counts against the maximum number of exports running at the same time and the maximum number of exports that a single user can run. For more information about export limits, see Export Content Search results.

Generate and download a Content Search report

The steps to generate and download a Content Search report are similar to actually exporting search results.

Step 1: Generate the report for export

The first step is to prepare the report for downloading to your computer exporting. When you the report, the report documents are uploaded to an Azure Storage area in the Microsoft cloud.

1. Go to https://protection.office.com.
2. Sign in using your work or school account.
3. In the left pane of the Security & Compliance Center, click **Search** > **Content search**.
4. On the **Content search** page, select a search.
5. In the details pane, under **Export report to a computer**, click **Generate report**.

   > **NOTE**
   > If the results for a search are older than 7 days, you are prompted to update the search results. If this happens, cancel the export, click **Update search results** in the details pane for the selected search, and then start the report export again after the results are updated.

6. On the **Export a report** page, under **Include these items from the search**, choose one of the following options:

   - Export only indexed items
   - Export indexed and unindexed items
   - Export only unindexed items

   For more information about unindexed items, see Partially indexed items in Content Search.

7. Choose to include search statistics for all versions of SharePoint documents. This option appears only if the content sources of the search include SharePoint or OneDrive for Business sites.
8. Click **Generate report**.

   The search results report is prepared for downloading, which means the report documents will be uploaded to the Azure Storage area in the Microsoft cloud. When the report is ready for download, the **Download report** link is displayed under **Export report to a computer** in the details pane.

> **NOTE**
> You can also export a report for a Content Search that's associated with an eDiscovery case. To do this, go to **eDiscovery** > **eDiscovery**, select a case, and click **Edit** ✎. On the **Searches** page, select a search, and then click **Export** ⬇ > **Export a report**.

Step 2: Download the report

The next step is to download the report from the Azure Storage area to your local computer.

1. In the details pane for the search that you generated the report for, under **Export report to a computer**, click **Download report**.

   The **Download report** page is displayed and contains the following information about the report that's downloaded to your computer.

- The number of items that will be downloaded.
- The estimated total size of the items that will be downloaded.
- Whether indexed or unindexed will be exported. Unindexed items are items that have a recognized format, are encrypted, or weren't indexed for other reasons.
- Whether versions of SharePoint documents will be downloaded.
- The status of the report export process. You can start downloading the report even if the preparation of the report isn't complete.

2. Under **Export key**, click **Copy to clipboard**. You use this key in step 5 to download the report.

   > **IMPORTANT**
   > Because anyone can install and start the eDiscovery Export tool, and then use this key to download the search report, be sure to take precautions to protect this key just like you would protect passwords or other security-related information.

3. Click **Download report**.
4. If you're prompted to install the **eDiscovery Export Tool**, click **Install**.
5. In the **eDiscovery Export Tool**, paste the export key that you copied in step 2 in the appropriate box.
6. Click **Browse** to specify the location where you want to download the report.
7. Click **Start** to download the search results to your computer.

   The **eDiscovery Export Tool** displays status information about the export process, including an estimate of the number (and size) of the remaining items to be downloaded. When the export process is complete, you can access the files in the location where they were downloaded.

> **NOTE**
> You can download the report for a Content Search that's associated with an eDiscovery case. To do this, go to **eDiscovery** > **eDiscovery**, select a case, and click **Edit** ✎. On the **Exports** page, select an report export, and then click **Download report** in the details pane.

What's included in the report

When you generate and export a report about the results of a Content Search, the following documents are downloaded:

- **Export Summary:** An Excel document that contains a summary of the export. This includes information such as the number of content sources that were searched, the number of search results from each content location, the estimated number of items, the actual number of items that would be exported, and the estimated and actual size of items that would be exported.

> **NOTE**
> If you include unindexed items when exporting the report, the number of unindexed items are included in the total number of estimated search results and in the total number of downloaded search results (if you were to export the search results) that are listed in the Export Summary report. In other words, the total number of items that would be downloaded is equal to the total number of estimated results and the total number of unindexed items.

- **Manifest:** A manifest file (in XML format) that contains information about each item included in the search results.
- **Results:** An Excel document that contains a row with information about each indexed item that would be exported with the search results. For email, the result log contains information about each message, including:
  - The location of the message in the source mailbox (including whether the message is in the primary or archive mailbox).
  - The date the message was sent or received.
  - The Subject line from the message.
  - The sender and recipients of the message.

    For documents from SharePoint and OneDrive for Business sites, the Results log contains information about each document, including:
  - The URL for the document.
  - The URL for the site collection where the document is located.
  - The date that the document was last modified.
  - The name of the document (which is located in the Subject column in the result log).

    > **NOTE**
    > The number of rows in the **Results** report should be equal to the total number of search results minus the total number of items listed in the **Unindexed Items** report.
- **Unindexed Items:** An Excel document that contains information about any unindexed items included in the search results. If you don't include unindexed items when you generate the search results report, this report will still be downloaded, but will be empty.

View keyword statistics for Content Search results

2/18/2021 • 6 minutes to read • Edit Online

After you create and run a Content Search, you can view statistics about the estimated search results. This includes a summary of the search results (similar to the summary of the estimated search results displayed in the details pane), the query statistics such as the number of content locations with items that match the search query, and the name of content locations that have the most matching items. You can display statistics for one or more content searches. This lets you to quickly compare the results for multiple searches and make decisions about the effectiveness of your search queries.

Additionally, you can configure new and existing searches to return statistics for each keyword in a search query. This lets you compare the number of results for each keyword in a query and to compare the keyword statistics from multiple searches.

You can also download the search statistics and keyword statistics to a CSV file. This lets you use the filtering and sorting features in Excel to compare results, and prepare reports for your search results.

Get statistics for Content Searches

To display statistics for Content searches:

1. In the Microsoft 365 compliance center, go to **Show all** > **Content search**.
2. In the list of searches, select two or more searches, and then click **Search statistics** on the **Bulk actions** flyout page.

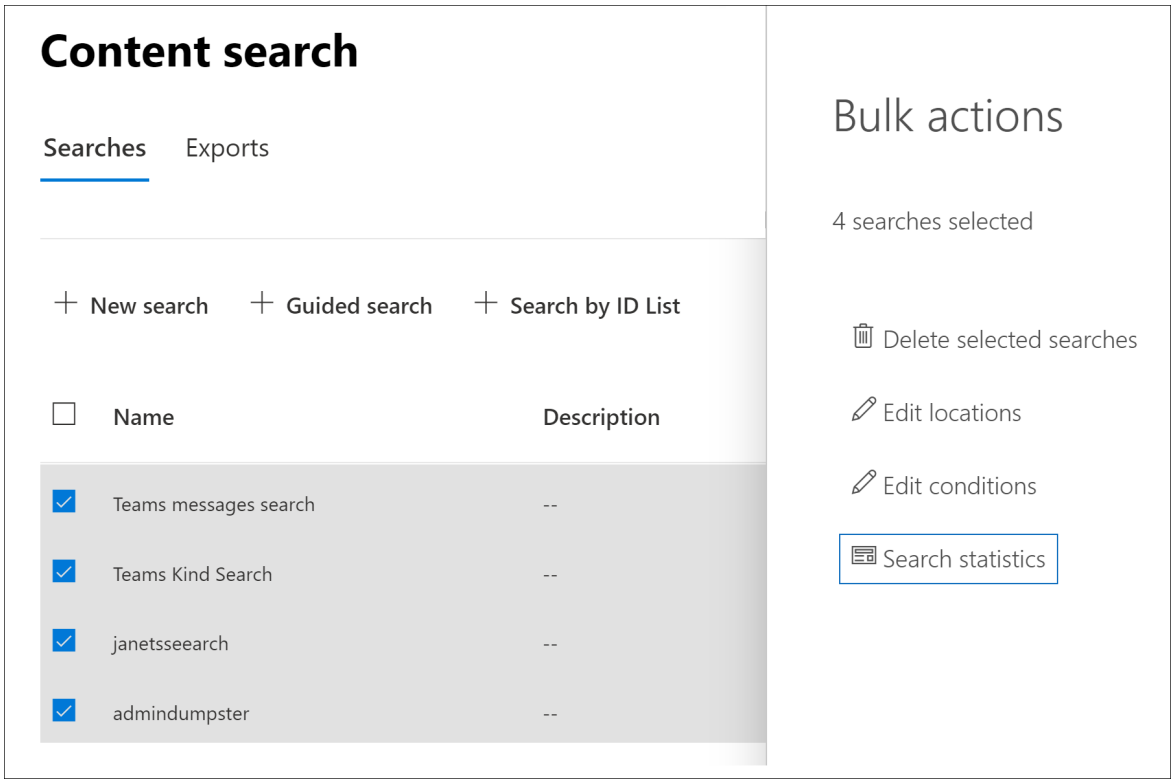

3. On the **Search statistics** page, click one of the following links to display statistics about the selected searches.

   Summary

   This page displays statistics similar to the ones displayed in the details pane on the **Content search**

page. Statistics for all selected searches are displayed. Note that you can also re-run the selected searches from this page to update the statistics.

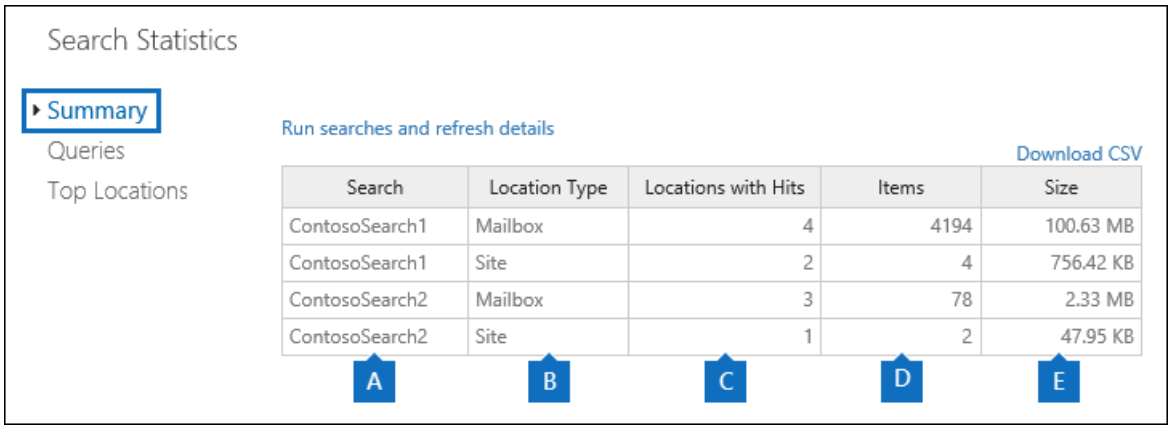

a. The name of the Content Search. As previously stated, you can display and compare statistics for multiple searches.

b. The type of content location that was searched. Each row displays statistics for mailboxes, sites, and public folders from the specified search.

c. The number of content locations containing items that match the search query. For mailboxes, this statistic also includes the number of archive mailboxes that contain items that match the search query.

d. The total number of items of all specified content locations that match the search query. Examples of item types include email messages, calendar items, and documents. If an item contains multiple instances of a keyword that is being searched for, it's only counted once in the total number of items. For example, if you're searching for words "stock" or "fraud" and an email message contains three instances of the word "stock", it's only counted once in the **Items** column.

e. The total size of all items that were found in the specified content location that match the search query.

Queries

This page displays statistics about the search query.

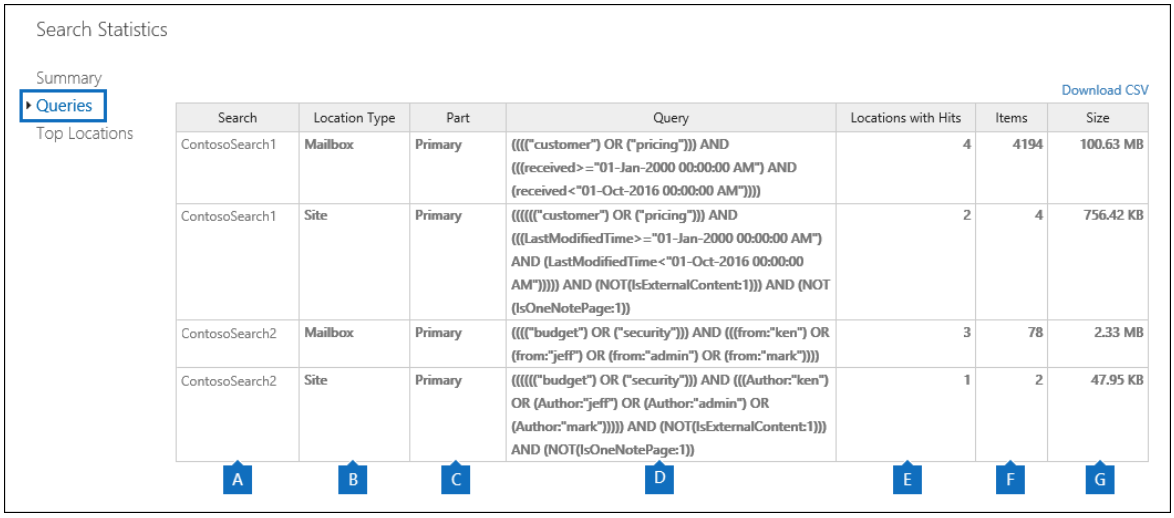

a. The name of the Content Search that the row contains query statistics for.

b. The type of content location that the query statistics are applicable to.

c. This column indicates which part of the search query the statistics are applicable to. **Primary** indicates the entire search query. If you use a keyword list when you create or edit a search query, statistics for each component of the query are included in this table. See the Get keyword statistics for Content

Searches section in this article for more information.

d. This column contains the actual search query that run by the Content Search tool. Note that the tool automatically adds a few additional components to the query that you create.

- When you search for all content in mailboxes (by not specifying any keywords), the actual key word query is `size>=0` so that all items are returned.
- When you search SharePoint Online and OneDrive for Business sites, the two following components are added:

  NOT IsExternalContent:1 - Excludes any content from an on-premises SharePoint organization.

  NOT IsOneNotePage:1 - Excludes all OneNote files because these would be duplicates of any document that matches the search query.

e. The number of the content locations (specified by the ** Location type ** column) that contain items that match the search query listed in the **Query** column.

f. The number of items (from the specified content location) that match the search query listed in the **Query** column. As previously explained, if an item contains multiple instances of a keyword that is being searched for, it's only counted once in the this column.

g. The total size of all items that were found (in the specified content location) that match the search query in the **Query** column.

Top locations

This page displays statistics about the number of items that match the search query in each content location that was searched. The top 1,000 locations are displayed. If you view statistics for multiple searches, the top 1,000 locations for each search are displayed. Note that a content location isn't included on this page if it doesn't contain any items that match the search query.

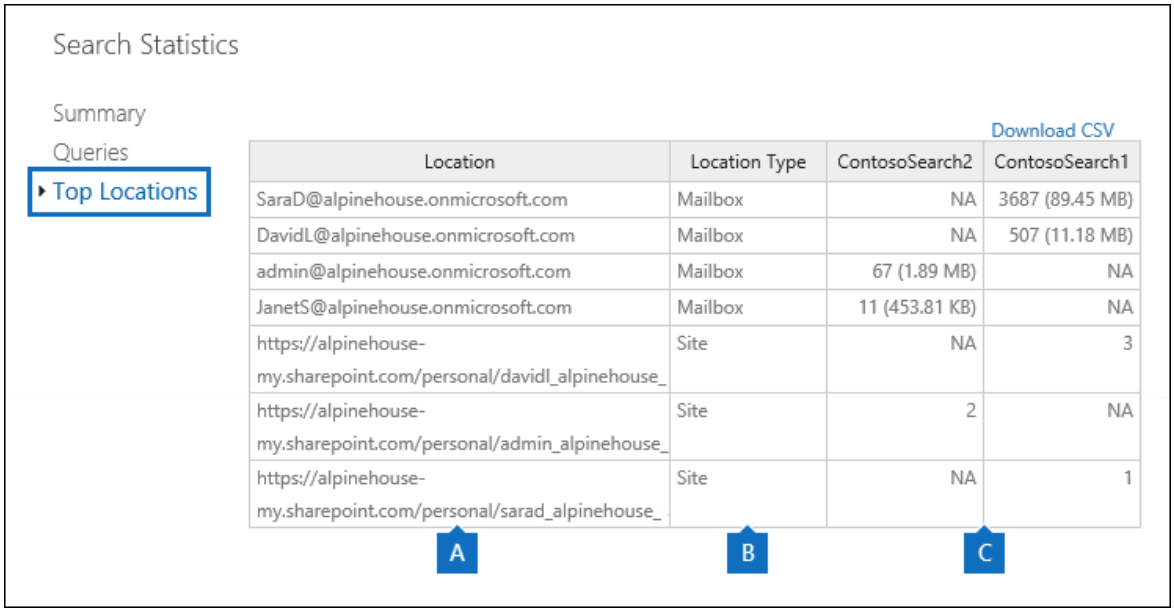

a. The name of the content location.

b. The type of content location that the location statistics are applicable to.

c. There are columns for each search that you're displaying statistics for. This column shows the number (and total size) of items that match the search query in each content location. Note that when you're displaying statistics for multiple searches, the "NA" in this column indicates that the content location wasn't included in that search.

Get keyword statistics for Content Searches

As previous explained, the **Queries** page shows the search query and the number (and size) of items that match the query. If you use a keyword list when you create or edit a search query, you can get enhanced statistics that show how many items match each keyword or keyword phrase. This can help you quickly identify which parts of the query are the most (and least) effective. For example, if a keyword returns a large number of items, you might choose to refine the keyword query to narrow the search results. You can set up a keyword list when you create or edit a Content Search.

To create a keyword list and view keyword statistics for a Content Search:

1. In the Microsoft 365 compliance center, go to **Show all** > **Content search**.
2. In the list of content searches, click and a search, and then click **Edit** .
3. Click **Query** and then do the following things:

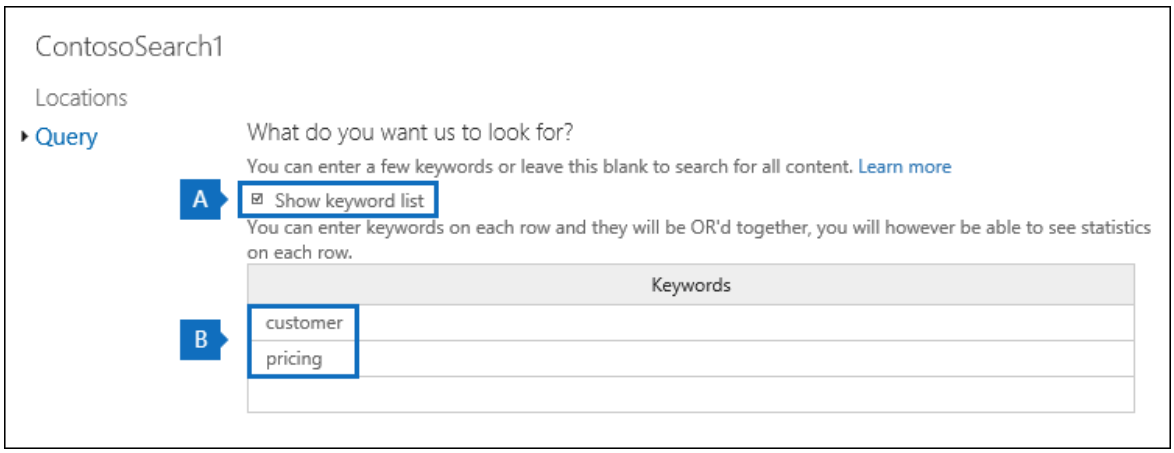

   a. Click the **Show keyword list** check box.

   b. Type a keyword or keyword phase in a row in the keywords table. For example, type **budget** in the first row and then type **security** in the second row.

4. After adding the keywords that you want to search and get statistics for, click **Search** to run the revised search.
5. When the search is completed, select it in the list of searches, and then click **Search statistics** . You can also display and compare keyword statistics for multiple searches.
6. On the **Search statistics** page, click **Query** to display the keyword statistics for the selected searches.

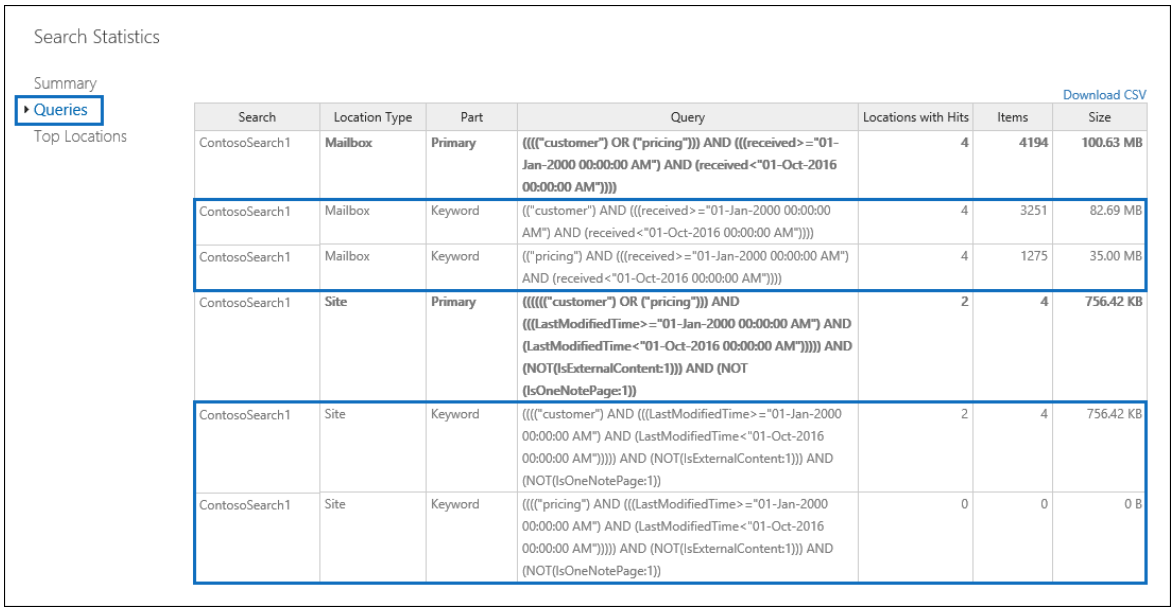

Search	Location Type	Part	Query	Locations with Hits	Items	Size
ContosoSearch1	**Mailbox**	**Primary**	**((("customer") OR ("pricing"))) AND (((received>="01-Jan-2000 00:00:00 AM") AND (received<"01-Oct-2016 00:00:00 AM"))))**	**4**	**4194**	**100.63 MB**
ContosoSearch1	Mailbox	Keyword	(("customer") AND (((received>="01-Jan-2000 00:00:00 AM") AND (received<"01-Oct-2016 00:00:00 AM"))))	4	3251	82.69 MB
ContosoSearch1	Mailbox	Keyword	(("pricing") AND (((received>="01-Jan-2000 00:00:00 AM") AND (received<"01-Oct-2016 00:00:00 AM"))))	4	1275	35.00 MB
ContosoSearch1	**Site**	**Primary**	**((((("customer") OR ("pricing"))) AND (((LastModifiedTime>="01-Jan-2000 00:00:00 AM") AND (LastModifiedTime<"01-Oct-2016 00:00:00 AM"))))) AND (NOT(IsExternalContent:1))) AND (NOT (IsOneNotePage:1))**	**2**	**4**	**756.42 KB**
ContosoSearch1	Site	Keyword	((((("customer") AND (((LastModifiedTime>="01-Jan-2000 00:00:00 AM") AND (LastModifiedTime<"01-Oct-2016 00:00:00 AM"))))) AND (NOT(IsExternalContent:1))) AND (NOT(IsOneNotePage:1))	2	4	756.42 KB
ContosoSearch1	Site	Keyword	((((("pricing") AND (((LastModifiedTime>="01-Jan-2000 00:00:00 AM") AND (LastModifiedTime<"01-Oct-2016 00:00:00 AM"))))) AND (NOT(IsExternalContent:1))) AND (NOT(IsOneNotePage:1))	0	0	0 B

As shown in the previous screenshot, the statistics for each keyword are displayed; this includes:

- The keyword statistics for each type of content location included in the search.
- The actual search query for each keyword, which includes any conditions from the search query.
- The complete search query (identified as **Primary** in the **Part** column) and the statistics for the complete query. Note these are the same statistics displayed on the **Summary** page.

> **NOTE**
>
> To help reduce issues caused by large keyword lists, you're now limited to a maximum of 20 rows in the keyword list of a search query.

eDiscovery solution series: Data spillage scenario - Search and purge

11/2/2020 • 14 minutes to read • Edit Online

What is data spillage and why should you care? Data spillage is when a confidential document is released into an untrusted environment. When a data spillage incident is detected, it's important to quickly assess the size and locations of the spillage, examine user activities around it, and then permanently purge the spilled data from the system.

Data spillage scenario

You're a lead information security officer at Contoso. You are informed of a data spillage situation where an employee unknowingly shared a highly confidential document with multiple people through email. You want to quickly assess who received this document internally and externally. Once identified, you would like to share case findings with other investigators to review, and then permanently remove the data from Office 365. After the investigation is complete, you want to generate a report with the evidence of permanent removal and other case details for any future reference.

Scope of this article

This document provides a list of instructions on how to permanently remove a message from Microsoft 365 so that it's not accessible or recoverable. To delete a message and make it recoverable until the deleted item retention period expires, see Search for and delete email messages in your organization.

Workflow for managing data spillage incidents

Here's a how to manage a data spillage incident:

(Optional) Step 1: Manage who can access the case and set compliance boundaries
Step 2: Create an eDiscovery case
Step 3: Search for the spilled data
Step 4: Review and validate case findings
Step 5: Use message trace log to check how spilled data was shared
Step 6: Prepare the mailboxes
Step 7: Permanently delete the spilled data
Step 8: Verify, provide a proof of deletion, and audit

Things to know before you start

- When a mailbox is on hold, a deleted message remains in the Recoverable Items folder until the retention period expires or the hold is released. Step 6 describes how to remove hold from the mailboxes. Check with your records management or legal departments before removing the hold. Your organization might have a policy that defines whether a mailbox on hold or a data spillage incident takes priority.
- To control which user mailboxes an data spillage investigator can search and manage who can access the case, you can set up compliance boundaries and create a custom role group, which is described in Step 1. To do this, you have to be a member of the Organization Management role group or be assigned the role

management role. If you or in administrator in your organization has already set compliance boundaries, you can skip Step 1.

- To create a case, you must be a member of the eDiscovery Manager role group or be a member of a custom role group that's assigned the Case Management role. If you're not a member, ask a Microsoft 365 administrator to add you to the eDiscovery manager role group.
- To create and run a Content Search, you have to be a member of the eDiscovery Manager role group or be assigned the Compliance Search management role. To delete messages, you have to be a member of the Organization Management role group or be assigned the Search And Purge management role. For information about adding users to a role group, see Assign eDiscovery permissions in the Security & Compliance Center.
- To search the audit log eDiscovery activities in Step 8, auditing must be turned on for your organization. You can search for activities that were performed within the last 90 days. To learn more about how to enable and use auditing, see the Auditing the data spillage investigation process section in Step 8.

(Optional) Step 1: Manage who can access the case and set compliance boundaries

Depending on your organizational practice, you need to control who can access the eDiscovery case used to investigate a data spillage incident and set up compliance boundaries. The easiest way to do this is to add investigators as members of an existing role group in the Security & Compliance Center and then add the role group as a member of the eDiscovery case. For information about the built-in eDiscovery role groups and how to add members to an eDiscovery case, see Assign eDiscovery permissions.

You can also create a new role group that aligns with your organizational needs. For example, you might want a group of data spillage investigators in the organization to access and collaborate on all data spillage cases. You can do this by creating a "Data Spillage Investigator" role group, assigning the appropriate roles (Export, RMS Decrypt, Review, Preview, Compliance Search, and Case Management), adding the data spillage investigators to the role group, and then adding the role group as a member of the data spillage eDiscovery case. See Set up compliance boundaries for eDiscovery investigations in Office 365 for detailed instructions on how to do this.

Step 2: Create an eDiscovery case

An eDiscovery case provides an effective way to manage your data spillage investigation. You can add members to the role group that you created in Step 1, add the role group as a member of new a eDiscovery case, perform iterative searches to find the spilled data, export a report to share, track the status of the case, and then refer back to the details of the case if needed. Consider establishing a naming convention for eDiscovery cases used for data spillage incidents, and provide as much information as you can in the case name and description so you can locate and refer to in the future if necessary.

To create a new case, you can use eDiscovery in the security and compliance center. See "Create a new case" in eDiscovery cases.

Step 3: Search for the spilled data

Now that you've created a case and managed access, you can use the case to iteratively search to find the spilled data and identify the mailboxes that contain the spilled data. You will use the same search query that you used to find the email messages to delete those same messages in Step 7.

To create a content search associated with an eDiscovery case, see "Create and run a Content Search associated with a case" in eDiscovery cases.

Important: The keywords that you use in the search query may contain the actual spilled data that you're

searching for. For example, if you searching for documents containing a social security number and you use the it as search keyword, you must delete the query afterwards to avoid further spillage. See Deleting the search query in Step 8.

Step 4: Review and validate case findings

After you create a content search, you need to review and validate that the search results and verify that they consist only of the email messages that must be deleted. In a content search, you can preview a random sampling of 1,000 email messages without exporting the search results to avoid further data spillage. You can read more about the preview limitations at Limits for Content Search.

If you have more than 1,000 mailboxes or more than 100 email messages per mailbox to review, you can divide the initial search into multiple searches by using additional keywords or conditions such as date range or sender/recipient and review the results of each search individually. Make sure to note down all search queries to use when you delete messages in Step 7.

If a custodian or end user is assigned an Office 365 E5 license, you can examine up to 10,000 search results at once using Advanced eDiscovery. If there are more than 10,000 email messages to review, you can divide the search query by date range and review each result individually as search results are sorted by date. In Advanced eDiscovery, you can tag search results using the **Label as** feature in the preview panel and filter the search result by the tag you labeled. This is helpful when you collaborate with a secondary reviewer. By using additional analytics tools in Advanced eDiscovery, such as optical character recognition, email threading, and predictive coding, you can quickly process and review thousands of messages and tag them for further review. See Quick setup for Advanced eDiscovery.

When you find an email message that contains spilled data, check the recipients of the message to determine if it was shared externally. To further trace an message, you can collect sender information and date range so you can use the message trace logs, which is described in Step 5.

After you verified the search results, you may want to share your findings with others for a secondary review. People who you assigned to the case in Step 1 can review the case content in both eDiscovery and Advanced eDiscovery and approve case findings. You can also generate a report without exporting the actual content. You can also use this same report as a proof of deletion, which is described in Step 8.

To generate a statistical report:

1. Go to the **Search** page in the eDiscovery case, and click the search that you want to generate a report for.
2. On the flyout page, click **More > Export report**.

   The Export report page is displayed.

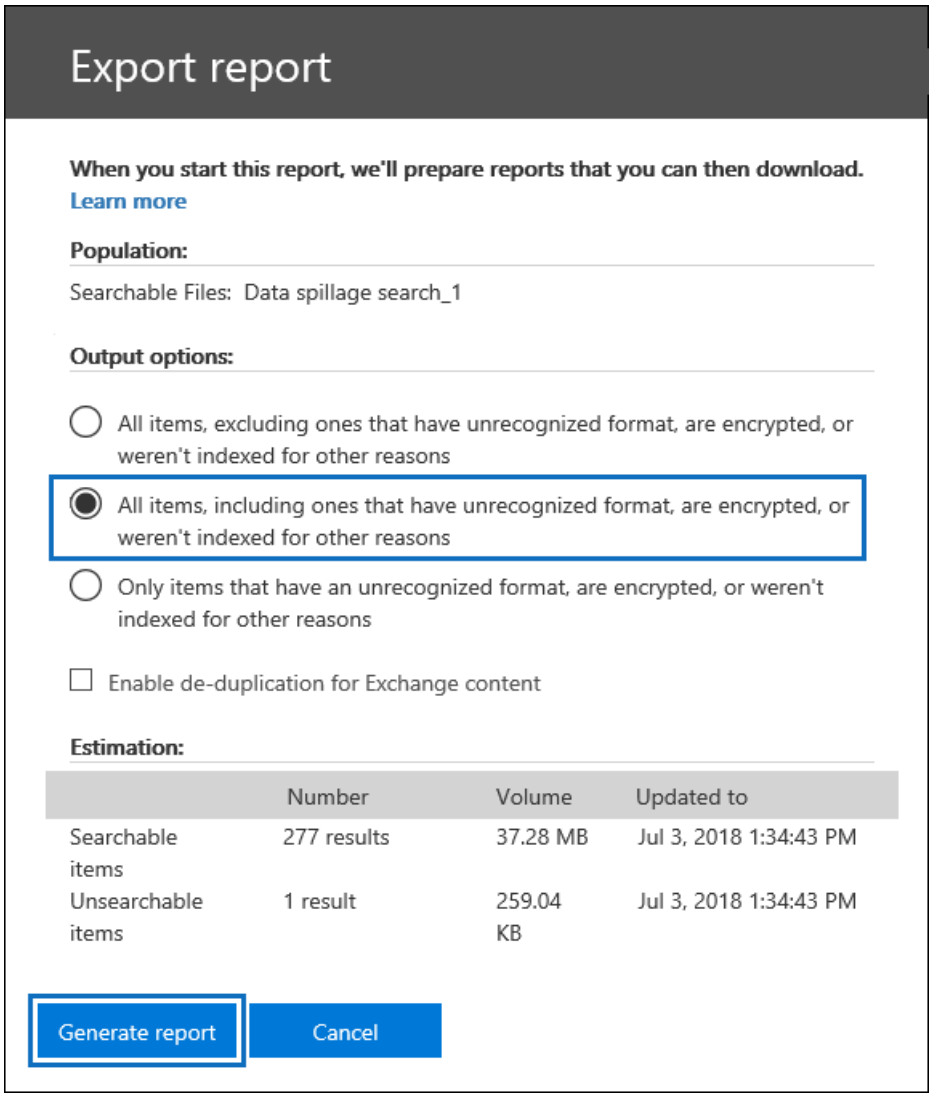

3. Select **All items, including ones that have unrecognized format, are encrypted, or weren't indexed for other reasons** and then click **Generate report**.

4. In the eDiscovery case, click **Export** to display the list of export jobs. You may have to click **Refresh** to update the list to display the export job you just created.

5. Click the export job, and then click **Download** report on the flyout page.

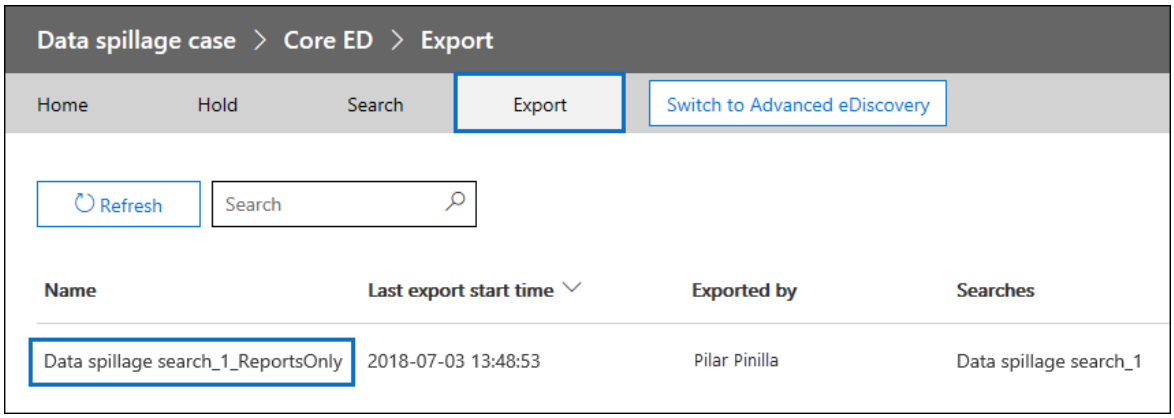

The **Export Summary** report contains the number of locations found with results and the size of the search results. You can use this to compare with the report generated after deletion and provide as a proof of deletion. The **Results** report contains a more detailed summary of the search results, including the subject, sender, recipients, if the email was read, dates, and size of each message. If any of the details in this report contains that actual spilled data, be sure to permanently delete the Results.csv file when the investigation is complete.

For more information about exporting reports, see Export a Content Search report.

Step 5: Use message trace log to check how spilled data was shared

To further investigate if email with spilled data was shared, you can optionally query the message trace logs with the sender information and the date range information that you gathered in Step 4. Note that the retention period for message trace is 30 days for real time data and 90 days for historical data.

You can use Message trace in the security and compliance center or use the corresponding cmdlets in Exchange Online PowerShell. It's important to note that message tracing doesn't offer full guarantees on the completeness of data returned. For more information about using Message trace, see:

- Message trace in the Security & Compliance Center
- New Message Trace in Security & Compliance Center

Step 6: Prepare the mailboxes

After you review and validate that the search results contains only the messages that must be deleted, you need to collect a list of the email addresses of the impacted mailboxes to use in Step 7 when you delete the spilled data. You may also have to prepare the mailboxes before you can permanently delete email messages depending on whether single item recovery is enabled on the mailboxes that contain the spilled data or if any of those mailboxes are on hold.

Get a list of addresses of mailboxes with spilled data

There are two ways to collect a list of email addresses of mailboxes with spilled data.

Option 1: Get a list of addresses of mailboxes with spilled data

1. Open the eDiscovery case, go to the **Search** page and select the appropriate content search.
2. On the flyout page, click **View results**.
3. In the **Individual results** drop down list, click **Search statistics**.
4. In the **Type** drop down list, click **Top locations**.

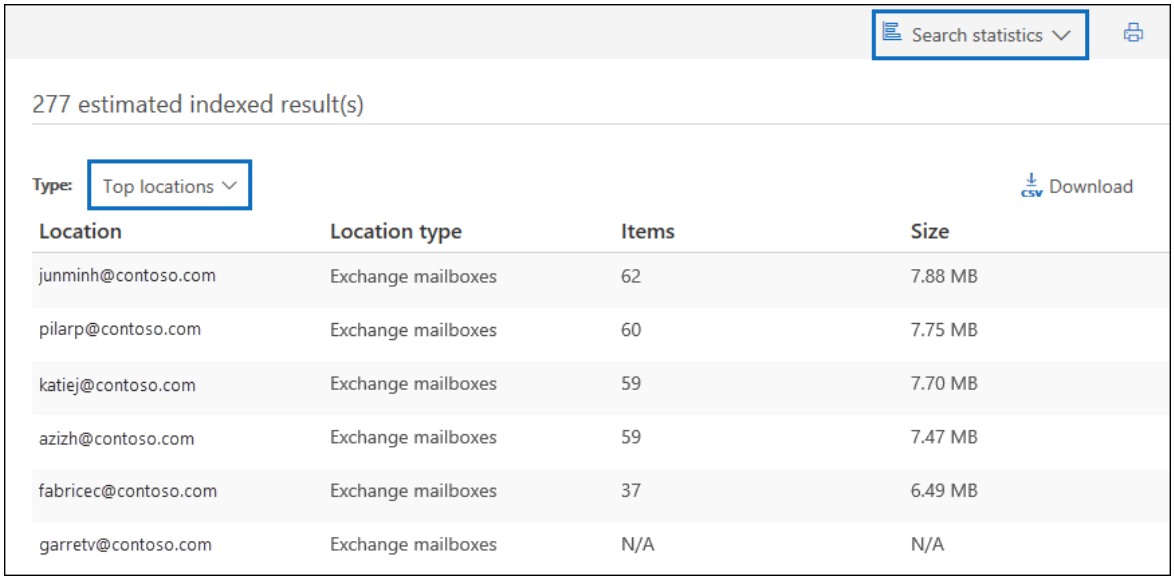

   A list of mailboxes that contain search results is displayed. The number of items in each mailbox that match the search query is also displayed.

5. Copy the information in the list and save it to a file or click **Download** to download the information to a CSV file.

Option 2: Get mailbox locations from the export report

Open the Export Summary report that you downloaded in Step 4. In the first column in the report, the email address of each mailbox is listed under **Locations**.

Prepare the mailboxes so you can delete the spilled data

If single item recovery is enabled or if a mailbox is placed on hold, a permanently deleted (purged) message will be retained in Recoverable Items folder. So before you can purge spilled data, you need to check the existing mailbox configurations and disable single item recovery and remove any hold or retention policy. Keep in mind that you can prepare one mailbox at a time, and then run the same command on different mailboxes or create a PowerShell script to prepare multiple mailboxes at the same time.

- See "Step 1: Collect information about the mailbox" in Delete items in the Recoverable Items folder of cloud-based mailboxes on hold for instructions about how to check if single item recovery is enabled or if the mailbox is placed on hold or it's assigned to a retention policy.
- See "Step 2: Prepare the mailbox" in Delete items in the Recoverable Items folder of cloud-based mailboxes on hold for instructions about disabling single item recovery.
- See "Step 3: Remove all holds from the mailbox" in Delete items in the Recoverable Items folder of cloud-based mailboxes on hold for instructions about how to remove a hold or retention policy from a mailbox.
- See "Step 4: Remove the delay hold from the mailbox" in Delete items in the Recoverable Items folder of cloud-based mailboxes on hold for instructions about removing the delay hold that is placed on the mailbox after any type of hold is removed.

> **IMPORTANT**
> Check with your records management or legal departments before removing a hold or retention policy. Your organization may have a policy that defines whether a mailbox on hold or a data spillage incident takes priority.

Be sure to revert the mailbox to previous configurations after you verify that the spilled data has been permanently deleted. See the details in Step 7.

Step 7: Permanently delete the spilled data

Using the mailbox locations that you collected and prepared in Step 6 and the search query that was created and refined in Step 3 to find email messages that contain the spilled data, you can now permanently delete the spilled data. As previously explained, to delete messages, you have to be a member of the Organization Management role group or be assigned the Search And Purge management role. For information about adding users to a role group, see Assign eDiscovery permissions in the Security & Compliance Center.

To delete the spilled messages, see steps 2 & 3 in Search for and delete email messages

> **IMPORTANT**
> Email items in a review set in an Advanced eDiscovery case can't be deleted by using the procedures in this article. That's because items in a review set are copies of items in the live service that are copied and stored in an Azure Storage location. This means they won't be returned by a content search that you create in Step 3. To delete items in a review set, you have to delete the Advanced eDiscovery case that contains the review set. For more information, see Close or delete an Advanced eDiscovery case.

Step 8: Verify, provide a proof of deletion, and audit

The final step in the workflow to manage a data spillage incident is to verify that the spilled data was permanently removed from the mailbox by going to the eDiscovery case and re-running the same search query

that was used to delete that data to confirm that no results are returned. After you confirm the spilled data has been permanently removed, you can export a report and include it (along with the original report) as a proof of deletion. Then you can close the case, which will allow you to re-open it if you have refer to it in the future. Additionally, you can also revert mailboxes to their previous state, delete the search query used to find the spilled data, and search for auditing records of tasks performed when managing the data spillage incident.

Reverting the mailboxes to their previous state

If you changed any mailbox configuration in Step 6 to prepare the mailboxes before the spilled data was deleted, you will need to revert them to their previous state. See "Step 6: Revert the mailbox to its previous state" in Delete items in the Recoverable Items folder of cloud-based mailboxes on hold.

Deleting the search query

If the keywords in the search query that you created and used in Step 3 contains some of all of the actual spilled data, you should delete the search query to prevent further data spillage.

1. In the security and compliance center, open the eDiscovery case, go to the **Search** page, and select the appropriate content search.
2. On the flyout page, click **Delete**.

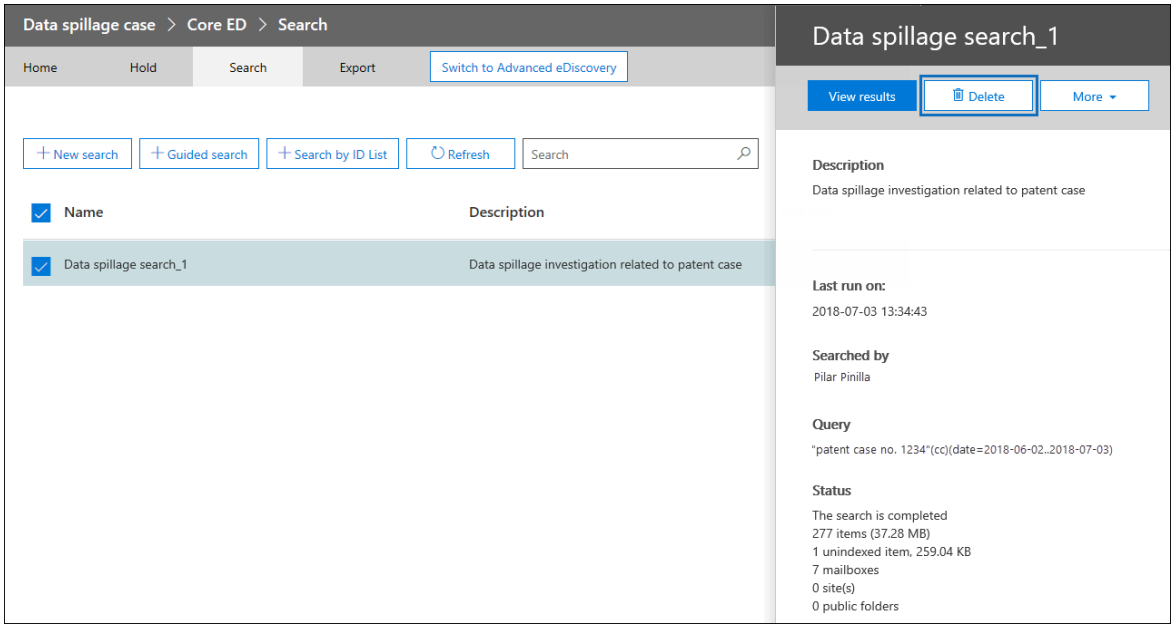

Auditing the data spillage investigation process

You can search the audit log for the eDiscovery activities that were performed during the investigation. You can also search the audit log to return the audit records for the **New-ComplianceSearchAction -Purge** command that you ran in Step 7 to delete the spilled data. For more information, see:

- Search the audit log
- Search for eDiscovery activities in the audit log

Preserve Bcc and expanded distribution group recipients for eDiscovery

2/18/2021 • 5 minutes to read • Edit Online

In-Place Hold, Litigation Hold, and Microsoft 365 retention policies (created in the Security & Compliance Center) allow you to preserve mailbox content to meet regulatory compliance and eDiscovery requirements. Information about recipients directly addressed in the To and Cc fields of a message is included in all messages by default. But your organization may require the ability to search for and reproduce details about all recipients of a message. This includes:

- **Recipients addressed using the Bcc field of a message:** Bcc recipients are stored in the message in the sender's mailbox, but not included in headers of the message delivered to recipients.
- **Expanded distribution group recipients:** Recipients who receive the message because they're members of a distribution group to which the message was addressed, either in the To, Cc or Bcc fields.

Exchange Online and Exchange Server 2013 (Cumulative Update 7 and later versions) retain information about Bcc and expanded distribution group recipients. You can search for this information by using an In-Place eDiscovery search in the Exchange admin center (EAC) or a Content Search in the Security & Compliance Center.

How Bcc recipients and expanded distribution group recipients are preserved

As stated earlier, information about Bcc'ed recipients is stored with the message in the sender's mailbox. This information is indexed and available to eDiscovery searches and holds.

Information about expanded distribution group recipients is stored with the message after you place a mailbox on In-Place Hold or Litigation Hold. In Office 365, this information is also stored when a Microsoft 365 retention policy is applied to a mailbox. Distribution group membership is determined at the time the message is sent. The expanded recipients list stored with the message is not impacted by changes to membership of the group after the message is sent.

INFORMATION ABOUT...	IS STORED IN...	IS STORED BY DEFAULT?	IS ACCESSIBLE TO...
To and Cc recipients	Message properties in the sender and recipients' mailboxes.	Yes	Sender, recipients, and compliance officers
Bcc recipients	Message property in the sender's mailbox.	Yes	Sender and compliance officers
Expanded distribution group recipients	Message properties in the sender's mailbox.	No. Expanded distribution group recipient information is stored after a mailbox is placed on In-Place Hold or Litigation Hold, or assigned to a Microsoft 365 retention policy.	Compliance officers

Searching for messages sent to Bcc and expanded distribution group

recipients

When searching for messages sent to a recipient, eDiscovery search results now include messages sent to a distribution group that the recipient is a member of. The following table shows the scenarios where messages sent to Bcc and expanded distribution group recipients are returned in eDiscovery searches.

Scenario 1: John is a member of the US-Sales distribution group. This table shows eDiscovery search results when Bob sends a message to John directly or indirectly via a distribution group.

WHEN YOU SEARCH BOB'S MAILBOX FOR MESSAGES SENT...	AND THE MESSAGE IS SENT WITH...	RESULTS INCLUDE MESSAGE?
To:John	John on TO	Yes
To:John	US-Sales on TO	Yes
To:US-Sales	US-Sales on TO	Yes
Cc:John	John on CC	Yes
Cc:John	US-Sales on CC	Yes
Cc:US-Sales	US-Sales on CC	Yes

Scenario 2: Bob sends an email to John (To/Cc) and Jack (Bcc directly, or indirectly via a distribution group). The table below shows eDiscovery search results.

WHEN YOU SEARCH...	FOR MESSAGES SENT...	RESULTS INCLUDE MESSAGE?	NOTES
Bob's mailbox	To/Cc:John	Yes	Presents an indication that Jack was Bcc'ed.
Bob's mailbox	Bcc:Jack	Yes	Presents an indication that Jack was Bcc'ed.
Bob's mailbox	Bcc:Jack (via distribution group)	Yes	List of members of the Bcc'ed distribution group, expanded when the message was sent, is visible in eDiscovery search preview, export, and logs.
John's mailbox	To/Cc:John	Yes	No indication of Bcc recipients.
John's mailbox	Bcc:Jack (directly or via distribution group)	No	Bcc information is not stored in the message delivered to recipients. You must search the sender's mailbox.
Jack's mailbox	To/Cc:John (directly or via distribution group)	Yes	To/Cc information is included in message delivered to all recipients.

WHEN YOU SEARCH...	FOR MESSAGES SENT...	RESULTS INCLUDE MESSAGE?	NOTES
Jack's mailbox	Bcc:Jack (directly or via distribution group)	No	Bcc information is not stored in the message delivered to recipients. You must search the sender's mailbox.

Frequently asked questions

Q. When and where is Bcc recipient information stored?

A. Bcc recipient information is preserved by default in the original message in sender's mailbox. If the Bcc recipient is a distribution group, distribution group membership is only expanded if the sender's mailbox is on hold or assigned to a Microsoft 365 retention policy.

Q. When and where is the list of expanded distribution group recipients stored?

A. Group membership is expanded at the time the message is sent. The list of expanded distribution group members is stored in the original message in the sender's mailbox. The sender's mailbox must be on In-Place Hold, Litigation Hold, or assigned to a Microsoft 365 retention policy.

Q. Can the To/Cc recipients see which recipients were Bcc'ed?

A. No. This information is not included in message headers, and isn't visible to To/Cc recipients. The sender can see the Bcc field stored in the original message stored in their mailbox. Compliance officers can see this information when searching the sender's mailbox.

Q. How can I ensure that expanded distribution group recipients are always preserved?

A. To ensure that expanded distribution group members are always preserved with a message, Place all mailboxes on hold or create an organization-wide Microsoft 365 retention policy.

Q. Which types of groups are supported?

A. Distribution groups, mail-enabled security groups, and dynamic distribution groups are supported.

Q. Is there a limit on the number of distribution group recipients that are expanded and stored in the message?

A. Up to 10,000 members of a distribution group is preserved.

Q. Are nested distribution groups supported?

A. Yes, 25 levels of nested distribution groups are expanded.

Q. Where is the Bcc and expanded distribution group recipient information visible?

A. Bcc and expanded distribution group recipients information is visible to Compliance officers when performing an eDiscovery search. Bcc and expanded distribution group recipients are included in search results copied to a Discovery mailbox or exported to a PST file and in the eDiscovery log included in search results. Bcc recipient information is also available in search preview.

Q. What happens if a member of a distribution group is hidden from the organization's global address list (GAL)?

A. There's no impact. If recipients are hidden from the GAL, they are still included in the list of recipients for the expanded distribution group.

Decryption in Microsoft 365 eDiscovery tools

2/18/2021 • 3 minutes to read • Edit Online

Encryption is an important part of your file protection and information protection strategy. Organizations of all types use encryption technology to protect sensitive content within their organization and ensure that only the right people have access to that content.

To execute common eDiscovery tasks on encrypted content, eDiscovery managers were required to decrypt email message content as it was exported from content searches, Core eDiscovery cases, and Advanced eDiscovery cases. Content encrypted with Microsoft encryption technologies wasn't available for review until after it was exported.

To make it easier to manage encrypted content in the eDiscovery workflow, Microsoft 365 eDiscovery tools now incorporate decryption of encrypted files that are attached to email messages and sent in Exchange Online. Additionally, encrypted documents stored in SharePoint Online and OneDrive for Business are decrypted in Advanced eDiscovery.

Prior to this new capability, only the content of an email message protected by rights management (and not attached files) were decrypted. Encrypted documents in SharePoint and OneDrive couldn't be decrypted during the eDiscovery workflow. Now, if a file that's encrypted with a Microsoft encryption technology is attached to an email message or located on a SharePoint or OneDrive account, those encrypted items are decrypted when the search results are prepared for preview, added to a review set in Advanced eDiscovery, and exported. This allows eDiscovery managers to view the content of encrypted email attachments and site documents when previewing search results, and review them after they have been added to a review set in Advanced eDiscovery.

Supported encryption technologies

Microsoft eDiscovery tools support items encrypted with Microsoft encryption technologies. These technologies include Office Message Encryption, Azure Rights Management, and Microsoft Information Protection (specifically sensitivity labels). For more information about Microsoft encryption technologies, see Encryption. Content encrypted by third-party encryption technologies isn't supported. For example, previewing or exporting content encrypted with non-Microsoft technologies isn't supported.

eDiscovery activities that support encrypted items

The following table identifies the supported tasks that can be performed in Microsoft 365 eDiscovery tools on encrypted files attached to email massages and encrypted documents in SharePoint and OneDrive. These supported tasks can be performed on encrypted files that match the criteria of a search. A value of `N/A` indicates the functionality isn't available in the corresponding eDiscovery tool.

EDISCOVERY TASK	CONTENT SEARCH	CORE EDISCOVERY	ADVANCED EDISCOVERY
Search for content in encrypted files in email and sites	Yes	Yes	Yes
Preview encrypted files attached to email	Yes	Yes	Yes

EDISCOVERY TASK	CONTENT SEARCH	CORE EDISCOVERY	ADVANCED EDISCOVERY
Preview encrypted documents in SharePoint and OneDrive	No	No	Yes
Review encrypted files in a review set	N/A	N/A	Yes
Export encrypted files attached to email	Yes	Yes	Yes
Export encrypted documents in SharePoint and OneDrive	No	No	Yes

Note: eDiscovery doesn't support encrypted files in SharePoint and OneDrive when a sensitivity label that applied the encryption is configured with either of the following settings:

- Users can assign permissions when they manually apply the label to a document. This is sometimes referred to as *user-defined permissions*.
- User access to the document has an expiration setting that is set to a value other than **Never**.

For more information about these settings, see the "Configure encryption settings" section in Restrict access to content by using sensitivity labels to apply encryption.

Documents encrypted with the previous settings can still be returned by an eDiscovery search. This may happen when a document property (such as the title, author, or modified date) matches the search criteria. Although these documents might be included in search results, they can't be previewed or reviewed. These documents will also remain encrypted when they're exported in Advanced eDiscovery.

Requirements for decryption in eDiscovery

You have to be assigned the RMS Decrypt role to preview, review, and export files encrypted with Microsoft encryption technologies. You also have to be assigned this role to review and query encrypted files that are added to a review set in Advanced eDiscovery.

This role is assigned by default to the eDiscovery Manager role group on the **Permissions** page in the Office 365 Security & Compliance Center. For more information about the RMS Decrypt role, see Assign eDiscovery permissions.

Collect eDiscovery diagnostic information

2/18/2021 • 3 minutes to read • Edit Online

Occasionally Microsoft Support engineers require specific information about your issue when you open a support case related to Core eDiscovery or Advanced eDiscovery. This article provides guidance on how to collect diagnostic information to help support engineers investigate and resolve issues. Typically, you don't need to collect this information until asked to do so by a Microsoft Support engineer.

> **IMPORTANT**
> The output from the cmdlets and diagnostic information described in this article may include sensitive information about litigation or internal investigations in your organization. Before sending the raw diagnostic information to Microsoft Support, you should review the information and redact any sensitive information (such as names or other information about parties to litigation or investigation) by replacing it with `xxxxxxx`. Using this method will also indicate to the Microsoft Support engineer that information was redacted.

Collect diagnostic information for Core eDiscovery

Collecting diagnostic information for Core eDiscovery is cmdlet-based, so you'll have to use Security & Compliance Center PowerShell. The following PowerShell examples will run cmdlets and then save the output to a specified text file. In most support cases, you should only have to run one of these commands.

To run the following cmdlets, connect to Security & Compliance Center PowerShell. After you're connected, run one or more of the following commands and be sure to replace placeholders with the actual object names.

After reviewing the generated text file and redacting sensitive information, send it to the Microsoft Support engineer working on your case.

> **NOTE**
> You can also run the commands in this section to collect diagnostic information for the searches and exports listed on the **Content search** page in the Microsoft 365 compliance center.

Collect information about searches

The following command collects information that's helpful when investigating issues with a Content search or a search associated with a Core eDiscovery case.

```
Get-ComplianceSearch "<Search name>" | FL > "ComplianceSearch.txt"
```

Collect information about search actions

The following command collects information to investigate problems with previewing, exporting, or purging the results of a Content search or a search associated with a Core eDiscovery case. You can identify the name of the search action by clicking an export that's listed on the **Exports** tab. To identify the names of preview and purge actions, you can run the **Get-ComplianceSearchAction** cmdlet to display a list of all actions. The format for the search action name is constructed by appending `_Preview`, `_Export`, or `_Purge` to the name of the corresponding search.

```
Get-ComplianceSearchAction "<Search action name>" | FL > "ComplianceSearchAction.txt"
```

Collect information about eDiscovery holds

When an eDiscovery hold associated with a Core eDiscovery case isn't functioning as expected, run the following command to collect information about the Case Hold Policy and associated Case Hold Rule for the eDiscovery hold. The *Case hold policy name* in the following command is the same as the name of the eDiscovery hold. You can identify this name on the **Holds** tabs in the Core eDiscovery case.

```
Get-CaseHoldPolicy "<Case hold policy name>" | %{"--CaseHoldPolicy--";$_|FL;"--CaseHoldRule--";Get-
CaseHoldRule -Policy $_.Name | FL} > "eDiscoveryCaseHold.txt"
```

Collect all case information

Sometimes, it's not apparent what information is required by Microsoft Support to investigate your issue. In this situation, you can collect all of the diagnostics information for a Core eDiscovery case. The *Core eDiscovery case name* in the following command is the same as the name of a case that's displayed on the **Core eDiscovery** page in the Microsoft 365 compliance center.

```
Get-ComplianceCase "<Core eDiscovery case name>"| %{"$($_.Name)";"`t==Searches==";Get-ComplianceSearch -Case
$_.Name | FL;"`t==Search Actions==";Get-ComplianceSearchAction -Case $_.Name |FL;"`t==Holds==";Get-
CaseHoldPolicy -Case $_.Name | %{$_|FL;"`t`t ==$($_.Name) Rules==";Get-CaseHoldRule -Policy $_.Name | FL}} >
"eDiscoveryCase.txt"
```

Collect diagnostic information for Advanced eDiscovery

The **Settings** tab in an Advanced eDiscovery case lets you quickly copy the diagnostic information for the case. The diagnostic information is saved to the clipboard so you can paste it to a text file and send to Microsoft Support.

1. Go to https://compliance.microsoft.com and then click **Show all > eDiscovery > Advanced**.
2. Select a case and then click the **Settings** tab.
3. Under **Case Information**, click **Select**.
4. On the flyout page, click **Copy diagnostic information** to copy the info to the clipboard.
5. Open a text file (in Notepad) and then paste the information in the text file.
6. Save the text file and name it something like `AeD Diagnostic Info YYYY.MM.DD` (for example, `AeD Diagnostic Info 2020.11.03`).

After reviewing the file and redacting sensitive information, send it to the Microsoft Support engineer working on your case.

Investigate, troubleshoot, and resolve common eDiscovery issues

2/18/2021 • 5 minutes to read • Edit Online

This topic covers basic troubleshooting steps you can take to identify and resolve issues you may encounter during an eDiscovery search or elsewhere in the eDiscovery process. Resolving some of these scenarios requires help from Microsoft Support. Information on when to contact Microsoft Support is included in the resolution steps.

Error/issue: Ambiguous location

If you try to add user's mailbox location to search and there are duplicate or conflicting objects with the same userID in the Exchange Online Protection (EOP) directory, you receive this error:

```
The compliance search contains the following invalid location(s):useralias@contoso.com. The location
"useralias@contoso.com" is ambiguous
```

.

Resolution

Check for duplicate users or distribution list with the same user ID.

1. Connect to Security & Compliance Center PowerShell.
2. Run the following command to retrieve all instances of the username:

   ```
   Get-Recipient <username>
   ```

   The output for 'useralias@contoso.com' would be similar to the following:

NAME	RECIPIENTTYPE
Alias, User	MailUser
Alias, User	User

3. If multiple users are returned, locate and fix the conflicting object.

Error/issue: Search fails on specific locations

An eDiscovery or content search may yield the following error:

```
This search completed with (#) errors. Would you like to retry the search on the failed locations?
```

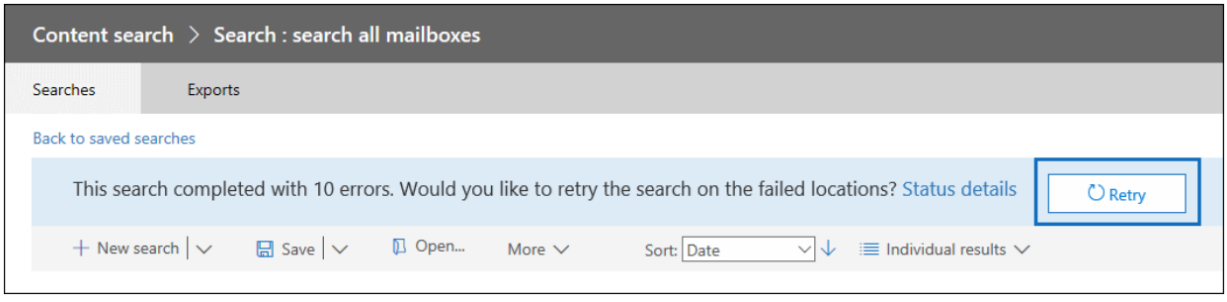

Resolution

If you receive this error, we recommend that you verify the locations that failed in the search then rerun the search only on the failed locations.

1. Connect to Security & Compliance Center PowerShell and then run the following command:

```
Get-ComplianceSearch <searchname> | FL
```

2. From the PowerShell output, view the failed locations in the errors field or from the status details in the error from the search output.
3. Retry the eDiscovery search on the failed locations only.
4. If you continue to receive these errors, see Retry failed locations for more troubleshooting steps.

Error/issue: File not found

When running an eDiscovery search that includes SharePoint Online and One Drive For Business locations, you may receive the error `File Not Found` although the file is located on the site. This error will be in the export warnings and errors.csv or skipped items.csv. This may occur if the file can't be found on the site or if the index is out of date. Here's the text of an actual error (with emphasis added).

> 28.06.2019 10:02:19_FailedToExportItem_Failed to download content. Additional diagnostic info : Microsoft.Office.Compliance.EDiscovery.ExportWorker.Exceptions.ContentDownloadTemporaryFailure: Failed to download from content 6ea52149-91cd-4965-b5bb-82ca6a3ec9be of type Document. Correlation Id: 3bd84722-937b-4c23-b61b-08d6fba9ec32. ServerErrorCode: -2147024894 ---> Microsoft.SharePoint.Client.ServerException: ***File Not Found***. at Microsoft.SharePoint.Client.ClientRequest.ProcessResponseStream(Stream responseStream) at Microsoft.SharePoint.Client.ClientRequest.ProcessResponse() --- End of inner exception stack trace ---

Resolution

1. Check location identified in the search to ensure the that the location of the file is correct and added in the search locations.
2. Use the procedures at Manually request crawling and re-indexing of a site, a library, or a list to reindex the site.

Error/issue: Search fails because recipient is not found

An eDiscovery search fails with error the `recipient not found`. This error may occur if the user object cannot be found in Exchange Online Protection (EOP) because the object has not synced.

Resolution

1. Connect to Exchange Online PowerShell.
2. Run the following command to check if the user is synced to Exchange Online Protection:

```
Get-Recipient <userId> | FL
```

3. There should be a mail user object for the user question. If nothing is returned, investigate the user object. Contact Microsoft Support if the object can't be synced.

Error/issue: Exporting search results is slow

When exporting search results from eDiscovery or Content Search in the Security and Compliance center, the download takes longer than expected. You can check to see the amount of data to be download and possibly increase the export speed.

Resolution

1. Connect to Security & Compliance Center PowerShell and then run the following command:

```
Get-ComplianceSearch <searchname> | FL
```

2. Find the amount of data to be downloaded in the SearchResults and SearchStatistics parameters.
3. Run the following command:

```
Get-ComplianceSearchAction | FL
```

4. In the results field, find the data that has been exported and view any errors encountered.
5. Check the trace.log file located in the directory that you exported the content to for any errors.
6. If you still have issues, consider dividing searches that return a large set of results into smaller searches. For example, you can use date ranges in search queries to return a smaller set of results that can be downloaded faster.

Error/issue: "Internal server error (500) occurred"

When running an eDiscovery search, if the search continually fails with error similar to "Internal server error (500) occurred", you may need rerun the search only on specific mailbox locations.

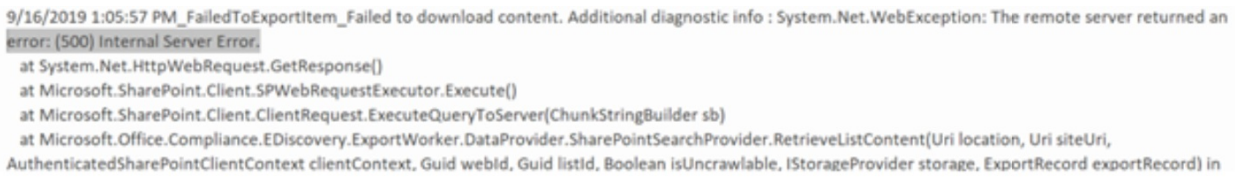

Resolution

1. Break the search into smaller searches and run the search again. Try using a smaller date range or limit the number of locations being searched.
2. Connect to Security & Compliance Center PowerShell and then run the following command:

```
Get-ComplianceSearch <searchname> | FL
```

3. Examine the output for results and errors.
4. Examine the trace.log file. It's located in the same folder that you exported the search results to.
5. Contact Microsoft Support.

Error/issue: Holds don't sync

eDiscovery Case Hold Policy Sync Distribution error. The error reads:

> "Resources: It's taking longer than expected to deploy the policy. It might take an additional 2 hours to update the final deployment status, so check back in a couple hours."

Resolution

1. Connect to Security & Compliance Center PowerShell and then run the following command for an eDiscovery case hold:

   ```
   Get-CaseHoldPolicy <policyname> - DistributionDetail | FL
   ```

   For a retention policy, run the following command:

   ```
   Get-RetentionCompliancePolicy <policyname> - DistributionDetail | FL
   ```

2. Examine the value in the DistributionDetail parameter for errors like the following:

   > Error: Resources: It's taking longer than expected to deploy the policy. It might take an additional 2 hours to update the final deployment status, so check back in a couple hours."

3. Try running the RetryDistribution parameter on the policy in question:

   For eDiscovery case holds:

   ```
   Set-CaseHoldPolicy <policyname> -RetryDistribution
   ```

   For retention policies:

   ```
   Set-RetentionCompliancePolicy <policyname> -RetryDistribution
   ```

4. Contact Microsoft Support.

Error: "The condition specified using HTTP conditional header(s) is not met"

When downloading search results using the eDiscovery Export Tool, it's possible you might receive the following error:

```
System.Net.WebException: The remote server returned an error: (412) The condition specified using HTTP
conditional header(s) is not met.
```

This is transient error, which typically occurs in the Azure Storage location.

Resolution

To resolve this issue, retry downloading the search results, which will restart the eDiscovery Export Tool.

Error/issue: Downloaded export shows no results

After a successful export, the completed download via the export tool shows zero files in the results.

Resolution

This is a client-side issue and in order to remediate it, please attempt the following steps:

1. Try using another client/machine to download.
2. Make sure to download to a local drive.
3. Make sure the virus scanner is not running.
4. Make sure that no other export is downloading to the same folder or any parent folder.

5. If the previous steps did not work, disable zipping and de-duplication.
6. If this works then the issue is due to a local virus scanner or a disk issue.

Retry a Content Search to resolve a content location error

11/2/2020 • 2 minutes to read • Edit Online

When you use Content Search in the security and compliance center to search a large number of mailboxes, you may get search errors that are similar to the error:

```
Error


The search on the following locations failed:

User1@contoso.com: Problem in processing the request. Please try again later. If you keep getting this
error, contact your admin. (CS008-009)

User2@contoso.com: Application error occurred. Please try again later. (CS012-002)
```

These errors (with error codes of CS001-002, CS003-002, CS008-009, CS012-002, and other errors of the form CS0XX-0XX) indicate that Content Search failed to search specific content locations; in this example, two mailboxes weren't searched. These errors are displayed on the status details flyout page of the Content Search.

Cause of content location errors

When searching a large number of mailboxes, the search is distributed across thousands of servers in a Microsoft datacenter. At any one time, specific servers could be in reboot state or in the process of failing over to redundant copies. In either of these cases, the Content Search's request to retrieve data will time out. In the previous example, the errors for the mailboxes that failed were the result of the search timing out.

Resolving content location errors

Restarting the search will often result in similar errors on different servers. Instead of restarting the search, click the **Retry** button that is displayed at the top of the search results page.

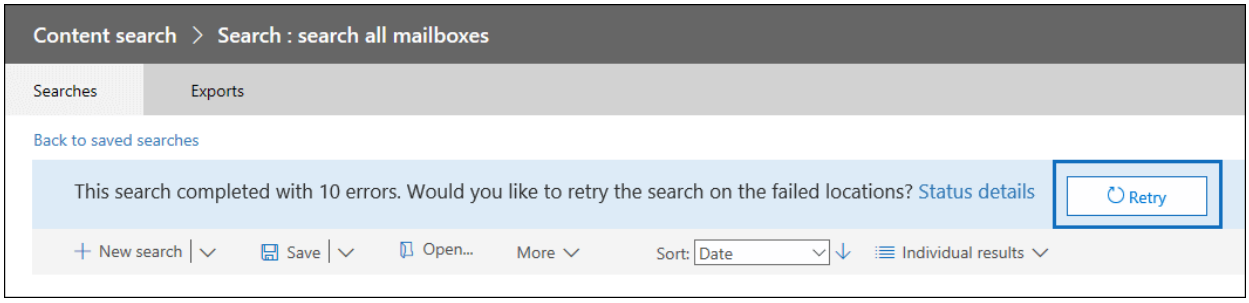

This will result in the retrying the search only for the mailboxes that failed. When you retry the search, the other results that were successfully returned are retained.

Tips to avoid content location errors

Here are some additional causes of content location errors and some tips to help you avoid them when searching large numbers of mailboxes.

- The mailbox being searched might be busy due to user activity. In this case, the search service might

throttle itself to prevent the mailbox from becoming unavailable. To avoid this, try running searches during non-business hours.

- The search query might be retrieving too much content from the mailbox. If possible, try to narrow the scope of the search by using keywords, date ranges, and search conditions.
- Too many keywords or keyword phrases when you create a search query using the keywords list. When you run a search query that uses the keywords list, the service essentially runs a separate search for each row in the keyword list so that statistics can be generated. If you're using the keywords list in search queries, minimize the number of rows in the keyword list or divide the number keywords into smaller lists and create a different search for each keyword list.

  > **NOTE**
  > To help reduce issues caused by large keyword lists, you're now limited to a maximum of 20 rows in the keyword list of a search query.

- Too many searches are being performed on the same mailbox at the same time. If possible, try to run one search at a time on any one mailbox.
- Searching too many mailboxes in a single search. The probability of content location errors increases when searching a large number of mailboxes. If possible, try to run multiple searches so that each search includes a subset of mailboxes in your organization.
- Required maintenance is being performed on the mailbox. Though this cause probably occurs infrequently, wait a little while after receiving the content location error and then retry the search.

Use the eDiscovery Export Tool in Microsoft Edge

11/2/2020 • 2 minutes to read • Edit Online

As a result of recent changes to the newest version of Microsoft Edge, ClickOnce support is no longer enabled by default. To continue using the eDiscovery Export Tool to download Content Search or eDiscovery search results, you either need to use Microsoft Internet Explorer or enable ClickOnce support in the newest version of Microsoft Edge.

Enable ClickOnce support in Microsoft Edge

1. In Microsoft Edge, go to **edge://flags/#edge-click-once**.
2. If the existing value is set to **Default** or **Disabled** in the dropdown list, change it to **Enabled**.

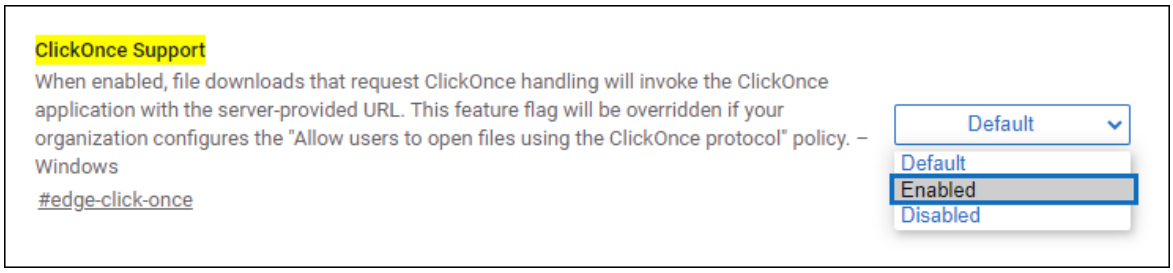

3. Scroll down to the bottom of the browser window and click **Restart** to restart Edge.

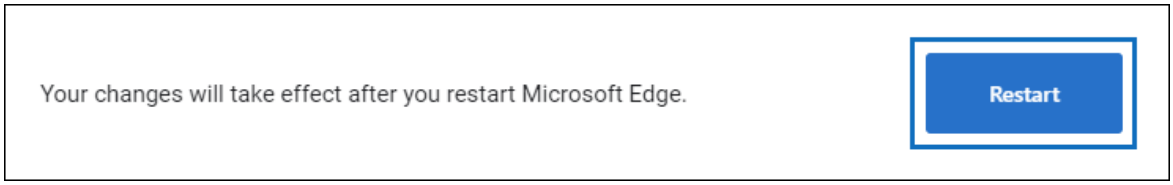

Note: Organizations can use Group Policy to disable ClickOnce support. To check if there is an organizational policy for ClickOnce support, go to **edge://policy**. The following screenshot shows that ClickOnce is enabled across the entire organization. If this policy value is set to **false**, you will need to contact an admin in your organization.

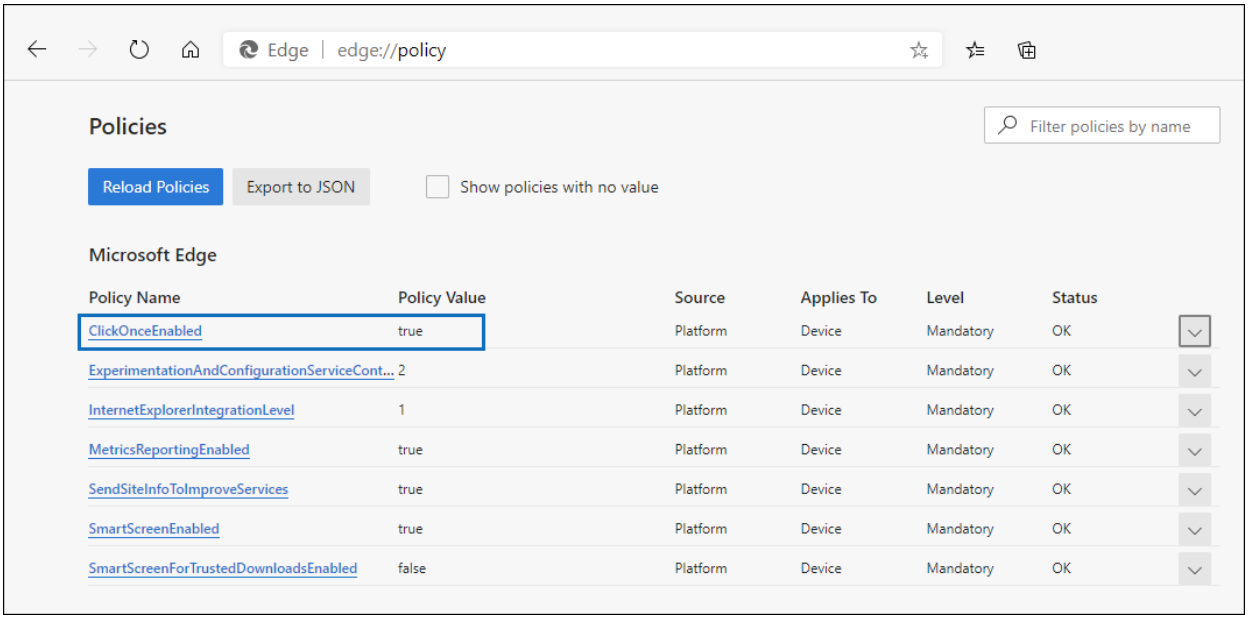

Install and run the eDiscovery Export Tool

1. Click **Download results** on the flyout page of an export in Content Search or an eDiscovery case.

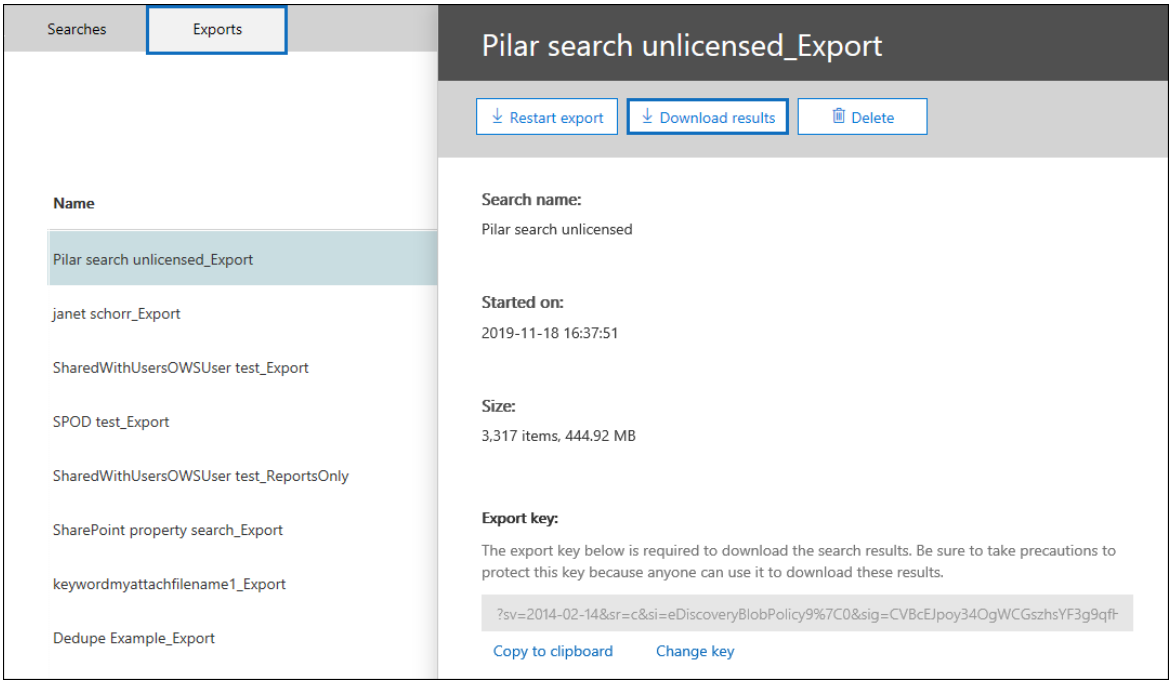

2. You will be prompted with a confirmation to launch the tool, Click **Open**.

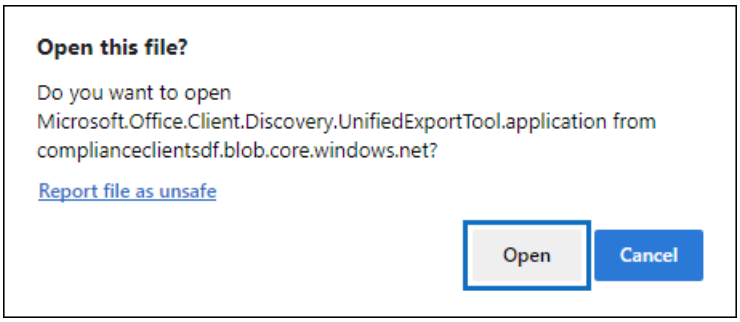

If the eDiscovery Export Tool isn't installed, you will be prompted with a Security Warning,

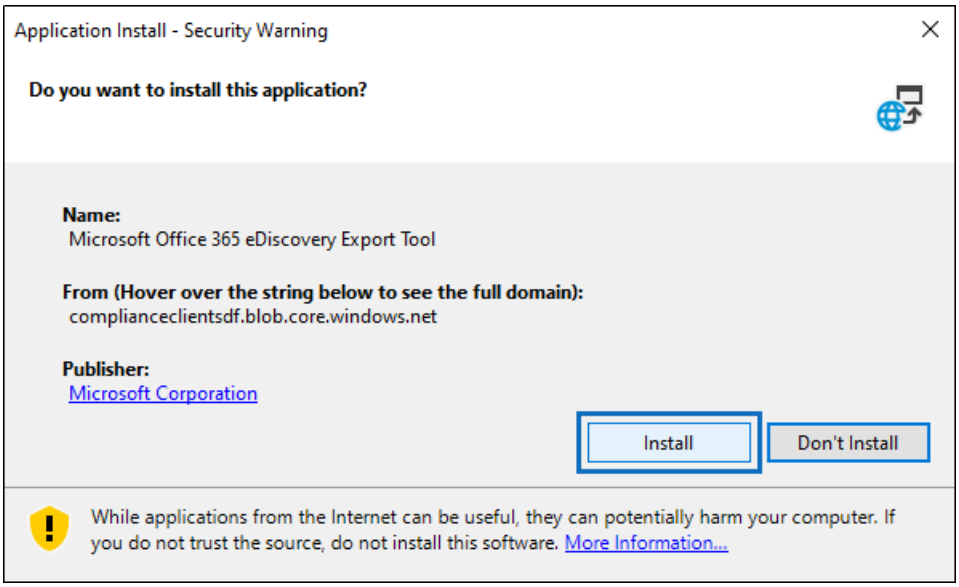

3. Click **Install**. After it's installed, the export tool will launch automatically.

For more information, see the following topics:

- Export Content Search results

- How to enable experiment flags in Microsoft Edge

Search for eDiscovery activities in the audit log

11/2/2020 • 18 minutes to read • Edit Online

Content Search and eDiscovery-related activities (for Core eDiscovery and Advanced eDiscovery) that are performed in Security & Compliance Center or by running the corresponding PowerShell cmdlets are logged in the audit log. Events are logged when administrators or eDiscovery managers (or any user assigned eDiscovery permissions) perform the following Content Search and Core eDiscovery tasks in the Security & Compliance Center:

- Creating and managing Core and Advanced eDiscovery cases
- Creating, starting, and editing Content Searches
- Performing Content Search actions, such as previewing, exporting, and deleting search results
- Managing custodians and review sets in Advanced eDiscovery
- Configuring permissions filtering for Content Search
- Managing the eDiscovery Administrator role

> **IMPORTANT**
> The activities described in this article are only the result of eDiscovery tasks performed by using the Security & Compliance Center. eDiscovery tasks that were performed by using the In-Place eDiscovery tool in Exchange Online or the eDiscovery Center in SharePoint Online aren't included.

For more information about searching the audit log, the permissions that are required, and exporting search results, see Search the audit log in the Security & Compliance Center.

How to search for and view eDiscovery activities

Currently, you have to do a few specific things to view eDiscovery activities in the audit log. Here's how.

1. Go to https://protection.office.com.
2. Sign in using your work or school account.
3. In the left pane, click **Search**, and then click **Audit log search**.
4. In the **Activities** drop-down list, under **eDiscovery activities** or **Advanced eDiscovery activities**, click one or more activities to search for.

   > **NOTE**
   > The **Activities** drop-down list also includes a group of activities named **eDiscovery cmdlet activities** that will return records from the cmdlet audit log.

5. Select a date and time range to display eDiscovery events that occurred within that period.
6. In the **Users** box, select one or more users to display search results for. Leave this box blank to return entries for all users.
7. Click **Search** to run the search using your search criteria.

8. After the search results are displayed, you can click **Filter results** to filter or sort the resulting activity records. Unfortunately, you can't use filtering to explicitly exclude certain activities.

9. To view details about an activity, click the activity record in the list of search results.

   A **Details** fly out page is displayed that contains the detailed properties from the event record. To display additional details, click **More information**. For a description of these properties, see the Detailed properties for eDiscovery activities section.

10. If desired, you can export the audit log search results to a CSV file, and then use the Excel Power Query feature to format and filter these records. For more information, see Export, configure, and view audit log records.

eDiscovery activities

The following table describes the Content Search and Core eDiscovery activities that are logged when an administrator or eDiscovery manager performs an eDiscovery-related activity using the Security & Compliance Center or running the corresponding cmdlet in Security & Compliance Center PowerShell. Note also that some activities performed in Advanced will be returned when you search for activities in this list.

> **NOTE**
> The eDiscovery activities described in this section provide similar information to the eDiscovery cmdlet activities described in the next section. We recommend that you use the eDiscovery activities described in this section because they will appear in the audit log search results within 30 minutes. It takes up to 24 hours for the eDiscovery cmdlet activities to appear in audit log search results.

FRIENDLY NAME	OPERATION	CORRESPONDING CMDLET	DESCRIPTION
Added member to eDiscovery case	CaseMemberAdded	Add-ComplianceCaseMember	A user was added as a member of an eDiscovery case. As a member of a case, a user can perform various case-related tasks depending on whether they have been assigned the necessary permissions.
Changed content search	SearchUpdated	Set-ComplianceSearch	An existing content search was changed. Changes can include adding or removing content locations or editing the search query.
Changed eDiscovery administrator membership	CaseAdminUpdated	Update-eDiscoveryCaseAdmin	The list of eDiscovery Administrators in your organization was changed. This activity is logged when the list of eDiscovery Administrators is replaced with a group of new users. If a single user is added or removed, the CaseAdminAdded operation is logged.

FRIENDLY NAME	OPERATION	CORRESPONDING CMDLET	DESCRIPTION
Changed eDiscovery case	CaseUpdated	Set-ComplianceCase	An eDiscovery case was changed. Changes include closing an open case or reopening a closed case.
Changed eDiscovery case membership	CaseMemberUpdated	Update-ComplianceCaseMember	The membership list of an eDiscovery case was changed. This activity is logged when all members are replaced with a group of new users. If a single member is added or removed, CaseMemberAdded or CaseMemberRemoved operation is logged.
Changed search permissions filter	SearchPermissionUpdated	Set-ComplianceSecurityFilter	A search permissions filter was changed.
Changed search query for eDiscovery case hold	HoldUpdated	Set-CaseHoldRule	A query-based hold associated with an eDiscovery case was changed. Possible changes include editing the query or date range for a query-based hold.
Content search preview item downloaded	PreviewItemDownloaded	N/A	A user downloaded an item to their local computer (by clicking the **Download original item** link) when previewing search results.
Content search preview item listed	PreviewItemListed	N/A	A user clicked **Preview search results** to display the preview search results page, which lists up to 1000 items from the results of a Content Search.
Content search preview item viewed	PreviewItemRendered	N/A	An eDiscovery manager viewed an item by clicking it when previewing search results.
Created content search	SearchCreated	New-ComplianceSearch	A new content search was created.
Created eDiscovery administrator	CaseAdminAdded	Add-eDiscoveryCaseAdmin	A user was added as an eDiscovery Administrator in the organization.

FRIENDLY NAME	OPERATION	CORRESPONDING CMDLET	DESCRIPTION
Created eDiscovery case	CaseAdded	New-ComplianceCase	An eDiscovery case was created. When a case is created, you only have to give it a name. Other case-related tasks such as adding members, creating holds, and creating content searches associated with the case result in additional events being logged.
Created search permissions filter	SearchPermissionCreated	New-ComplianceSecurityFilter	A search permissions filter was created.
Created search query for eDiscovery case hold	HoldCreated	New-CaseHoldRule	A query-based hold associated with an eDiscovery case was created.
Deleted content search	SearchRemoved	Remove-ComplianceSearch	An existing content search was deleted.
Deleted eDiscovery administrator	CaseAdminRemoved	Remove-eDiscoveryCaseAdmin	An eDiscovery Administrator was deleted from your organization.
Deleted eDiscovery case	CaseRemoved	Remove-ComplianceCase	An eDiscovery case was deleted. Any hold associated with the case has to be removed before the case can be deleted.
Deleted search permissions filter	SearchPermissionRemoved	Remove-ComplianceSecurityFilter	A search permissions filter was deleted.
Deleted search query for eDiscovery case hold	HoldRemoved	Remove-CaseHoldRule	A query-based hold associated with an eDiscovery case was deleted. Removing the query from the hold is often the result of deleting a hold. When a hold or a hold query is deleted, the content locations that were on hold are released.
Downloaded export of content search	SearchExportDownloaded	N/A	A user downloaded the results of a content search to their local computer. A **Started export of content search** activity has to be initiated before search results can be downloaded.
Previewed results of content search	SearchPreviewed	N/A	A user previewed the results of a content search.

FRIENDLY NAME	OPERATION	CORRESPONDING CMDLET	DESCRIPTION
Purged results of content search	SearchResultsPurged	New-ComplianceSearchAction	A user purged the results of a Content Search by running the **New-ComplianceSearchAction -Purge** command.
Removed analysis of content search	RemovedSearchResultsSentToZoom	Remove-ComplianceSearchAction	A content search prepare action (to prepare search results for Advanced eDiscovery) was deleted. If the preparation action was less than two weeks old, the search results that were prepared for Advanced eDiscovery were deleted from the Microsoft Azure storage area. If the preparation action was older than 2 weeks, then this event indicates that only the corresponding preparation action was deleted.
Removed export of content search	RemovedSearchExported	Remove-ComplianceSearchAction	A content search export action was deleted. If the export action was less than two weeks old, the search results that were uploaded to the Microsoft Azure storage area were deleted. If the export action was older than 2 weeks, then this event indicates that only the corresponding export action was deleted.
Removed member from eDiscovery case	CaseMemberRemoved	Remove-ComplianceCaseMember	A user was removed as a member of an eDiscovery case.
Removed preview results of content search	RemovedSearchPreviewed	Remove-ComplianceSearchAction	A content search preview action was deleted.
Removed purge action performed on content search	RemovedSearchResultsPurged	Remove-ComplianceSearchAction	A content search purge action was deleted.
Removed search report	SearchReportRemoved	Remove-ComplianceSearchAction	A content search export report action was deleted.
Started analysis of content search	SearchResultsSentToZoom	New-ComplianceSearchAction	The results of a content search were prepared for analysis in Advanced eDiscovery.

FRIENDLY NAME	OPERATION	CORRESPONDING CMDLET	DESCRIPTION
Started content search	SearchStarted	Start-ComplianceSearch	A content search was started. When you create or change a content search by using the Security & Compliance Center GUI, the search is automatically started. If you create or change a search by using the **New-ComplianceSearch** or **Set-ComplianceSearch** cmdlet, you have to run the **Start-ComplianceSearch** cmdlet to start the search.
Started export of content search	SearchExported	New-ComplianceSearchAction	A user exported the results of a content search.
Started export report	SearchReport	New-ComplianceSearchAction	A user exported a content search report.
Stopped content search	SearchStopped	Stop-ComplianceSearch	A user stopped a content search.
(none)	CaseViewed	Get-ComplianceCase	A user viewed the list of cases on the **eDiscovery** page in the security and compliance center or by running the cmdlet.
(none)	SearchViewed	Get-ComplianceSearch	A user viewed the list on content searches (listed on the **Searches** tab) in the security and compliance center or by running the cmdlet. This activity is also logged when a user views the list of content searches associated with an eDiscovery case (by clicking the **Searches** tab in a case) or by running the **Get-ComplianceSearch -Case** command.
(none)	ViewedSearchExported	Get-ComplianceSearchAction -Export	A user viewed the list of content search export jobs (listed on the **Exports** tab) in the security and compliance center or by running the cmdlet. This activity is also logged when a user views the list of export jobs in an eDiscovery case (listed on the **Exports** tab in a case) or by running the **Get-ComplianceSearchAction -Case -Export** command.

FRIENDLY NAME	OPERATION	CORRESPONDING CMDLET	DESCRIPTION
(none)	ViewedSearchPreviewed	Get-ComplianceSearchAction -Preview	A user previews the results of a content search in the security and compliance center or by running the cmdlet.

Advanced eDiscovery activities

The following table describes the Advanced eDiscovery activities logged in the audit log. These activities (in addition to relevant eDiscovery activities can be used to help you track the progression of activity in an Advanced eDiscovery case.

FRIENDLY NAME	OPERATION	DESCRIPTION
Added data to another review set	AddWorkingSetQueryToWorkingSet	User added documents from one review set to a different review set.
Added data to review set	AddQueryToWorkingSet	User added the search results from a content search associated with an Advanced eDiscovery case to a review set.
Added non-Microsoft 365 data to review set	AddNonOffice365DataToWorkingSet	User added non-Microsoft 365 data to a review set.
Added remediated documents to review set	AddRemediatedData	User uploads documents that had indexing errors that were fixed to a review set.
Analyzed data in review set	RunAlgo	User ran analytics on the documents in a review set.
Annotated document in review set	AnnotateDocument	User annotated a document in a review set. Annotation includes redacting content in a document.
Compared load sets	LoadComparisonJob	User compared two different load sets in a review set. A load set is when data from a content search that associated with the case is added to a review set.
Converted redacted documents to PDF	BurnJob	User converted all the redacted documents in a review set to PDF files.
Created review set	CreateWorkingSet	User created a review set.
Created review set search	CreateWorkingSetSearch	User created a search query that searches the documents in a review set.

FRIENDLY NAME	OPERATION	DESCRIPTION
Created tag	CreateTag	User created a tag group in a review set. A tag group can contain one or more child tags. These tags are then used to tag documents in the review set.
Deleted review set search	DeleteWorkingSetSearch	User deleted a search query in a review set.
Deleted tag	DeleteTag	User deleted a tag or a tag group in a review set.
Downloaded document	DownloadDocument	User downloaded a document from a review set.
Edited tag	UpdateTag	User changed a tag in a review set.
Exported documents from review set	ExportJob	User exported documents from a review set.
Modified case setting	UpdateCaseSettings	User modified the settings for a case. Case settings include case information, access permissions, and settings that control search and analytics behavior.
Modified review set search	UpdateWorkingSetSearch	User edited a search query in a review set.
Previewed review set search	PreviewWorkingSetSearch	User previewed the results of a search query in a review set.
Remediated error documents	ErrorRemediationJob	User fixes files that contained indexing errors.
Tagged document	TagFiles	User tags a document in a review set.
Tagged results of a query	TagJob	User tags all of the documents that match the criteria of search query in a review set.
Viewed document in review set	ViewDocument	User viewed a document in a review set.

eDiscovery cmdlet activities

The following table lists the cmdlet audit log records that are logged when an administrator or user performs an eDiscovery-related activity by using the Security & Compliance Center or by running the corresponding cmdlet in remote PowerShell that's connected to your organization's Security & Compliance Center. The detailed information in the audit log record is different for the cmdlet activities listed in this table and the eDiscovery activities described in the previous section.

As previously stated, it takes up to 24 hours for eDiscovery cmdlet activities to appear in the audit log search results.

TIP

The cmdlets in the **Operation** column in the following table are linked to the corresponding cmdlet help topic on TechNet. Go to the cmdlet help topic for a description of the available parameters for each cmdlet. The parameter and the parameter value that were used with a cmdlet are included in the audit log entry for each eDiscovery cmdlet activity that's logged.

FRIENDLY NAME	OPERATION (CMDLET)	DESCRIPTION
Created hold in eDiscovery case	New-CaseHoldPolicy	A hold was created for an eDiscovery case. A hold can be created with or without specifying a content source. If content sources are specified, they'll be identified in the audit log entry.
Deleted hold from eDiscovery case	Remove-CaseHoldPolicy	A hold that is associated with an eDiscovery case was deleted. Deleting a hold releases all of the content locations from the hold. Deleting the hold also results in deleting the case hold rules associated with the hold (see **Remove-CaseHoldRule** below).
Changed hold in eDiscovery case	Set-CaseHoldPolicy	A hold that is associated with an eDiscovery was changed. Possible changes include adding or removing content locations or turning off (disabling) the hold.
Created search query for eDiscovery case hold	New-CaseHoldRule	A query-based hold associated with an eDiscovery case was created.
Deleted search query for eDiscovery case hold	Remove-CaseHoldRule	A query-based hold associated with an eDiscovery case was deleted. Removing the query from the hold is often the result of deleting a hold. When a hold or a hold query is deleted, the content locations that were on hold are released.
Changed search query for eDiscovery case hold	Set-CaseHoldRule	A query-based hold associated with an eDiscovery case was changed. Possible changes include editing the query or date range for a query-based hold.
Created eDiscovery case	New-ComplianceCase	An eDiscovery case was created. When a case is created, you only have to give it a name. Other case-related tasks such as adding members, creating holds, and creating content searches associated with the case result in additional events being logged.
Deleted eDiscovery case	Remove-ComplianceCase	An eDiscovery case was deleted. Any hold associated with the case has to be removed before the case can be deleted.

FRIENDLY NAME	OPERATION (CMDLET)	DESCRIPTION
Changed eDiscovery case	Set-ComplianceCase	An eDiscovery case was changed. Changes include closing an open case or reopening a closed case.
Added member to eDiscovery case	Add-ComplianceCaseMember	A user was added as a member of an eDiscovery case. As a member of a case, a user can perform various case-related tasks depending on whether they have been assigned the necessary permissions.
Removed member from eDiscovery case	Remove-ComplianceCaseMember	A user was removed as a member of an eDiscovery case.
Changed eDiscovery case membership	Update-ComplianceCaseMember	The membership list of an eDiscovery case was changed. This activity is logged when all members are replaced with a group of new users. If a single member is added or removed, the **Add-ComplianceCaseMember** or **Remove-ComplianceCaseMember** operation is logged.
Created content search	New-ComplianceSearch	A new content search was created.
Deleted content search	Remove-ComplianceSearch	An existing content search was deleted.
Changed content search	Set-ComplianceSearch	An existing content search was changed. Changes can include adding or removing content locations that are searched and editing the search query.
Started content search	Start-ComplianceSearch	A content search was started. When you create or change a content search by using the Security & Compliance Center GUI, the search is automatically started. If you create or change a search by using the **New-ComplianceSearch** or **Set-ComplianceSearch** cmdlet, you have to run the **Start-ComplianceSearch** cmdlet to start the search.
Stopped content search	Stop-ComplianceSearch	A content search that was running was stopped.
Created content search action	New-ComplianceSearchAction	A content search action was created. Content search actions include previewing search results, exporting search results, preparing search results for analysis in Advanced eDiscovery, and permanently deleting items that match the search criteria of a content search.
Deleted content search action	Remove-ComplianceSearchAction	A content search action was deleted.

FRIENDLY NAME	OPERATION (CMDLET)	DESCRIPTION
Created search permissions filter	New-ComplianceSecurityFilter	A search permissions filter was created.
Deleted search permissions filter	Remove-ComplianceSecurityFilter	A search permissions filter was deleted.
Changed search permissions filter	Set-ComplianceSecurityFilter	A search permissions filter was changed.
Created eDiscovery administrator	Add-eDiscoveryCaseAdmin	A user was added as an eDiscovery Administrator in your organization.
Deleted eDiscovery administrator	Remove-eDiscoveryCaseAdmin	An eDiscovery Administrator was deleted from your organization.
Changed eDiscovery administrator membership	Update-eDiscoveryCaseAdmin	The list of eDiscovery Administrators in your organization was changed. This activity is logged when the list of eDiscovery Administrators is replaced with a group of new users. If a single user is added or removed, the **Add-eDiscoveryCaseAdmin** or **Remove-eDiscoveryCaseAdmin** operation is logged.

Detailed properties for eDiscovery activities

The following table describes the properties that are included when you click **More information** on the **Details** page for an eDiscovery activity listed in the search results. These properties are also included in the CSV file when you export the audit log search results. An audit log record for an eDiscovery activity won't include every detailed property listed below.

> **TIP**
> When you export the search results, the CSV file contains a column named **Detail**, which contains the detailed properties described in the following table in a multi-value property. You can use the Power Query feature in Excel to split this column into multiple columns so that each property will have its own column. This will let you sort and filter on one or more of these properties. For more information, see the "Export the search results to a file" section in Search the audit log.

PROPERTY	DESCRIPTION
Case	The identity (GUID) of the eDiscovery case that was created, changed, or deleted.
ClientApplication	eDiscovery cmdlet activities have a value of **EMC** for this property. This indicates the activity was performed by using the Security & Compliance Center GUI or running the cmdlet in PowerShell.
ClientIP	The IP address of the device that was used when the activity was logged. The IP address is displayed in either an IPv4 or IPv6 address format.
ClientRequestId	For eDiscovery activities, this property is typically blank.

PROPERTY	DESCRIPTION
CmdletVersion	The build number for the version of the Security & Compliance Center running in your organization.
CreationTime	The date and time in Coordinated Universal Time (UTC) when the eDiscovery activity was completed.
EffectiveOrganization	The name of the Microsoft 365 organization.
ExchangeLocations	The Exchange Online mailboxes that are included in a content search or placed on hold in an eDiscovery case.
Exclusions	Mailbox or site locations that are excluded from a content search or a hold in an eDiscovery case.
ExtendedProperties	Additional properties from a content search, a content search action, or hold in an eDiscovery case, such as the object GUID and the corresponding cmdlet and cmdlet parameters that were used when the activity was performed.
Id	The ID of the report entry. The ID uniquely identifies the audit log entry.
NonPIIParameters	A list of the parameters (without any values) that were used with the cmdlet identified in the Operation property. The parameters listed in this property are the same as those listed in the Parameters property.
ObjectId	The GUID or name of the object (for example, a Content Search or an eDiscovery case) that was created, changed, or deleted by the activity listed in the Operation property. This object is also identified in the Item column in the audit log search results.
ObjectType	The type of eDiscovery object that the user created, deleted, or modified; for example, a content search action (preview, export, or purge), an eDiscovery case, or a content search.
Operation	The name of the operation that corresponds to the eDiscovery activity that was performed.
OrganizationId	The GUID for your Microsoft 365 organization.
Parameters	The name and value for the parameters that were used with the corresponding cmdlet.
PublicFolderLocations	The public folder locations in Exchange Online that are included in a content search or placed on hold in an eDiscovery case.
Query	The search query associated with the activity, such as a content search or a query-based hold.

PROPERTY	DESCRIPTION
RecordType	The type of operation indicated by the record. The value of **18** indicates an event related to an activity listed in the eDiscovery cmdlet activities section. A value of **24** indicates an event related to an activity listed in the How to search for and view eDiscovery activities section.
ResultStatus	Indicates whether the action (specified in the Operation property) was successful or not.
SecurityComplianceCenterEventType	Indicates that the activity was a Security & Compliance Center event. All eDiscovery activities will have a value of **0** for this property.
SharepointLocations	The SharePoint Online sites that are included in a content search or placed on hold in an eDiscovery case.
StartTime	The date and time in Coordinated Universal Time (UTC) when the eDiscovery activity was started.
UserId	The user who performed the activity (specified in the Operation property) that resulted in the record being logged. Records for eDiscovery activity performed by system accounts (such as NT AUTHORITY\SYSTEM) are also included in the audit log.
UserKey	An alternative ID for the user identified in the UserId property. For eDiscovery activities, the value for this property is typically the same as the UserId property.
UserServicePlan	The subscription used by your organization. For eDiscovery activities, this property is typically blank.
UserType	The type of user that performed the operation. The following values indicate the user type. 0 A regular user. 2 An administrator in your organization. 3 A Microsoft datacenter administrator or datacenter system account. 4 A system account. 5 An application. 6 A service principal.
Version	Indicates the version number of the activity (identified by the Operation property) that's logged.
Workload	Theservice where the activity occurred. For eDiscovery activities, the value is **SecurityComplianceCenter**.

Overview of Microsoft 365 Advanced eDiscovery

2/18/2021 • 8 minutes to read • Edit Online

The Advanced eDiscovery solution in Microsoft 365 builds on the existing Microsoft eDiscovery and analytics capabilities. Advanced eDiscovery provides an end-to-end workflow to preserve, collect, analyze, review, analyze, and export content that's responsive to your organization's internal and external investigations. It also lets legal teams manage the entire legal hold notification workflow to communicate with custodians involved in a case.

Advanced eDiscovery can help your organization respond to legal matters or internal investigations by discovering data where it lives. You can seamlessly manage eDiscovery workflows by identifying persons of interest and their data sources, seamlessly apply holds to preserve data, and then manage the legal hold communication process. By collecting data from the source, you can search the live Microsoft 365 platform to quickly find what you need. Intelligent, machine learning capabilities such as deep indexing, email threading, and near duplicate detection also help you reduce large volumes of data to a relevant data set.

The following sections describe how these Advanced eDiscovery capabilities can help your organization.

Discover and collect data in-place

Traditionally, organizations that rely on multiple third-party eDiscovery solutions require copying large volumes of data out of Microsoft 365 to process and having to host duplicate data. This necessity increases the time to find relevant data and the risk, cost, and complexity of managing multiple solutions.

Advanced eDiscovery in Microsoft 365 lets you discover data at the source and staying within your Microsoft 365 security and compliance boundary. By collecting data in-place from the live system, Advanced eDiscovery reduces the friction of going back to the source and reduces unnecessary work of having to find missing content, which often happens when journaling lags in traditional eDiscovery solutions.

Native search and collection capabilities for data in Teams, Yammer, SharePoint Online, OneDrive for Business, and Exchange Online further enhances data discovery. For example, Advanced eDiscovery:

- Reconstructs Teams conversations (instead of returning individual messages from conversations).
- Collects cloud-based content shared with users by use of links or modern attachments in email message and Teams chats.
- Has built-in support for hundreds of non-Microsoft 365 file types.
- Collects data from third-party sources (such as Bloomberg, Facebook, Slack, and Zoom Meetings) that's imported and archived in Microsoft 365 by data connectors.

Manage eDiscovery workflow in one platform

Advanced eDiscovery can help you reduce the number of eDiscovery solutions you need to rely on. It provides a streamlined, end-to-end workflow, all which occurs within Microsoft 365. Advanced eDiscovery helps reduce the friction of identifying and collecting potential sources of relevant information by automatically mapping unique and shared data sources to the person of interest (known as a *custodian*), and by providing reporting and analytics on potentially relevant data prior to collecting it for analysis and review.

Additionally, Microsoft Graph APIs can help you automate the eDiscovery workflow and extend Advanced eDiscovery for custom solutions.

Cull data intelligently

Intelligent, machine learning capabilities in Advanced eDiscovery help you reduce the amount of data to review. These intelligent capabilities help you reduce and cull large volumes of data to a relevant set. For example, a built-in review set query helps filter only for unique content by identifying near duplicates. This capability can substantially reduce the amount of data to review.

Additional machine learning capabilities can further refine and identify relevant data using smart tags and technology assisted review tools like the Relevance modules.

Advanced eDiscovery architecture

Here's an Advanced eDiscovery architecture diagram that shows the end-to-end workflow in a single-geo environment and in a multi-geo environment, and the end-to-end data flow that's aligned with the Electronic Discovery Reference Model (EDRM).

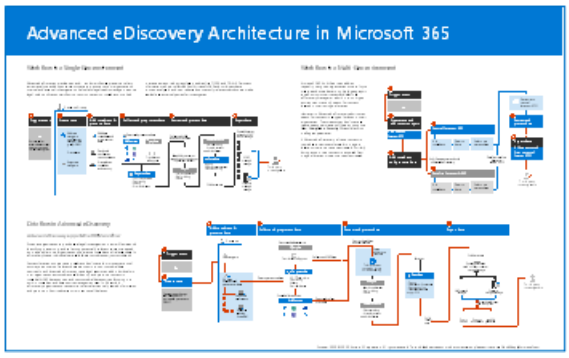

View as an image

Download as a PDF file

Download as a Visio file

For more information about the end-to-end workflow in Advanced eDiscovery, see this Microsoft Mechanics video.

Advanced eDiscovery workflow

The following sections describe each step in the built-in workflow in the Advanced eDiscovery tool in the Microsoft 365 compliance center. The following screenshot shows the **Overview** tab of a case named *2020.11.03 - Contoso v. Fabrikam*.

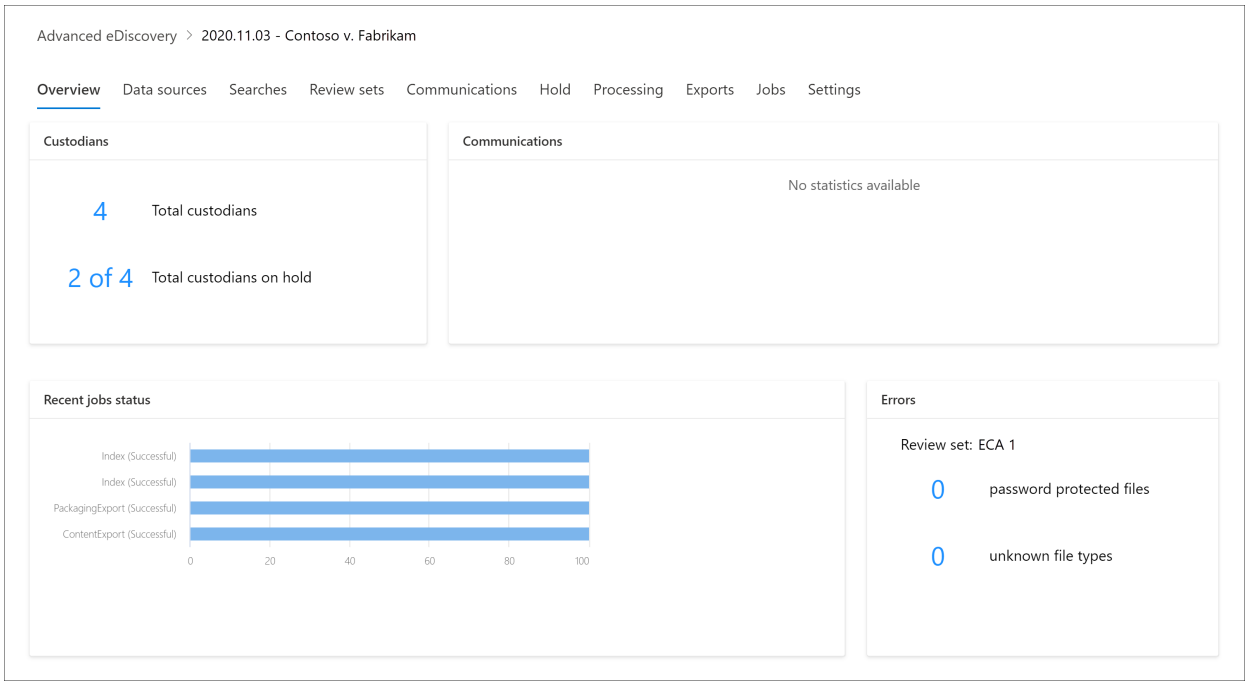

For more detailed information, see Manage the Advanced eDiscovery workflow.

Managing custodians and non-custodial data sources

Use the **Data sources** tab to add and manage the people that you've identified as persons of interest in the case and other data sources that may not be associated with a custodian. When you add custodians or non-custodial data sources, you can quickly perform actions like placing a legal hold on custodian and non-custodial data sources, communicating with custodians, and searching custodian and non-custodial data sources to collect content that's relevant to the case. As the case progresses, it's easy to add new custodians or non-custodial date sources or release them from the case. For more information, see Work with custodians.

Managing legal hold notifications

Use the **Communications** tab to manage the process of communicating with the custodians in the case. A legal hold notice instructs custodians to preserve any content that's relevant to the case. Legal teams must be able to track the notices that have been received, read, and acknowledged by custodians. The communications workflow in Advanced eDiscovery allows you to create and send initial notifications, reminders, release notices, and escalations if custodians fail to acknowledge a hold notification. For more information, see Work with communications.

Managing content preservation

When you add a custodian to a case, you can place a hold on custodial data. Use the **Hold** tab to manage the hold created when you add custodians, and to manage other legal holds associated with the case; for example, you can identify and place a hold on non-custodial data sources. You can also edit any hold in the case and make it a query-based hold to preserve only the content that matches the query. For example, you could add a date range to the hold so that only content created within a specific date ranged in preserved. You can also get statistics on content that's on hold, remove the hold after it's no longer relevant to the case, or delete it. For more information, see Manage holds.

Indexing custodian data

When you add a custodian and the corresponding custodial data sources to a case, any partially indexed item from a custodian data source is reindexed by a process called *Advanced indexing*. This allows custodial content such as images, unsupported file types, and other potentially unindexed content to be fully searchable when you run searches to collect data for the case. Use the **Processing** tab to monitor the status of Advanced indexing and fix processing errors by using a process called *error remediation*. For more information, see Fix processing errors.

Collecting case data

Use the **Searches** tab to create searches to search the in-place custodial and non-custodial data sources for content relevant to the case. You can create and run query-based searches (using keywords and conditions) to identify a set of email messages and documents that are relevant to the case and that you want to further review and analyze in subsequent steps in the eDiscovery workflow. You can create one or more searches associated with the case. You can also use the search tool to preview sample documents and view search statistics to help you refine and improve the search results. After you're satisfied the search results contain the all data relevant to the case, you add the search results to a review set for further review, analysis, and culling. For more information, see Collect data for a case.

Reviewing and analyzing case data

Use the **Review sets** tab to review and analyze the content you've collected from the live system and added to a review set. A *review set* is a static collection of that data (in other words, an offline copy of data) of custodial data (and, if applicable, non-custodial data) that you collected in the previous phase of the eDiscovery workflow. When you add search results to a review set, a process is triggered to extract files from containers, extract metadata, and extract text. When this process is complete, the system builds a new index of all the data collected from custodians and adds it to the review set. After the data is added to the review set, you can run more queries to narrow the case data, view data as text or in the native file format, and annotate, redact, and tag documents in the review set. You can also perform advanced analytics, such as identifying document duplication, email threading, and themes. After you've culled the data to only what is relevant to the case, you can either download documents directly or export them along with file metadata, annotations, and any tags. For more information, see:

- View documents in a review set
- Query the data in a review set
- Tag documents in a review set
- Analyze data in a review set

Exporting data for review and presentation

After you export the data from a review set, use the **Exports** tab to manage an export job and download data from a review set. When you export a review set, the data is uploaded to a Microsoft-provided Azure Storage location (or an Azure Storage location managed by your organization). After it's uploaded to Azure, it's then and available to download to a local computer. You can obtain the storage assess key necessary to download the exported data on the **Exports** tab. For more information, see Export case data.

Managing jobs

Use the **Jobs** tab to monitor long-running processes for case-related tasks that you've initiated. Examples of jobs include ones related to reindexing, searching, and exporting case data. For example, if you create a search on the **Searches** tab that includes many data sources, the status of this search process will be displayed on the **Jobs** tab. For more information, see Manage jobs.

Configuring case settings

Use the **Settings** tab to configure case-wide settings. This includes adding members to a case, closing or deleting a case, and configuring search and analytics settings. For more information, see:

- Add members to a case
- Close or delete a case
- Configure search and analytics settings

Set up Microsoft 365 Advanced eDiscovery

2/18/2021 • 4 minutes to read • Edit Online

Advanced eDiscovery in Microsoft 365 provides an end-to-end workflow to preserve, collect, review, analyze, and export data that's responsive to your organization's internal and external investigations. Nothing is needed to deploy Advanced eDiscovery, but there are some prerequisite tasks that an IT admin and eDiscovery manager have to complete before your organization can start to create and use Advanced eDiscovery cases to manage your investigations.

This article discusses the steps necessary to set up Advanced eDiscovery. This includes ensuring the proper licensing required to access Advanced eDiscovery and add custodians to cases, and assigning permissions to your legal and investigation team so they can access and manage cases.

Step 1: Verify and assign appropriate licenses

Licensing for Advanced eDiscovery requires the appropriate organization subscription and per-user licensing.

- **Organization subscription:** To access Advanced eDiscovery in the Microsoft 365 compliance center or the Security & Compliance Center, your organization must have one of the following:
  - Microsoft 365 E5 or Office 365 E5 subscription
  - Microsoft 365 E3 subscription with E5 Compliance add-on
  - Microsoft 365 E3 subscription with E5 eDiscovery and Audit add-on

  If you don't have an existing Microsoft 365 E5 plan and want to try Advanced eDiscovery, you can add Microsoft 365 to your existing subscription or sign up for a trial of Microsoft 365 E5.

- **Per-user licensing:** To add a user as a custodian in an Advance eDiscovery case, that user must be assigned one of the following licenses, depending on your organization subscription:
  - Microsoft 365: Users must be assigned a Microsoft 365 E5 license, an E5 Compliance add-on license, or an E5 eDiscovery and Audit add-on license.
  - Office 365: Users must be assigned an Office 365 E5 license.

  For information about how to assign licenses, see Assign licenses to users.

> **NOTE**
> Users only need an E5 license (or the appropriate add-on license) to be added as custodians to an Advanced eDiscovery case. IT admins, eDiscovery managers, lawyers, paralegals, or investigators who use Advanced eDiscovery to manage cases and review case data don't need an E5 or add-on license.

Step 2: Assign eDiscovery permissions

To access Advanced eDiscovery or added as a member of an Advanced eDiscovery case, a user must be assigned the appropriate permissions. Specifically, a user must be added as a member of the eDiscovery Manager role group in the Security & Compliance Center. Members of this role group can create and manage Advanced eDiscovery cases. They can add and remove members, place custodians and content locations on hold, manage legal hold notifications, create and edit searches associated in a case, add search results to a review set, analyze data in a review set, and export and download from an Advanced eDiscovery case.

Complete the following steps to add users to the eDiscovery Manager role group:

1. Go to https://protection.office.com/permissions and sign in using the credentials for an admin account in your Microsoft 365 organization.
2. On the **Permissions** page, select the **eDiscovery Manager** role group.
3. On the eDiscovery Manager flyout page, click **Edit** next to the **eDiscovery Manager** section.
4. On the **Choose eDiscovery Manager** page in the edit role group wizard, click **Choose eDiscovery Manager**.
5. Click **Add** then select the checkbox for all users you want to add to the role group.
6. Click **Add** to add the selected users, and then click **Done**.
7. Click **Save** to add the users to the role group, and then click **Close** to complete the step.

More information about the eDiscovery Manager role group

There are two subgroups in the eDiscovery Manager role group. The difference between these subgroups is based on scope.

- **eDiscovery Manager:** Can view and manage the Advanced eDiscovery cases they create or are a member of. If another eDiscovery Manager creates a case but doesn't add a second eDiscovery Manager as a member of that case, the second eDiscovery Manager won't be able to view or open the case on the Advanced eDiscovery page in the compliance center. In general, most people in your organization can be added to the eDiscovery Manager subgroup.
- **eDiscovery Administrator:** Can perform all case management tasks that an eDiscovery Manager can do. Additionally, an eDiscovery Administrator can:
  - View all cases that are listed on the Advanced eDiscovery page.
  - Manage any case in the organization after they add themselves as a member of the case.
  - Access and export case data for any case in the organization.

  Because of the broad scope of access, an organization should have only a few admins who are members of the eDiscovery Administrators subgroup.

For more information about eDiscovery permissions and a description of each role that's assigned to the eDiscovery Manager role group, see Assign eDiscovery permissions.

Step 3: Configure global settings for Advanced eDiscovery

The last step to complete before people in your organization start to create and use cases is to configure global settings that apply to all cases in your organization. At this time, the only global setting is *attorney-client privilege detection* (more global settings will be available in the future). This setting enables the attorney-client privilege model to run when you analyze data in a review set. The model uses machine learning to determine the likelihood that a document contains content that is legal in nature. It also compares the participants of documents with an attorney list (that you submit when setting up the model) to determine if a document has at least one participant who is an attorney.

For more information about setting up and using the attorney-client privilege detection model, see Set up attorney-client privilege detection in Advanced eDiscovery.

> **NOTE**
>
> This is an optional step that you can perform anytime. Not implementing the attorney-client privilege detection model doesn't prevent you from creating and using Advanced eDiscovery cases.

Next steps

After you set up Advanced eDiscovery, you're ready to create a case.

Create and manage an Advanced eDiscovery case

2/18/2021 • 6 minutes to read • Edit Online

After setting up Advanced eDiscovery and assigning permissions to eDiscovery managers in your organization that will manage cases, the next step is to create and manage a case.

This article also provides a high-level overview of using cases to manage the Advanced eDiscovery workflow for a legal investigation.

Create a case

Complete the following steps to create a case and add members. The user who creates the case is automatically added as a member.

1. Go to https://compliance.microsoft.com and sign in using the credentials for user account that has been assigned eDiscovery permissions. Members of the Organization Management role group can also create Advanced eDiscovery cases.
2. In the left navigation pane of the Microsoft 365 compliance center, click **Show all**, and then click **eDiscovery > Advanced**.
3. On the **Advanced eDiscovery** page, click the **Cases** tab, and then click **Create a case**.
4. On the **New eDiscovery case** flyout page, give the case a name (required), and then type an optional case number and description. The case name must be unique in your organization.
5. Click **Save** to create the case.

   The new case is created and the **Settings** tab in the new case is displayed.
6. In the **Access & permissions** tile on the **Settings** tab, click **Select**, and then click **Update**.
7. Click **Update**.
8. On the **Manage this case** flyout page, under **Manage members**, click **Add** to add members to the case.
9. In the list of people, select the check box next to the names of the people that you want to add to the case. As previously explained, be sure that the people you add to the case have been assigned the appropriate eDiscovery permissions.
10. After you've selected the people to add as members of the case, click **Add**.
11. In the **Manage this case** flyout page, click **Save** to save the new list of case members.
12. Click the **Home** tab to go to the case home page.

Manage the workflow

To get you started using Advanced eDiscovery, here's a basic workflow that aligns with common eDiscovery practices. In each of these steps, we'll also highlight some extended Advanced eDiscovery functionality that you can explore.

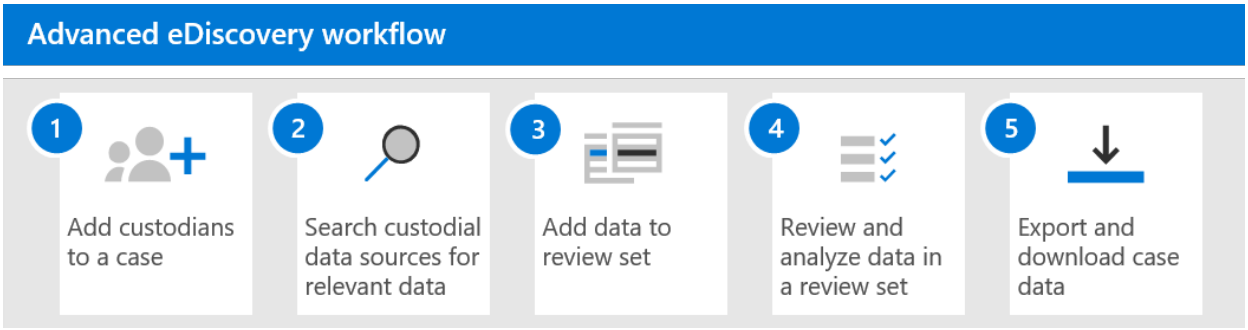

1. **Add custodians and non-custodial data sources to the case**. The first step after creating a case is to add custodians. A *custodian* is a person having administrative control of a document or electronic file that may be relevant to the case. Additionally, you can add data sources that aren't associated with a specific user but may be relevant to the case.

   Here are some things that happen (or that you can do) when you add custodians to a case:

   - Data in the custodian's Exchange mailbox, OneDrive account, and any Microsoft Teams or Yammer groups that the custodian is a member of can be "marked" as custodial data in the case.
   - Custodian data is reindexed (by a process called *Advanced indexing*). This helps optimize searching for it in the next step.
   - You can place a hold on custodian data. This preserves data that may be relevant to the case during the investigation.
   - You can associate other data sources with a custodian (for example, you can associate a SharePoint site or Microsoft 365 Group with a custodian) so this data can be reindexed, placed on hold, and searched, just like the data in the custodian's mailbox or OneDrive account.
   - You can use the communications workflow in Advanced eDiscovery to send a legal hold notification to custodians.

2. **Search data sources for data relevant to the case**. After you add custodians and non-custodial data sources to a case, use the built-in search tool to search these data sources for data that may be relevant to the case. You use keywords, properties, and conditions to build search queries that return search results with the data that's most likely relevant to the case. You can also:

   - View search statistics that may help you refine a search query to narrow the results.
   - Preview the search results to quickly verify whether the relevant data is being found.
   - Revise a query and rerun the search.

3. **Add data to a review set**. Once you've configured and verified that a search returns the desired data, the next step is to add the search results to a review set. When you add data to a review set, items are copied from their original location to a secure Azure Storage location. The data is reindexed again to optimize it for thorough and fast searches when reviewing and analyzing items in the review set. Additionally, you can also add non-Office 365 data into a review set.

   There's also a special kind of review set that you can add data to, called a *conversation review set*. These types of reviews sets provide conversation reconstruction capabilities to reconstruct, review, and export threaded conversations like those in Microsoft Teams. For more information, see Review conversations in Advanced eDiscovery.

4. **Review and analyze data in a review set**. Now that data is in a review set, you can use a wide-variety of tools and capabilities to view and analyze the case data with the goal of reducing the data set to what is most relevant to the case you're investigating. Here's a list of some tools and capabilities that you can use during this process.

- View documents. This includes viewing the metadata for each document in a review set, and viewing the document in its native version or text version.
- Create queries and filters. You create search queries using various search criteria (including the ability to search all file metadata properties) to further refine and cull the case data to what is most relevant to the case. You can also use review set filters to quickly apply other conditions to the results of a search query to further refine those results.
- Create and use tags. You can apply tags to documents in a review set to identify which are responsive (or non-responsive to the case) and then use those tags when creating search queries to include or exclude the tagged documents. You can also tagging to determine which documents to export.
- Annotate and redact documents. You can use the annotation tool in a review to annotate documents and redact content in documents as work product. We generate a PDF version of an annotated or redacted document during review to reduce the risk of exporting the unredacted native version of the document.
- Analyze case data. The analytics functionality in Advanced eDiscovery is powerful. After you run analytics on the data in review set, we perform analysis such as near duplicate detection, email threading, and themes that can help reduce the volume of documents that you have to review. We also generate an Analytics reports that summarize the result of running analytics. As previously explained, running analytics also runs the attorney-client privilege detection model.

5. **Export and download case data**. A final step after collecting, reviewing, and analyzing case data is to export it out of Advanced eDiscovery for external review or for review by people outside of the investigation team. Exporting data is a two-step process. The first step is to export data out of the review set and copy it to a different Azure Storage location (one provided by Microsoft or one managed by your organization). Then you use Azure Storage Explorer to download the data to a local computer. In addition to the exported data files, the contains of the export package also contains an export report, a summary report, and an error report.

Work with custodians and non-custodial data sources in Advanced eDiscovery

11/2/2020 • 2 minutes to read • Edit Online

When an organization responds to a legal investigation, the workflow around identifying, preserving, and collecting potentially relevant content is based on the people in the organization who are the custodians of relevant data. In eDiscovery, these individuals are called *data custodians* (or just *custodians*) and are defined as "persons having administrative control of a document or electronic file". For example, the custodian of an email message could be the owner of the mailbox that contains the relevant message.

Additionally, there may be content located in mailboxes and sites that aren't associated with a custodian but that's relevant to the case. Content locations where case custodians don't have administrative control but may be owners of relevant data, are known as *non-custodial data sources*.

In an Advanced eDiscovery case, legal teams can add individuals in their organization as custodians, and identify and preserve custodial data sources such as Exchange mailboxes, OneDrive accounts, and SharePoint and Teams sites. They can also identify and preserve non-custodial data sources. By using the built-in custodian and data source management tool in Advanced eDiscovery, organizations can secure electronically stored information from inadvertent (or intentional) deletion. This lets you eliminate the time-consuming and error-prone process of manually having to perform the legal hold processes.

For more information about working with custodians, see the following articles:

- Add custodians to a case
- Bulk-add custodians to a case
- Manage custodians in a case
- View custodian activity
- Add non-custodial data sources to a case

Add custodians to an Advanced eDiscovery case

2/18/2021 • 5 minutes to read • Edit Online

Use the built-in custodian management tool in Advanced eDiscovery to coordinate your workflows around managing custodians and identifying relevant, custodial data sources associated with a case. When you add a custodian, the system can automatically identify and place a hold on their Exchange mailbox and OneDrive for Business account. During the discovery process of your investigation, you might also identify other data sources (such as mailboxes, sites, or Teams) that a custodian accessed or contributed to. In this situation, you can use the custodian management tool to associate those data sources will a specific custodian. After you add custodians to a case and associate other data source with them, you can quickly preserve data and search the custodial data.

You can add and manage custodians in Advanced eDiscovery cases in four steps:

1. Identify the custodians.
2. Choose custodian data locations.
3. Configure hold settings.
4. Review the custodians and complete the process.

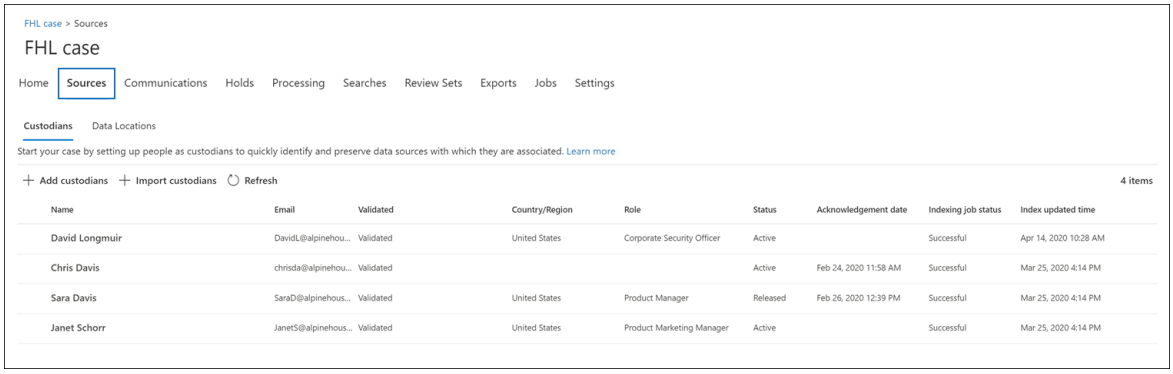

Make sure you have the necessary permissions

To add custodians to a case, you must be a member of the eDiscovery Manager role group. This provides you with the necessary permissions to add custodians to a case and place a hold on the custodial data sources. For more information, see Assign eDiscovery permissions.

Step 1: Identify custodians

1. Go to https://compliance.microsoft.com and sign in with a user account that has been assigned the appropriate eDiscovery permissions.
2. In the left navigation pane of the Microsoft 365 compliance center, click **Show all**, and then click **eDiscovery > Advanced**.
3. On the **Advanced eDiscovery** page, click the **Cases** tab, and then select the case that you want to add custodians to.
4. Click the **Data sources** tab and then click **Add data source** > **Add new custodians**.
5. Add one or more users in your organization as custodians to the case by typing the first part of a person's name or alias. After you find the correct person, select their name to add them to the list.

Step 2: Choose custodian data locations

After you select custodians, the system automatically attempts to identify and verify these users and their data sources. After adding custodians to the list, the tool automatically includes the primary mailbox and OneDrive account for each custodian. You can choose not to include these data sources when adding custodians to the case.

In addition to a custodian's mailbox and OneDrive account, you can also associate other data locations to a custodian, such as SharePoint site or a Microsoft Team the custodian is a member of. This allows you to preserve, collect, analyze, and review content in other data sources associated with the custodians of the case.

To deselect the primary mailbox and OneDrive account for a custodian:

1. Expand the custodian to view the primary data locations that have been automatically associated to each custodian.
2. Select **Clear** next to **Mailbox** or **OneDrive** to remove a custodian's mailbox or OneDrive account from being associated as a data location for this custodian.

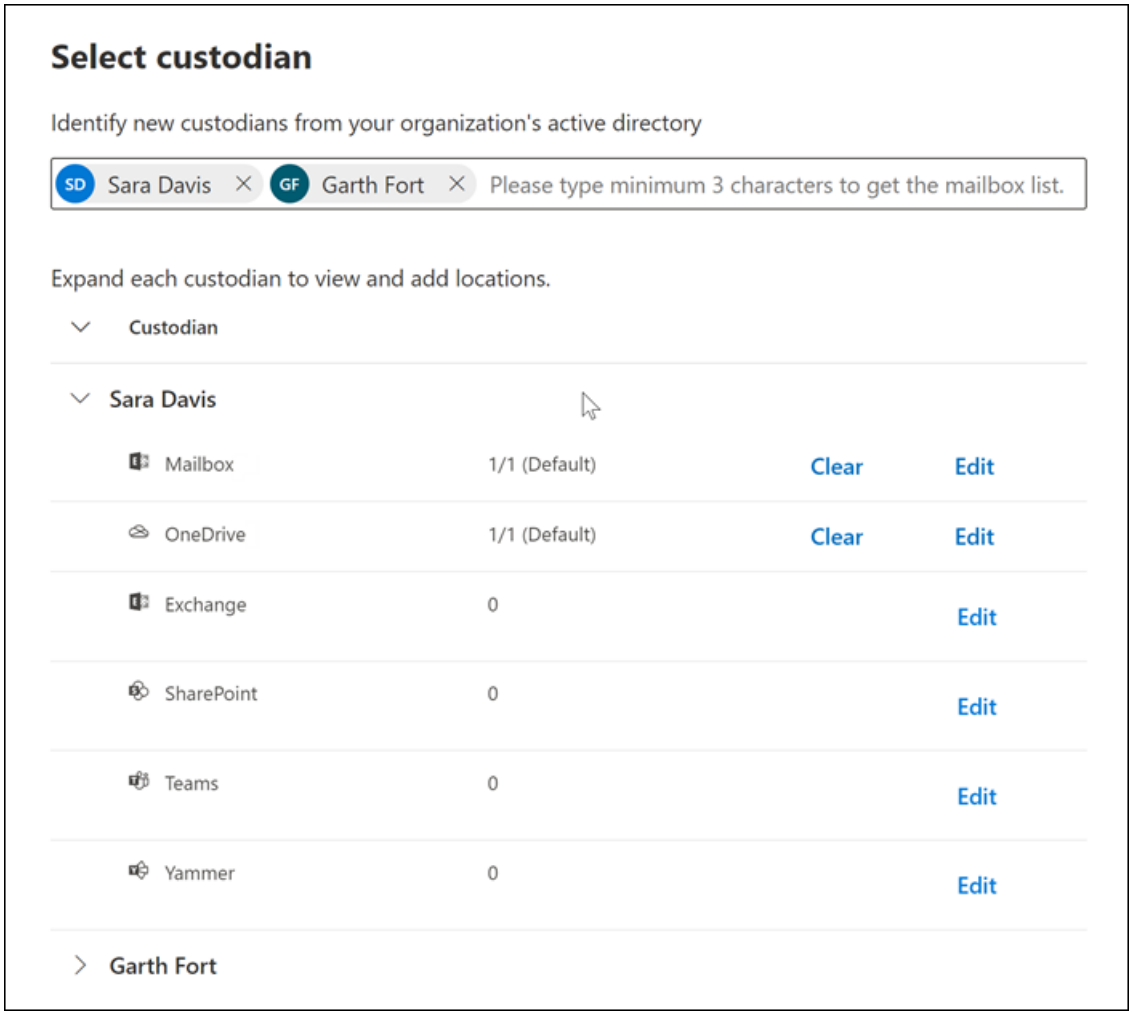

To associate other mailboxes, sites, Teams, or Yammer groups to a specific custodian:

1. Expand a custodian to display the following services to associate data locations with the custodian. Click **Edit** next to a service to add a data location.
   - **Exchange**: Use to associate other mailboxes to the custodian. Type into the search box the name or alias (a minimum of three characters) of user mailboxes or distribution groups. Select the mailboxes to assign to the custodian and then click **Add**.
   - **SharePoint**: Use to associate SharePoint sites to the custodian. Select a site in the list or search for a site by typing a URL in the search box. Select the sites to assign to the custodian and then click

Add.

- **Teams**: Use to assign the Microsoft Teams the custodian is currently a member of. Select the teams to assign to the custodian and then click **Add**. After you add a team, the system automatically identifies and locates the SharePoint site and group mailbox associated to that team and assigns them to the custodian.
- **Yammer**: Use to assign the Yammer groups the custodian is currently a member of. Select the groups to assign to the custodian and then click **Add**. After you add a team, the system automatically identifies and locates the SharePoint site and group mailbox associated to that group and assigns them to the custodian.

> **NOTE**
> You can use the **Exchange** and **SharePoint** location pickers to associate other teams or Yammer groups (that a custodian is not a member of) to a custodian. To do this, you have to add both the mailbox and site associated with each team or Yammer group.

2. You can view the total number of mailboxes, sites, Teams, and Yammer groups assigned to each custodian by expanding each custodian in the table. When you've finalized the assigned data locations for each custodian, these associations will be maintained and used during the collection, processing, and review stages in the Advanced eDiscovery workflow.
3. After adding custodians and configuring their data locations, click **Next** to go to the **Hold settings** page.

Step 3: Configure hold settings

After you've finalized the custodians and their data locations, you can place some or all of the custodians on hold. When you place a custodian on hold, all content in all content locations that are associated with the custodian is preserved until you remove the hold or release the custodian from the hold. In some cases, you may want to add custodians to a case without placing them on hold.

To place the custodians and data sources on hold:

1. On the **Hold settings** page, you can apply a hold to individual custodians by selecting the checkbox under the **Hold** column.

   Alternatively, you can place all custodians on hold by selecting the **Hold** checkbox at the top of the column.

2. Verify the custodian hold selections and then click **Next**.

> **NOTE**
> If you don't place a hold on a custodian, the custodian and their associated data sources will be added to the case but the content in those data sources won't preserved by the hold that associated with the case.

Step 4: Review the custodians and complete the process

Before you actually add the custodians to the case, you can review the list of custodians, the data locations assigned to them, and the hold settings.

1. Verify and review all the data sources count and the hold setting associated with each custodian in the table. If necessary, go back to the **Identify custodian** or **Hold settings** pages to make any changes.
2. Click **Submit** to add custodians and their data locations to the case and apply all custodial hold settings.

The new custodians are added to the case and displayed on the **Data sources** tab.

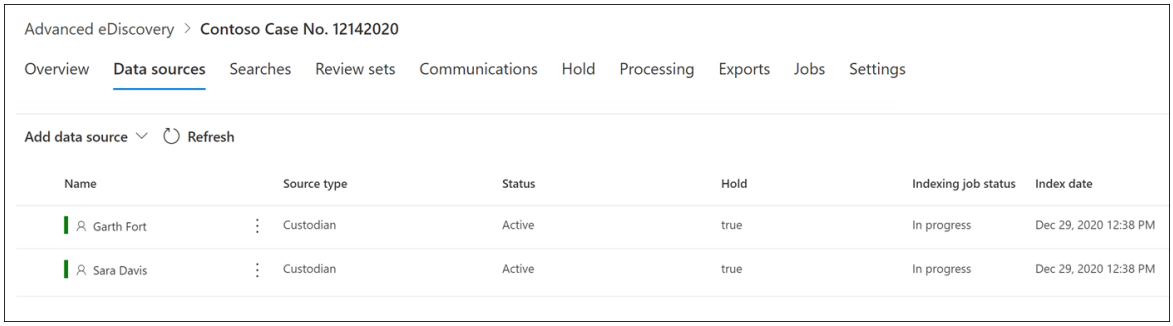

Import custodians to an Advanced eDiscovery case

2/18/2021 • 4 minutes to read • Edit Online

For Advanced eDiscovery cases that involve many custodians, you can import multiple custodians at once by using a CSV file that contains the information necessary to add them to a case.

Import custodians

1. Open the Advanced eDiscovery case and select the **Data sources** tab.
2. Click **Add data source** > **Import custodians**.
3. On the **Import custodians** flyout page, click **Download a blank template** to download a custodian template CSV file.

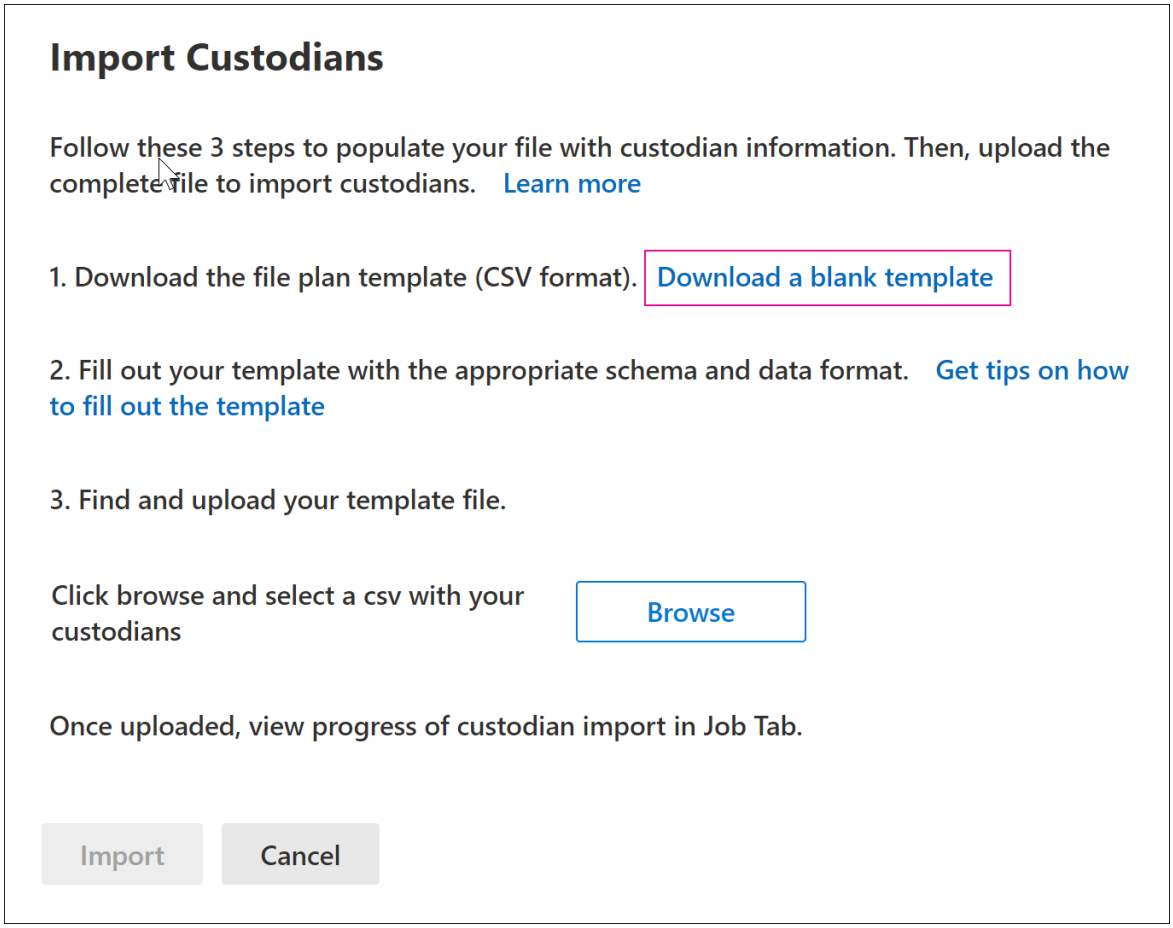

4. Add the custodial information to the CSV file and save it to your local computer. See the Custodian CSV file section for information about the required properties in the CSV file.
5. After you've prepared the CSV file with the custodian information, go back to the **Data sources** tab, and click **Add data source** > **Import custodians** again.
6. On the **Import custodians** flyout page, click **Browse** and then upload the CSV file that contains the custodian information.

   After the CSV file is uploaded, a job named **BulkAddCustodian** is created and displayed on the **Jobs** tab. The job validates the custodians and their associated data sources and then adds them to the **Data sources** page of the case.

Custodian CSV file

After you download the CSV custodian template, you can add custodians and their data source in each row. Be sure not to change the column names in the header row. Use the workload type and workload location columns to associate other data sources to a custodian.

COLUMN NAME	DESCRIPTION
Custodian contactEmail	The custodian's UPN email address. For example, sarad@contoso.onmicrosoft.com.
Exchange Enabled	TRUE/FALSE value to include or not include the custodian's mailbox.
OneDrive Enabled	TRUE/FALSE value to include or not included the custodian's OneDrive for Business account.
Is OnHold	TRUE/FALSE value to indicate whether to place the custodian data sources on hold.
Workload1 Type	String value indicating the type of data source to associate with the custodian. Possible values include: - ExchangeMailbox - SharePointSite - TeamsMailbox - TeamsSite - YammerMailbox - YammerSite
Workload1 Location	Depending on your workload type, this would be the location of the data source. For example, the email address for an Exchange mailbox or the URL for a SharePoint site.

Here's an example of a CSV file with custodian information:

CUSTODIAN CONTACTEMAIL	EXCHANGE ENABLED	ONEDRIVE ENABLED	IS ONHOLD	WORKLOAD1 TYPE	WORKLOAD1 LOCATION
robinc@onmicrosoft.contoso.com	TRUE	TRUE	TRUE	SharePointSite	https://contoso.sharepoint.com
pillarp@onmicrosoft.contoso.com	TRUE	TRUE	TRUE		

Custodian and data source validation

After you upload the custodian CSV file, Advanced eDiscovery does the following things:

1. Validates the custodians and their data sources.
2. Indexes all data sources for each custodian and places them on hold (if the **Is OnHold** property in the CSV file is set to TRUE).

Custodian validation

Currently, we only support importing custodians that are included in your organization's Azure Active Directory (Azure AD).

The custodian import tool finds and validates custodians using the UPN value in the **Custodian contactEmail** column in the CSV file. Custodians that are validated are automatically added to the case and listed on the **Data sources** tab of the case. If a custodian can't be validated, they are listed in the error log for the BulkAddCustodian job that is listed on the **Jobs** tab in the case. Unvalidated custodians are not added to the case or listed on the **Data sources** tab.

Data source validation

After custodians are validated and added to the case, each primary mailbox and OneDrive account that's associated with a custodian is added.

However, if any of the other data sources (such as SharePoint sites, Microsoft Teams, Microsoft 365 Groups, or Yammer groups) associated with a custodian can't be found, none of them are assigned to the custodian and the value **Not validated** is displayed in the **Status** column next to the custodian on the **Data sources** tab.

To add validated data sources for a custodian:

1. On the **Data sources** tab, select a custodian that contains data sources that aren't validated.
2. On the custodian flyout page, scroll to the **Custodial locations** section to view both validated and unvalidated data sources that are associated with custodian.
3. Click **Edit** at the top of the flyout page to remove invalid data sources or add new ones.
4. After you remove unvalidated data sources or add a new one, the value **Active** is displayed in **Status** column for the custodian on the **Data sources** tab. To add sources that previously appeared to be invalid, follow the remediation steps below to manually add them to a custodian.

Remediating invalid data sources

To manually add and associate a data source that was previously invalid:

1. On the **Data sources** tab, select a custodian to manually add and associate a data source that was previously invalid.
2. Click **Edit** at the top of the flyout page to associate mailboxes, sites, Teams, or Yammer groups to the custodian. Do this by clicking **Edit** next to the appropriate data location type.
3. Click **Next** to display the **Hold settings** page and configure the hold setting for the data sources you added.
4. Click **Next** to display the **Review custodians** page, and then click **Submit** to save your changes.

Manage custodians in an Advanced eDiscovery case

2/18/2021 • 5 minutes to read • Edit Online

The Custodians page on the **Sources** tab in an Advanced eDiscovery case contains a list of all custodians that have been added to the case. After you add custodians to a case, details about each custodian are automatically collected from Azure Active Directory and are viewable in Advanced eDiscovery.

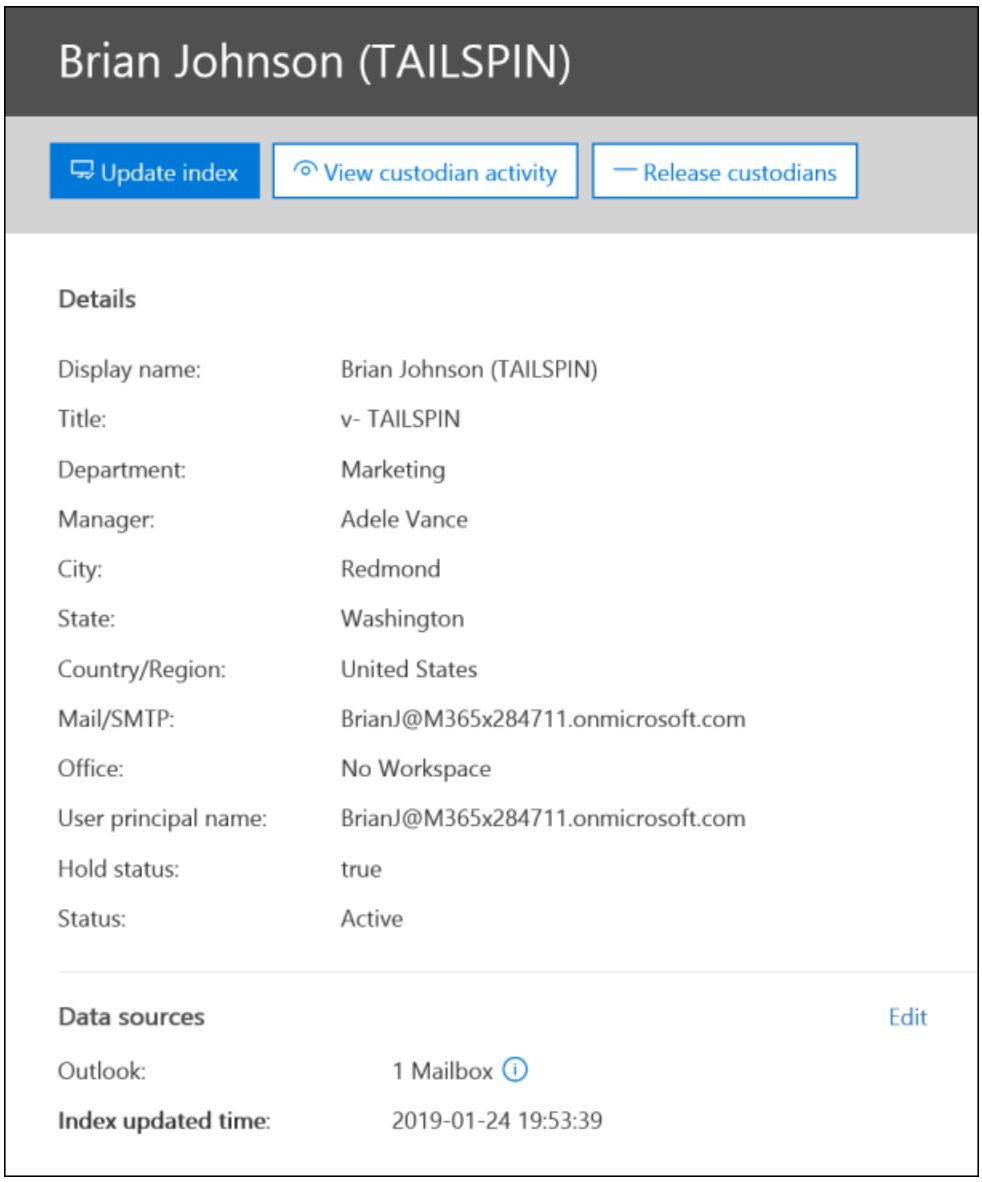

View custodian details

To view the details about a custodian, click the custodian from the list on the **Custodians** tab. A flyout page is displayed and contains the following information about the custodian:

- Contact information
  - **Display Name** - The name displayed in the address book for the custodian. This is usually the combination of the custodian's first name, middle initial, and last name.
  - **Mail/SMTP** - The primary SMTP address for the custodian, for example,

brianj@contoso.onmicrosoft.com. The custodian's user principal name (UPN) is also listed.

  - **Title** - The custodian's job title.
  - **Department** - The name for the department in which the custodian works.
  - **Manager** - The custodian's manager. The designated manager will receive any escalation communications for this custodian.

- Location information
  - **City** - The city in which the custodian is located.
  - **State** - The state or province in the custodian's address.
  - **Country/Region** - The country/region where the custodian is located.
  - **Office** - The office location in the custodian's place of business.

- Case information
  - **Hold status** - Indicates if the custodian has been placed on hold.
  - **Communication status**: Indicates if the custodian has been issued a hold notice. If the custodian has been issued a notice, this value of this property is **Published**. If the custodian has not been issued a notice, the status is **Un-published**.
  - **Status** - The status of the custodian within the case. A status of **Active** indicates that the custodian is part of the case. If a custodian is released from a case, the status is changed to **Released**.

- Data sources and indexing information
  - **Data sources** - Shows the count and type of data sources (mailboxes, sites, and Teams) that are associated with the custodian and are part of the case.
  - **Index updated time** - Indicates the time and date for when the advanced indexing job was last triggered. This property will also indicate when the advanced indexing process is currently in progress.

Edit a custodian

As your case progresses, you may discover that there may be additional data sources relevant to a specific custodian & your case. In other scenarios, you may want to remove certain data sources that have been reviewed and deemed as not relevant.

To update the data sources that are associated with a custodian:

1. Go to **eDiscovery > Advanced eDiscovery** and open the case.
2. Click the **Sources** tab.
3. On the **Custodians** page, select a custodian from the list and click **Edit** on the flyout page.

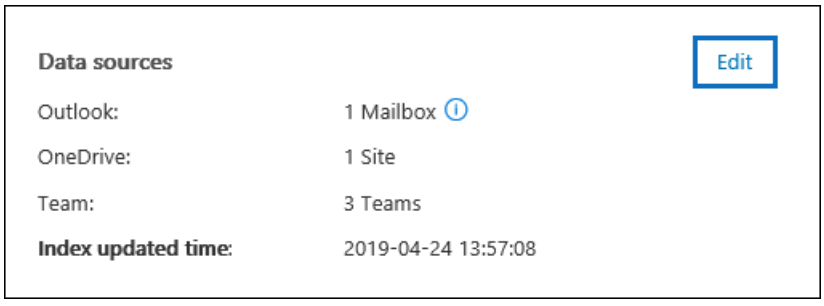

4. Click **Choose data sources** tab to change the settings for the custodian's Exchange mailbox and OneDrive account, click **Choose data sources**.
5. Click the **Select additional data sources** tab to add or remove Teams, SharePoint, or Exchange mailboxes associated with the custodian.

   For more information about data sources associated with a custodian, see Add custodians to a case.
6. Click **Place custodial holds** to enable or disable the hold for the custodian.

Re-index custodian data

In most eDiscovery workflows for legal investigations, a subset of a custodian's data is searched after the custodian is added to a legal case. Because of very large file sizes or possible data corruption, some items in the data sources associated with a custodian may be partially indexed. Using the advanced indexing capability in the Advanced eDiscovery, most partially indexed items can be automatically remediated by re-indexing these items on demand.

When a custodian is added to a case, the data located in the data sources associated with the custodian is automatically re-indexed (by the advanced indexing process). This means you can leave the data in-place instead of having to download and remediate it and then search it offline). However, during the lifecycle of a legal case new data sources might be associated with a custodian. In this case, you can re-index the custodian's data by re-running the advanced indexing process to remediate any partially indexed items and update the index for the custodian's data.

To trigger the re-indexing process to address partially indexed items:

1. Go to **eDiscovery > Advanced eDiscovery** and open the case.
2. Click the **Sources** tab.
3. On the **Custodians** page, select a custodian whose data must be reindexed.
4. On the flyout page, click **Update index**.

   A dialog is displayed saying the index job has been created.

Re-indexing custodian data is a long-running process; the corresponding job that's created is named **Re-indexing custodian data**. You can track the progress on the **Jobs** tab or on the **Custodians** tab by monitoring the status in the **Indexing job status** column.

For more information, see:

- Work with processing errors
- Manage jobs

Release a custodian from a case

A custodian is released in situations where a case is closed, the custodian is no longer under obligation to preserve content for a case, or when the custodian is deemed to no longer be relevant to the case.

If you release a custodian after a hold notice was published, a release notice will be sent to the custodian. Additionally, any holds placed on data sources that were associated with the custodian are removed. If the custodian was placed on a *silent hold*, where they weren't issued any legal hold notifications, a release notice will not be sent but any holds placed on data sources that were associated with that custodian are removed.

To release a custodian:

1. Go to **eDiscovery > Advanced eDiscovery** and open the case.

2. Click the **Sources** tab.
3. On the **Custodians** page, and then select the custodian who is being released from the case.
4. On the flyout page, click **Release custodian**.

   A warning page is displayed explaining that if a hold is placed on a data source associated with the custodian, the hold will be removed, and that any other hold associated with a different Advanced eDiscovery case will still apply. That includes other types of preservation and retention features (such as a Microsoft 365 retention policy).
5. Click **Yes** to confirm that you want to release the custodian.

   The status for this user on the **Custodians** tab is set to **Released** and the **Hold status** on the flyout page is changed to **False**.

> **NOTE**
> A custodian might be simultaneously involved in several legal cases. When a custodian is released from a case, the holds and notifications across other matters won't be impacted.

Bulk-edit custodians

You can use the bulk editor to edit multiple custodians as the same time. To do this, just select two or more custodians on the **Custodians** tab to display the bulk editor and then click one of tasks.

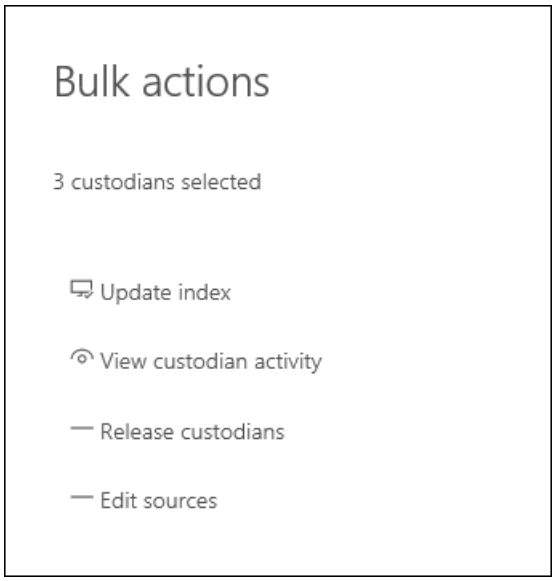

View custodian audit activity

11/2/2020 • 6 minutes to read • Edit Online

Need to find if a user viewed a specific document or purged an item from their mailbox? Advanced eDiscovery is now integrated with the existing audit log search tool in the Security & Compliance Center. Using this embedded experience, you can use the Advanced eDiscovery Custodian Management tool to facilitate your investigation by easily accessing and searching the activity for custodians within your case.

Get permissions

You have to be assigned the View-Only Audit Logs or Audit Logs role in Exchange Online to search the audit log. By default, these roles are assigned to the Compliance Management and Organization Management role groups on the Permissions page in the Exchange admin center. To give a user the ability to search the Advanced eDiscovery audit log with the minimum level of privileges, you can create a custom role group in Exchange Online, add the View-Only Audit Logs or Audit Logs role, and then add the user as a member of the new role group. For more information, see Manage role groups in Exchange Online.

> **IMPORTANT**
> If you assign a user the View-Only Audit Logs or Audit Logs role on the Permissions page in the Security & Compliance Center, they won't be able to search the audit log. You have to assign the permissions in Exchange Online. This is because the underlying cmdlet used to search the audit log is an Exchange Online cmdlet.

Step 1: Search the audit log for activities performed by a custodian

1. Go to **eDiscovery > Advanced eDiscovery** and open the case.
2. Click the **Sources** tab.
3. On the **Custodians** page, select a custodian from the list, and then click **View custodian activity** on the flyout page.

   The Custodian activities search page is displayed. Note the custodian you selected in the previous step is displayed in the **Custodian** drop-down box. You can select different custodians in the drop-down box, but you can only search for activities for one custodian at a time.

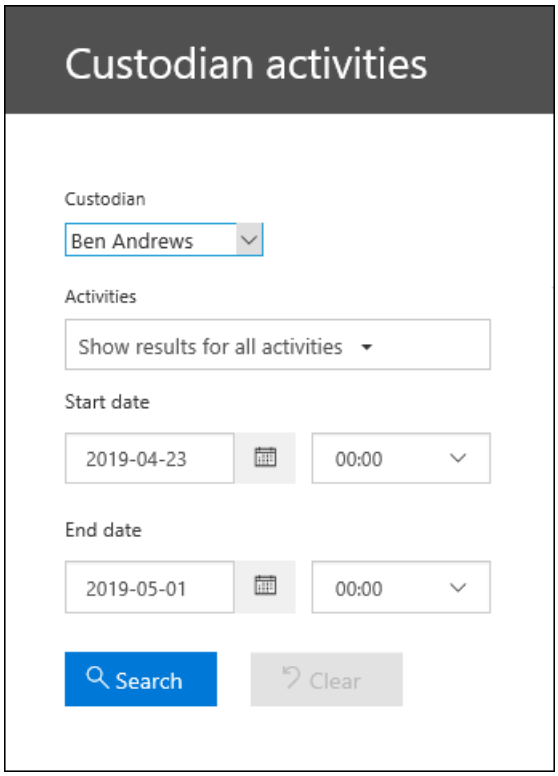

4. Configure the following search criteria:

   a. **Activities** - Click the drop-down list to display the activities that you can search for. After you run the search, only the audit records for the selected activities are displayed. Selecting **Show results for all activities** will display results for all activities performed by the custodian that match the other search criteria.

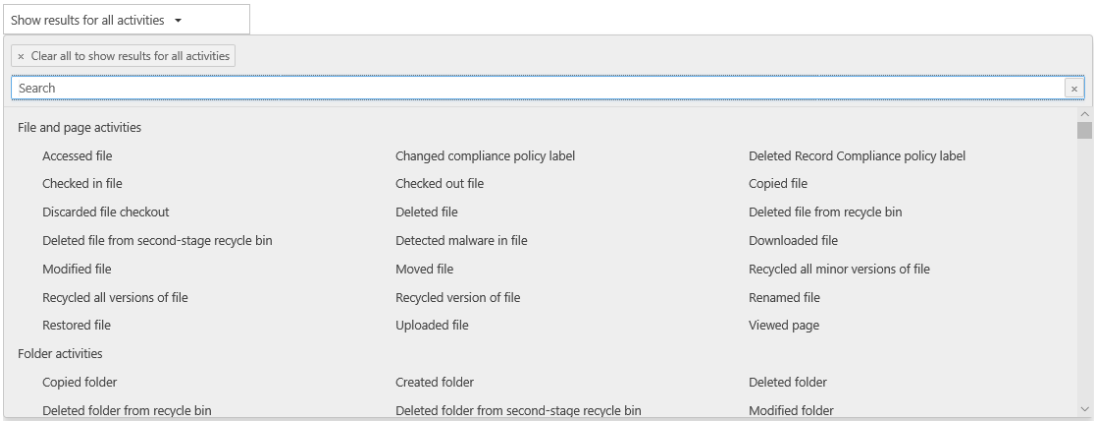

   b. **Start date and End date** - Select a date and time range to display the events that occurred within that period. The last seven days are selected by default. The date and time are presented in Coordinated Universal Time (UTC) format. The maximum date range that you can specify is one year.

   c. **Custodians** - Click in this box and then select a specific custodian to display search results for. Audit records for the selected activity performed by the users you select in this box are displayed in the list of results.

5. Click Search to run the search using your search criteria. The search results are loaded, and after a few moments they are displayed under Results on the Custodian Activities search page.

Step 2: View the audit log search results

The results of an audit log search are displayed under Results on the Custodian Audit log page. A maximum of

5,000 (newest) events are displayed in increments of 150 events. To display more events you can use the scroll bar in the Results pane or you can press Shift + End to display the next 150 events.

The results contain the following information about each event returned by the search.

- **Date**: The date and time (in UTC format) when the event occurred.
- **IP address**: The IP address of the device that was used when the activity was logged. The IP address is displayed in either an IPv4 or IPv6 address format.
- **User**: The user (or service account) who performed the action that triggered the event.
- **Activity**: The activity performed by the user. This value corresponds to the activities that you selected in the Activities drop down list. For an event from the Exchange admin audit log, the value in this column is an Exchange cmdlet.
- **Item**: The object that was created or modified as a result of the corresponding activity. For example, the file that was viewed or modified or the user account that was updated. Not all activities have a value in this column.
- **Detail**: Additional detail about an activity. Again, not all activities will have a value.

Step 3: Filter the search results

In addition to sorting, you can also filter the results of an audit log search. This can help you quickly filter the results for a specific user or activity.

To filter the results:

1. Create and run an audit log search.
2. When the results are displayed, click **Filter results**.
3. Keyword boxes are displayed under each column header.
4. Click one of the boxes under a column header and type a word or phrase, depending on the column you're filtering on. The results will dynamically readjust to display the events that match your filter.
5. To clear a filter, click the **X** in the filter box or just click **Hide filtering**.

Export the search results to a file

You can export the results of an audit log search to a comma separated value (CSV) file on your local computer. You can open this file in Microsoft Excel and use features such as search, sorting, filtering, and splitting a single column (that contains multi-value cells) into multiple columns.

1. Run an audit log search, and then revise the search criteria until you have the desired results.
2. Click Export results and select one of the following options:
   - **Save loaded results:** Choose this option to export only the entries that are displayed under **Results** on the **Custodian Audit log search** page. The CSV file that is downloaded contains the same columns (and data) displayed on the page (Date, User, Activity, Item, and Details). An additional column (titled **More**) is included in the CSV file that contains more information from the audit log entry. Because you're exporting the same results that are loaded (and viewable) on the Audit log search page, a maximum of 5,000 entries are exported.
   - **Download all results:** Choose this option to export all entries from the audit log that meet the search criteria. For a large set of search results, choose this option to download all entries from the audit log in addition to the 5,000 results that can be displayed on the **Custodian Audit log**

search page. This option will download the raw data from the audit log to a CSV file, and contains additional information from the audit log entry in a column named AuditData. It may take longer to download the file if you choose this export option because the file may be much larger than the one that's downloaded if you choose the other option.

> **IMPORTANT**
> You can download a maximum of 50,000 entries to a CSV file from a single audit log search. If 50,000 entries are downloaded to the CSV file, you can probably assume there are more than 50,000 events that met the search criteria. To export more than this limit, try using a date range to reduce the number of audit log entries. You might have to run multiple searches with smaller date ranges to export more than 50,000 entries.

3. After you select an export option, a message is displayed at the bottom of the window that prompts you to open the CSV file, save it to the Downloads folder, or save it to a specific folder

For more information about viewing, filtering, or exporting audit log search results, see Search the audit log in the Security & Compliance Center.

Add non-custodial data sources to an Advanced eDiscovery case

2/18/2021 • 3 minutes to read • Edit Online

In Advanced eDiscovery cases, it doesn't always meet your needs to associate a Microsoft 365 data source with a custodian in the case. But you may still need to associate that data with a case so that you can search it, add it to a review set, and analyze and review it. The feature in Advanced eDiscovery is called *non-custodial data sources* and lets you add data to a case without having to associate it to a custodian. It also applies the same Advanced eDiscovery functionality to non-custodial data that's available for data associated with custodian. Two of the most useful things that you can apply to non-custodial data is placing it on hold and processing it using Advanced indexing.

Add a non-custodial data source

Follow these steps to add and manage non-custodial data sources in an Advanced eDiscovery case.

1. On the **Advanced eDiscovery** home page, click the case that you want to add the data to.
2. Click the **Data sources** tab and then click **Add data source** > **Add data locations**.
3. On the **New non-custodial data locations** flyout page, choose the data sources that you want to add to the case. You can add multiple mailboxes and sites by expanding the **SharePoint** or **Exchange** sections and then clicking **Edit**.

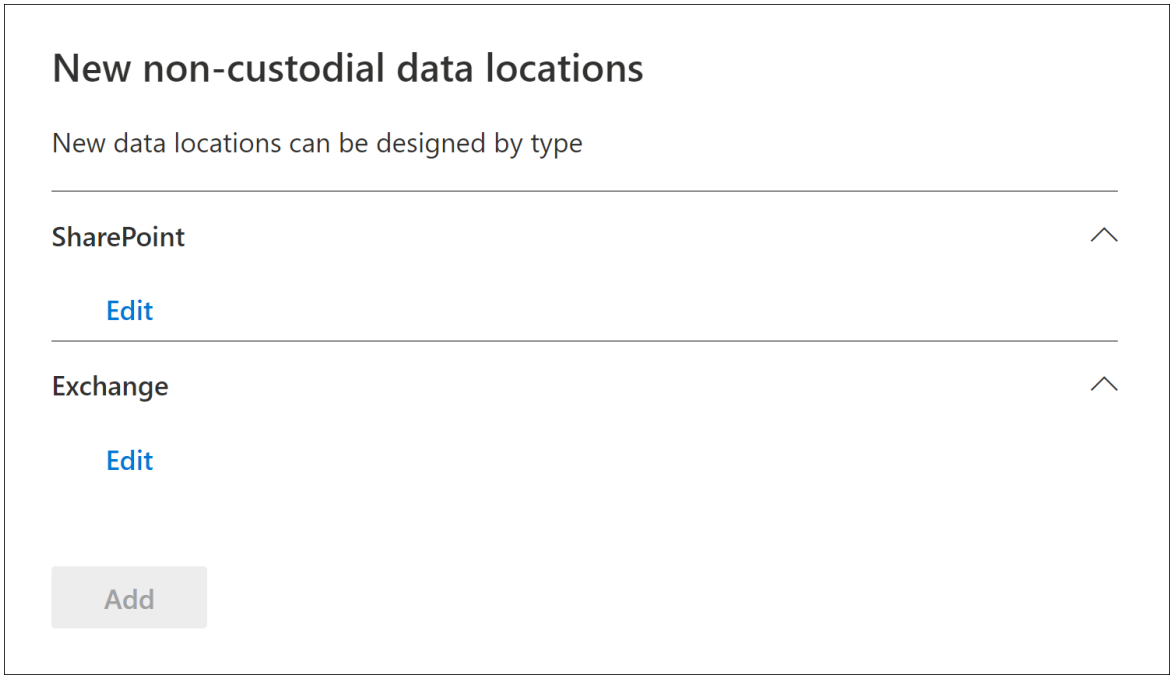

- **SharePoint** - Click **Edit** to add sites. Select a site in the list or you can search for a site by typing the URL of the site in the search bar. Select the sites that you want to add as non-custodian data sources and click **Add**.
- **Exchange** - Click **Edit** to add mailboxes. Type a name or alias (a minimum of three characters) in the search box for mailboxes or distribution groups. Select the mailboxes that you want to add as non-custodian data sources and click **Add**.

> **NOTE**
>
> You can use the **SharePoint** and **Exchange** sections to add sites and mailboxes associated with a Team or Yammer group as non-custodial data sources. You have to separately add the mailbox and site associated with a Team or Yammer group.

4. After you add non-custodial data sources, you have the option to place those locations on hold or not. Select or unselect the **Hold** checkbox next to the data source to place it on hold.

5. Click **Add** at the bottom of the **New non-custodial data locations** flyout page to add the data sources to the case.

   Each non-custodial data source that you added is listed on the **Data sources** page. Non-custodial data sources are identified by the **Data location** value in the **Source type** column.

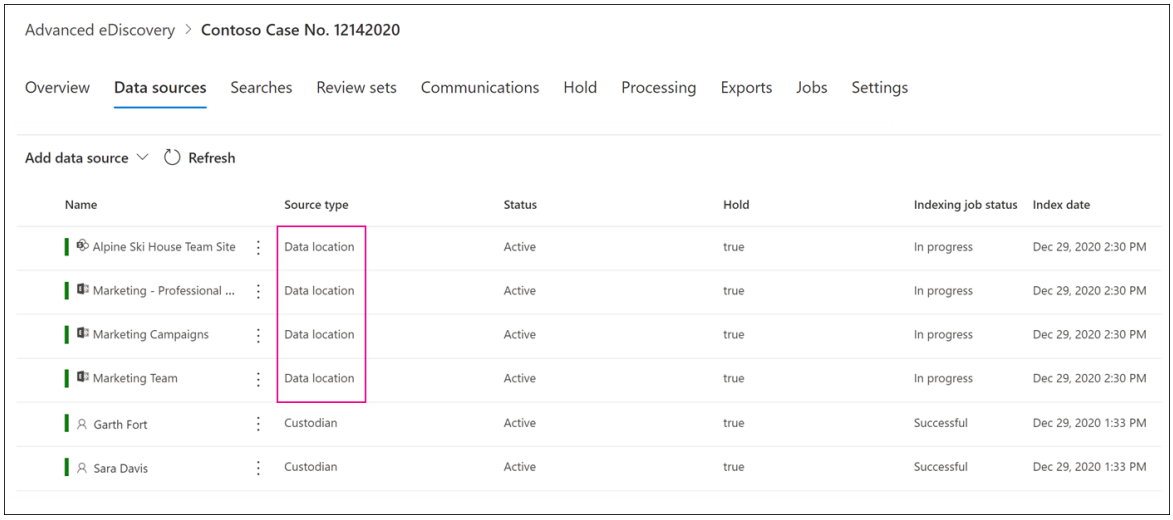

After you add non-custodial data sources to the case, a job named *Reindexing non-custodial data* is created and displayed on the **Jobs** tab of the case. After the job is created, the Advanced indexing process in initiated and the data sources are reindexed.

Manage the hold for non-custodial data sources

After you place a hold on a non-custodial data source, a hold policy that contains the non-custodial data sources for the case is automatically created. When you place other non-custodial data sources on hold, they are added to this hold policy.

1. Open the Advanced eDiscovery case and select the **Hold** tab.

2. Click **NCDSHold-<GUID>**, where the GUID value is unique to the case.

   The flyout page display information and statistics about the non-custodial data sources on hold.

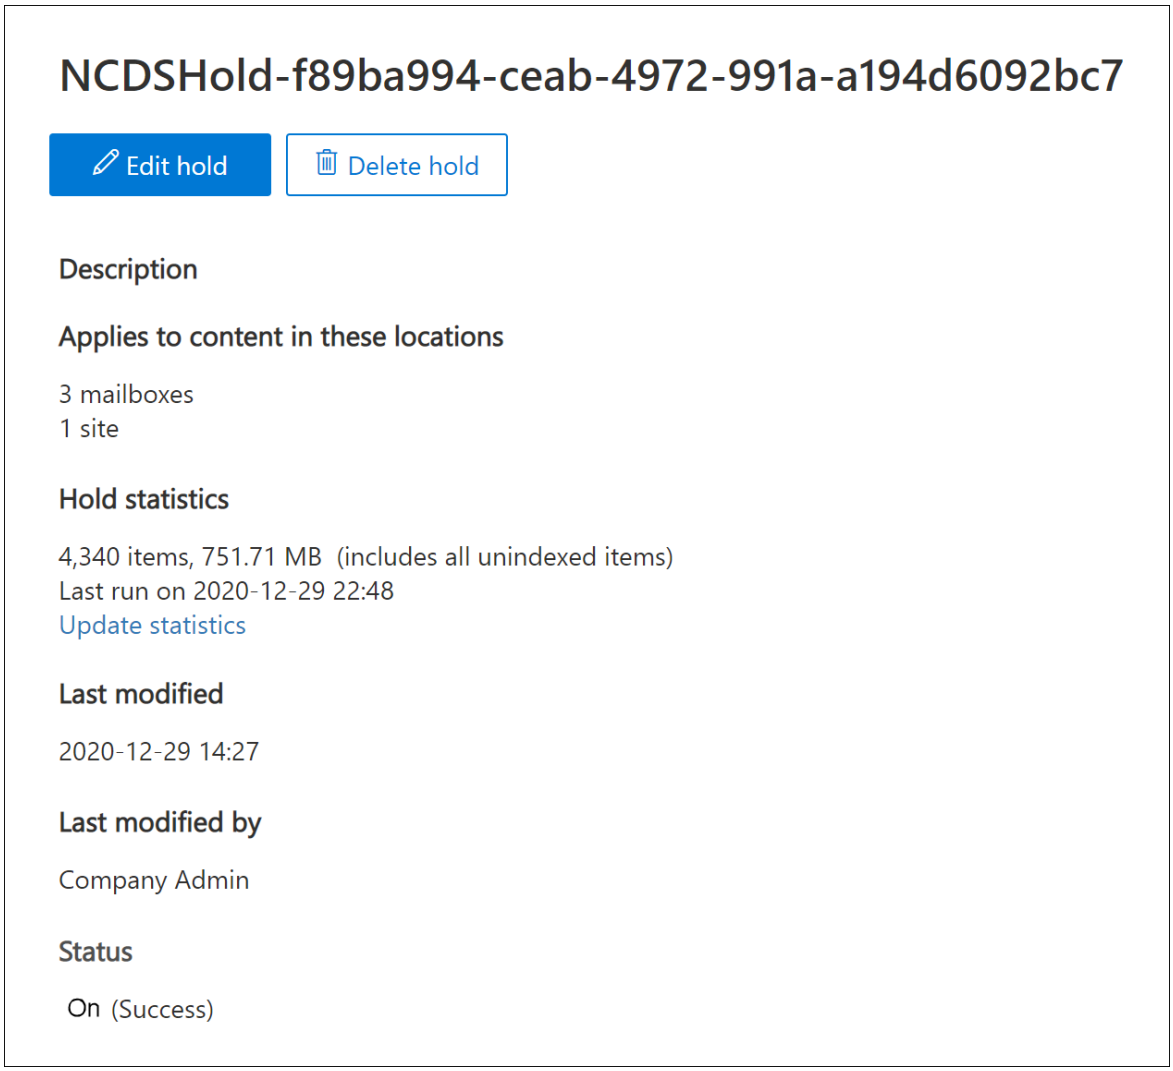

3. Click **Edit hold** to view the non-custodial data sources placed on hold and perform the following management tasks:

   - On the **Locations** page, you can release a non-custodial data source by removing it from the hold. Releasing a data source doesn't remove the non-custodial data source from the case. It only removes the hold that was placed on the data source.
   - On the **Query** page, you can edit the hold to create a query-based hold that is applied to all tha non-custodial data sources in the case.

Work with communications in Advanced eDiscovery

4/18/2020 • 2 minutes to read • Edit Online

Advanced eDiscovery allows legal departments to simplify their processes around tracking and distributing legal hold notifications. The custodian communications tool enables legal departments to manage and automate the entire legal hold process, from initial notifications, to reminders, and to escalations, all in one location.

What is a legal hold notification?

A legal hold (also known as a *litigation hold*) notice is a notification sent from an organization's legal department to employees, contingent staff, or custodians of data that may be relevant to a legal investigation. These notifications instruct custodians to preserve electronically stored information as well as any content that may be relevant to an active or impending legal matter. Legal teams must know that each custodian has received, read, understood, and has agreed to comply with the given instructions.

The legal hold notification process

An organization has a duty to preserve relevant information when it learns about an impending litigation or regulatory investigation. To comply with the preservation requirements of an investigation, the organization should immediately inform potential custodians about their duty to preserve relevant information.

With Advanced eDiscovery, legal teams can create and customize their legal hold notification workflow. The custodian communications tool lets legal teams to configure the following notices and workflows:

1. **Issuance notice:** A legal hold notice is issued (or initiated) by a notification from the legal department to custodians who may have relevant information about the case matter. This notice instructs the custodians to preserve any information that may be needed for discovery.
2. **Re-Issuance notice:** During a case, custodians may be required to preserve additional content (or less content) than was previously requested. For this scenario, you can update the existing hold notice and re-issue it to custodians.
3. **Release notice:** Once a matter is resolved and the custodian is no longer subject to a preservation requirement, the custodian can be released from the case. Additionally, you can notify the custodian that they are no longer required to preserve content, and provide instructions about how to resume their normal work activity with regard to their data.
4. **Reminders and escalations:** In some instances, just issuing a notice isn't enough to satisfy legal discovery requirements. With each notification, legal teams can schedule a set of reminder and escalation workflows to automatically follow up with unresponsive custodians.
    - **Reminders:** After a legal hold notice has been issued or re-issued to a set of custodians, an organization can set up reminders to alert unresponsive custodians.
    - **Escalations:** In some cases, if a custodian remains unresponsive even after a set of reminders over a period of time, the legal team can set up an escalation workflow to notify unresponsive custodians and their manager.

For more information about managing the custodian communication process, see the following:

- Create a legal hold notice
- Use the communications editor

- Manage hold notifications
- Acknowledge a hold notification

Create a legal hold notice

2/18/2021 • 7 minutes to read • Edit Online

Using Advanced eDiscovery custodian communications, organizations can manage their workflow around communicating with custodians. Through the Communications tool, legal teams can systematically send, collect, and track legal hold notifications. The flexible creation process also allows teams to customize the hold notification workflow and the content in the notices sent to custodians.

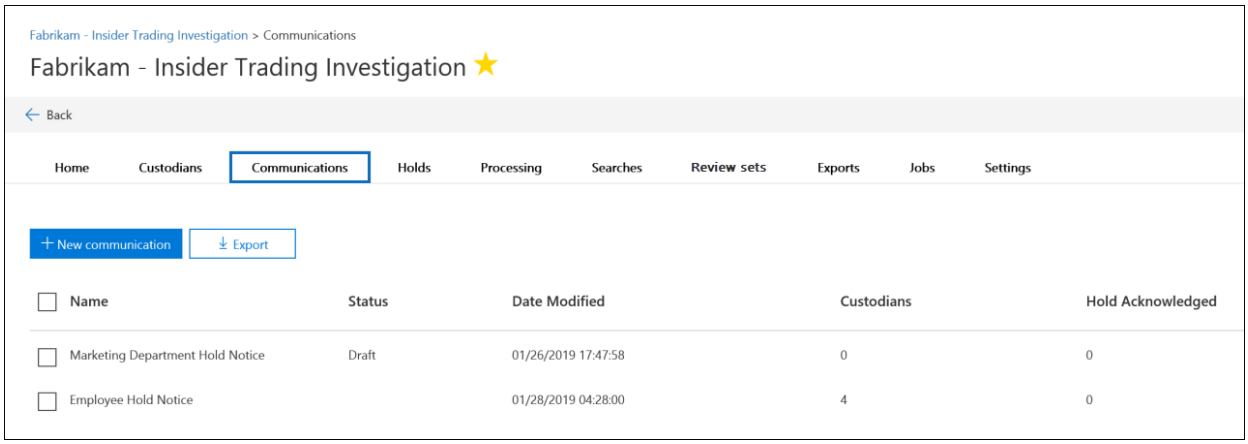

The article outlines the steps in the hold notification workflow.

Step 1: Specify communication details

The first step is to specify the appropriate details for legal hold notices or other custodian communications.

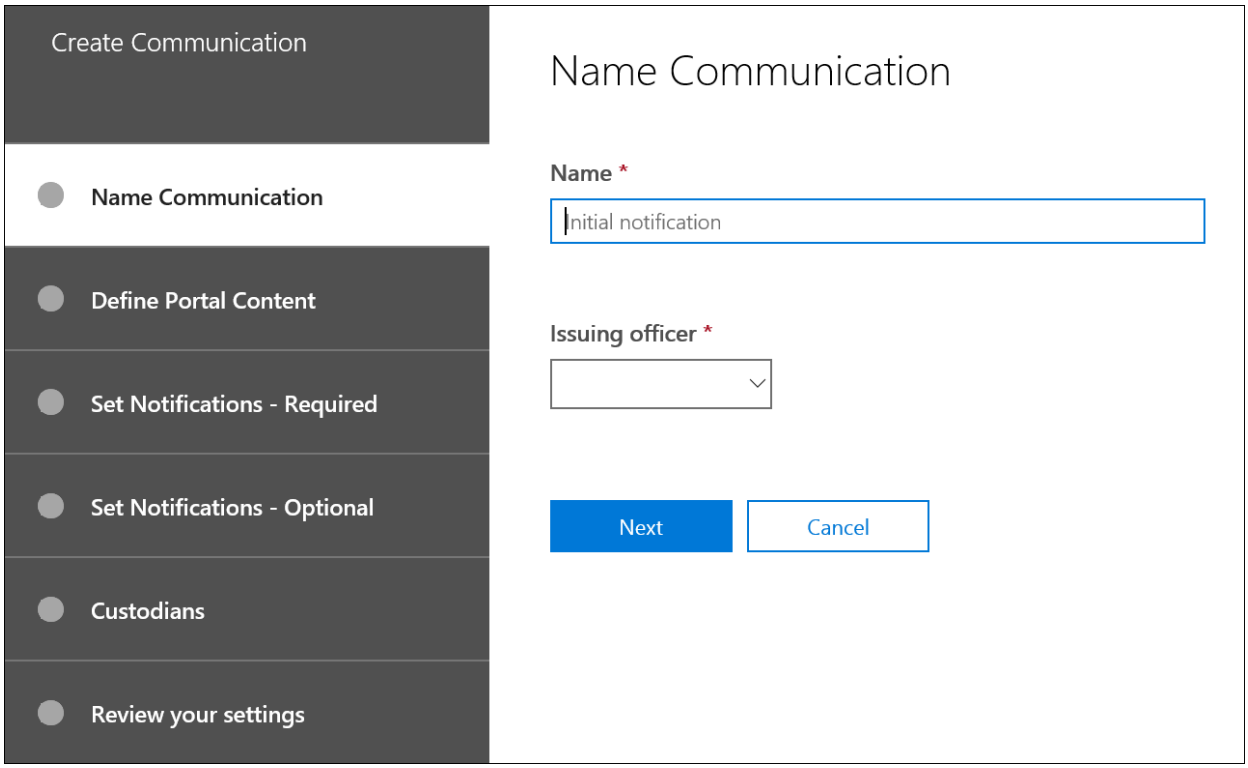

1. In the Security & Compliance Center, go to **eDiscovery** > **Advanced eDiscovery** to display the list of cases in your organization.
2. Select a case, click the **Communications** tab, and then click **New communication**.

3. On the **Name communication** page, specify the following (required) communication details.

   - **Name**: This is the name for the communication.
   - **Issuing officer**: The dropdown list displays a list of case members. For more information on how to add new members to a case, see Create an Advanced eDiscovery case. Each notice sent to custodians will be sent on behalf of the specified issuing officer.

> **NOTE**
> The issuing officer must have an **active mailbox** to show up in the Issuing Officer dropdown

4. Click **Next**.

Step 2: Define the portal content

Next, you can create and add the content of the hold notice. On the **Define portal content** page in the **Create communication** wizard, specify the contents of the hold notice. This content will be automatically appended to the Issuance, Re-Issue, Reminder, and Escalation notices. Additionally, this content will appear in the custodian's Compliance Portal.

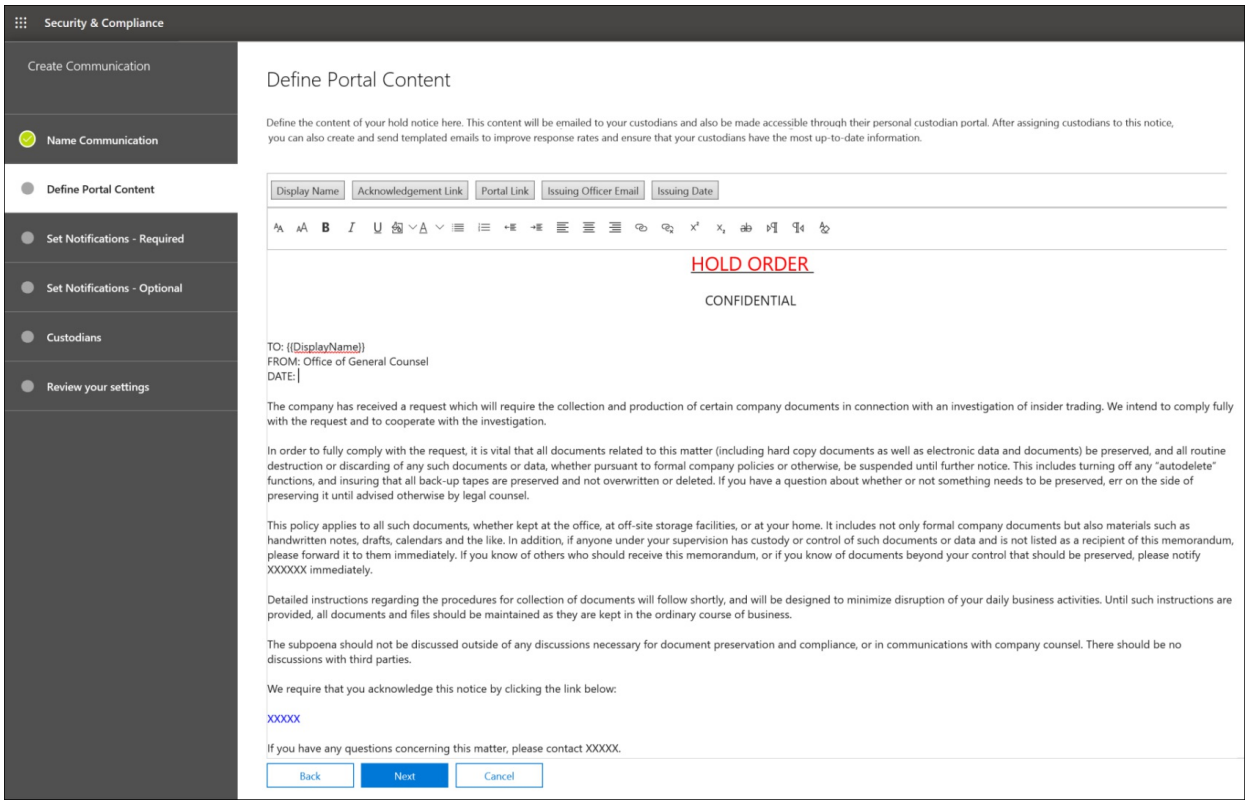

To create the portal content:

1. Type (or cut and paste from another document) your hold notice in the textbox for the portal content.
2. Insert merge variables into your notice to customize the notice and share the Custodian Compliance Portal.
3. Click **Next**.

TIP

To learn more about how to can customize the content and format of the portal content, see Use the Communications Editor.

Step 3: Set the required notifications

After you've defined the contents of the hold notice, you can set up the workflows around sending and managing the notification process. Notifications are email messages that are sent to notify and follow up with custodians. Every custodian added to the communication will receive the same notification.

To set up and send a hold notice, you must include Issuance, Re-Issuance, and Release notifications.

Issuance notification

After the communication is created, the **Issuance Notification** is initiated by the specified Issuing Officer. The Issuance notification is the first communication sent to the custodian to inform them about their preservation obligations.

To create an issuance notification:

1. In the **Issuance** tile, click **Edit**.
2. If necessary, add additional case members or staff to the **Cc** and **Bcc** fields. To add multiple users to these fields, separate email addresses with a semi-colon.
3. Specify the **Subject** for the notice (required).
4. Specify the contents or additional instructions that you would like to provide to the custodian (required). The portal content you defined in Step 2 is added to the end of the issuance notice.
5. Click **Save**.

Re-Issuance notification

As the case progresses, custodians may be required to preserve additional or less data than was previously instructed. After you update the portal content, the re-issuance notification is sent and alerts custodians about any changes to their preservation obligations.

To create a re-issuance notification:

1. In the **Reissue** tile, click **Edit**.
2. If necessary, add additional case members or staff to the **Cc** and **Bcc** fields. To add multiple users to these fields, separate email addresses with a semi-colon.
3. Specify the **Subject** for the notice (required).
4. Specify the contents or additional instructions that you would like to provide to the custodian (required). The portal content you defined in Step 2 is added to the end of the re-issuance notice.
5. Click **Save**.

> **NOTE**
>
> If the portal content is modified (on the **Define Portal Content** page in the **Edit communication** wizard), the re-issuance notification will be automatically sent to all custodians assigned to the notice. After the notification is sent, custodians will be asked to re-acknowledge their hold notice. If you have set up any reminder or escalation workflows, these will also re-start. For more information about what other case management events trigger communications, see Events that trigger notifications.

Release notification

After a matter is resolved or if a custodian is no longer subject to preserve content, you can release the custodian from a case. If the custodian was previously issued a hold notice, the release notification can be used to alert custodians that they have been released from their obligation.

To create a release notification:

1. In the **Release** tile, click **Edit**.
2. If necessary, add additional case members or staff to the **Cc** and **Bcc** fields. To add multiple users to these fields, separate email addresses with a semi-colon.
3. Specify the **Subject** for the notice (required).
4. Specify the contents or additional instructions that you would like to provide to the custodian (required).
5. Click **Save** and go to the next step.

(Optional) Step 4: Set the optional notifications

Optionally, you can simplify the workflow for following up with unresponsive custodians by creating and scheduling automated reminder and escalation notifications.

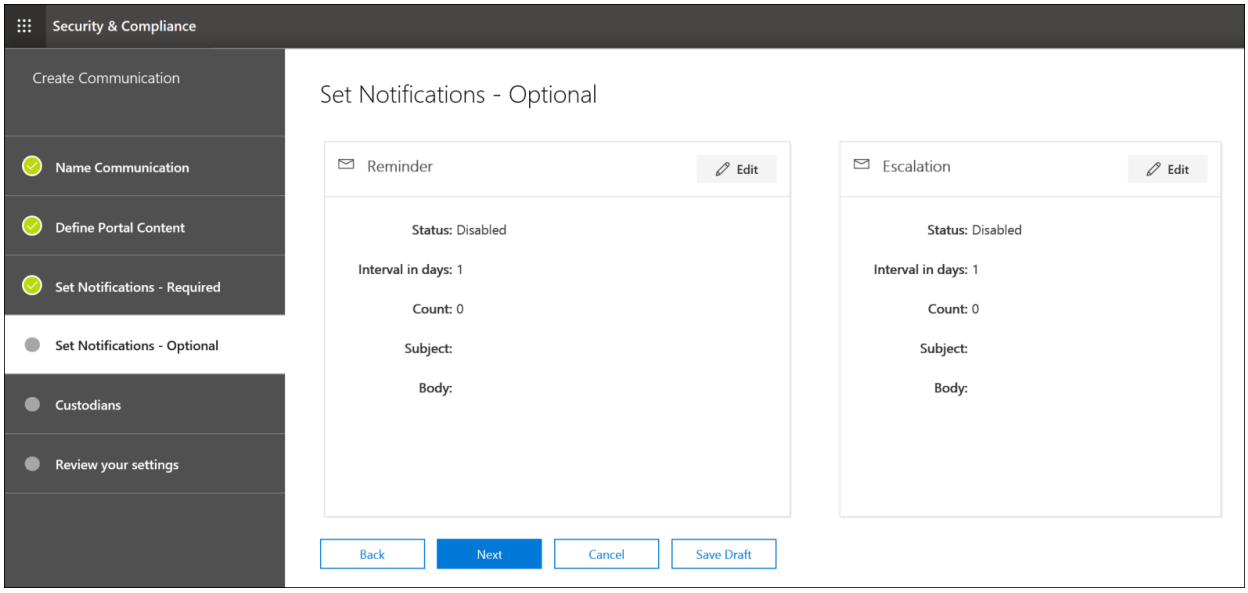

Reminders

After you have sent a hold notification, you can follow up with unresponsive custodians by defining a reminder workflow.

To schedule reminders:

1. In the **Reminder** tile, click **Edit**.
2. Enable the **Reminder** workflow by turning on the **Status** toggle (required).

3. Specify the **Reminder interval (in days)** (required). This is the number of days to wait before sending the first and follow-up reminder notifications. For example, if you set the reminder interval to seven days, then the first reminder would be sent seven days after the hold notification was initially issued. All subsequent reminders would also be sent every seven days.
4. Specify the **Number of reminders** (required). This field specifies how many reminders to send to unresponsive custodians. For example, if you set the number of reminders to 3, then a custodian would receive a maximum of three reminders. After a custodian acknowledges the hold notification, reminders will no longer be sent to that user.
5. Specify the **Subject** for the notice (required).
6. Specify the contents or additional instructions that you would like to provide to the custodian (required). The portal content you defined in Step 2 is added to the end of the reminder notice.
7. Click **Save** and go the next step.

Escalations

In some situations, you may need additional ways to follow up with unresponsive custodians. If a custodian doesn't acknowledge a hold notification after receiving the specified number of reminders, the legal team can specify a workflow to automatically send an escalation notice to the custodian and their manager.

To schedule escalations:

1. In the **Escalation** tile, click **Edit**.
2. Enable the **Escalation** workflow by turning on the **Status** toggle.
3. Specify the **Escalation interval (in days)** (required).
4. Specify the **Number of escalations** (required). This field specifies how many escalations to send to unresponsive custodians. For example, if you set the number of escalations to 3, then an escalation notice would be sent to the custodian and their manager a maximum of three times. After a custodian acknowledges the hold notification, escalations will no longer be sent.
5. Specify the **Subject** for the notice (required).
6. Specify the contents or additional instructions that you would like to provide to the custodian (required). The portal content you defined in Step 2 is added to the end of the escalation notice.
7. Click **Save** and go the next step.

Step 5: Assign custodians to receive notifications

After you have finalized the content for notifications, select the custodians that you would like to send notifications to.

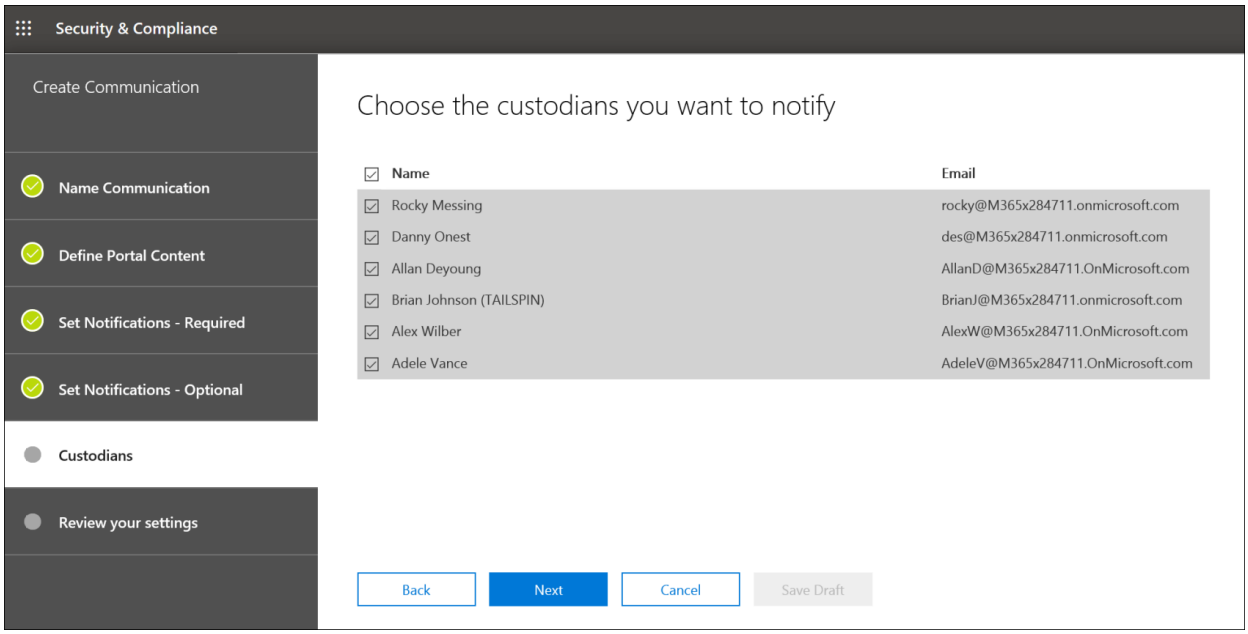

To add custodians:

1. Assign custodians to the communication by clicking the checkbox next to their name.

   After the communication is created, the notification workflow will automatically apply to the selected custodians.

2. Click **Next** to review the communication settings and details.

> **NOTE**
> You can only add custodians who have been added to the case and haven't been sent another notification within the case.

Step 6: Review settings

After you review the settings and click **Send** to complete the communication, the system will automatically start the communication workflow by sending the issuance notice.

Events that trigger notifications

The following table describes events in the case management process that trigger when the different types of notifications are sent to custodians.

TYPE OF COMMUNICATION	TRIGGER
Issuance notices	The initial creation of the notification. You can also manually resend a hold notification.
Re-issuance notices	Updating the portal content on the **Define Portal Content** page in the **Edit communication** wizard.
Release notices	The custodian is released from the case.
Reminders	The interval and number of reminders configured for the reminder.

TYPE OF COMMUNICATION	TRIGGER
Escalations	The interval and number of reminders configured for the escalation.

Use the communications editor

2/18/2021 • 2 minutes to read • Edit Online

As you define the content of your portal content, legal hold notifications, and related reminders/escalations, you can use the Communications Editor to format and dynamically customize your content.

Rich text editor

The Communications Editor allows user to customize the text using the editor options. For example, users can change font types, create bulleted lists, highlight content, and more.

Merge field variables

You can use email merge variables from the Communications Editor to embed customized custodian attributes into a communication's body text. When sent to the custodian, the merge field will be populated with the corresponding field. For example, when sent to custodian John Smith, the merge field [Custodian Name] would be translated with the corresponding name.

You can use email merge fields by selecting the **Merge field** icons on the top of the rich-text editor control. The placeholder will be added based off the location of the users' cursor.

List of merge field variables

FIELD NAME	FIELD DETAILS	
Display Name	The custodian's first and last name.	
Acknowledgment Link	A customized link to record each custodian's acknowledgment.	
Portal Link	A customized link for the custodian's Compliance Portal.	
Issuing Officer	The email address of the specified issuing officer.	
Issuing Date	The date that the notice was issued (UTC).	

Manage hold notifications

2/25/2020 • 2 minutes to read • Edit Online

After you have initiated your legal hold notification workflow, you can use the communications workflow in Advanced eDiscovery to track the status of your communications. The Communications tab contains a list of all notifications within your Advanced eDiscovery case. You can see details such as the number of custodians that have been assigned or have acknowledged the notice.

Monitor acknowledgments

After you select a communication from the **Communications** tab, you can view a list of custodians that have acknowledged a hold notice.

1. In the compliance center, go to **eDiscovery > Advanced eDiscovery**.
2. Select a case and then click the **Communications** tab.
3. Select a communication to display the **Custodian communication** flyout page.

A list of custodians associated with the selected communication is displayed on the communication flyout page. This page also displays insights and about how many acknowledgments were received and how many are outstanding. The page also shows which custodians have sent an acknowledgment that they received the hold notification.

Re-send a hold notice

Occasionally, custodians lose track of email messages in their day-to-day work. Or for a long-running litigation case, a custodian may contact you or others and request that you re-send a notice. As you manage the communications workflow for legal hold notices, you may need to re-send a notice to bring it back to the "top of a user's mailbox".

To re-send a hold notice to a custodian:

1. In Advanced eDiscovery, select a case and then click the **Communications** tab.
2. Select a communication to display the **Custodian communication** flyout page.
3. Click **More > Re-send hold notice**.
4. On the **Re-send hold notice** flyout page, select the custodians that you what to re-send the notice and type an optional reason.
5. Click **Re-send** to send the notice to the selected custodians.

If a custodian hasn't acknowledged the hold notification, the reminder and escalation workflow is restarted. If a custodian has acknowledged the hold notice, the custodian will receive a copy of the original hold notice.

> **NOTE**
> You can only resend a legal hold notification to custodians that are assigned to the communication.

Update preservation requirements

As the case progresses, custodians may be required to preserve additional or less data than was previously instructed. In eDiscovery terms, you need to re-issue the hold notice with updated content.

To update the contents of the initial hold notice:

1. In Advanced eDiscovery, select a case and then click the **Communications** tab.
2. Select the hold notice that you want to update and click **Edit** on the **Custodian communication** flyout page.
3. In the **Edit Communication** wizard, click **Define Portal Content** in the left pane of the wizard, and update the contents of the notice.
4. Click **Save**.

The re-issuance notice will be sent to all the custodians assigned to the legal hold notification. In addition, if the Reminder or Escalation notice is enabled, then the workflows for those types of notices will restart.

Update legal hold notifications and settings

When you update the content or settings of the Issuance, Release, Reissue, Reminder, or Escalation notice, these changes will apply to all future communications generated by the workflow.

More information

- Add custodians to a case
- Create a legal hold notice
- Acknowledge a hold notification

Acknowledge a hold notification

5/5/2020 • 2 minutes to read • Edit Online

When responding to a regulatory request or investigation, you may be required to inform custodians of their obligation to preserve electronically stored information (ESI) and any material that may be relevant to an active or imminent legal matter. Once sent, legal teams must know that each custodian has received, read, understood, and agreed to follow the given instructions.

To help reduce the time, cost, and effort of following up with your custodians, Advanced eDiscovery allows you to send and follow up on legal hold notifications through email. In addition to email notices, each custodian will have access to an individualized Compliance Portal, allowing custodians to be kept informed of changes to their obligation status.

Email notifications

After a Legal Hold Notification has been issued, each custodian will receive a unique and personalized email containing your defined legal hold notice and added instructions.

> **TIP**
> See how you can use the built-in Communication Editor to allow your custodians to acknowledge their notice or access their Compliance Portal directly from their email.

Based on the configuration of your legal hold notification, your custodians may receive the following notices:

- **Issuance notice:** The first notice sent to your custodian. This notice will contain your issuance instructions and the hold notice appended to the end of your message.
- **Reminder notice:** If enabled, a reminder notice will be sent to your custodians based on the specified frequency and interval. The reminders will continue to be sent either until the custodian has acknowledged their notice or until the number of reminders have been exhausted.
- **Escalation notice:** If enabled, an escalation notice will be sent to your custodian and their manager after the reminder notices have been exhausted. The system will automatically send escalation notices until the specified number of escalations have been completed or until the custodian acknowledges their hold notification.
- **Reissue notice:** During the course of an investigation, if the contents of the hold notice are updated, then the updated notice will automatically be sent to the custodian.
- **Release notice:** When a custodian is released from the case, they'll be sent the release notice.

Compliance Portal

In addition to the email notifications, each custodian will have access to a unique Compliance Portal. Through the portal, each custodian can view, access, and acknowledge their active hold notifications.

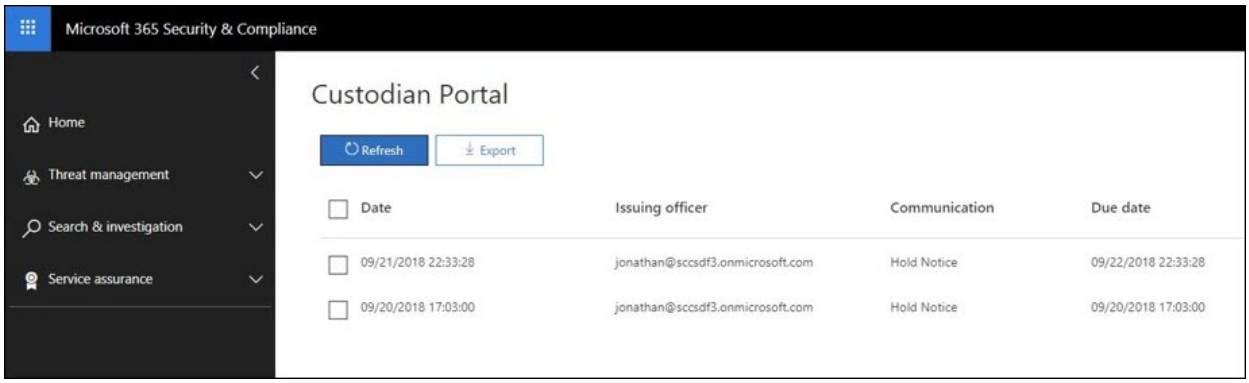
Microsoft 365 Security & Compliance
Home
Threat management
Search & investigation
Service assurance
Custodian Portal
Refresh
Export
Date
Issuing officer
Communication
Due date
09/21/2018 22:33:28
jonathan@sccsdf3.onmicrosoft.com
Hold Notice
09/22/2018 22:33:28
09/20/2018 17:03:00
jonathan@sccsdf3.onmicrosoft.com
Hold Notice
09/20/2018 17:03:00

Manage holds in Advanced eDiscovery

11/2/2020 • 10 minutes to read • Edit Online

You can use an Advanced eDiscovery case to create holds to preserve content that might be relevant to your case. Using the Advanced eDiscovery hold capabilities, you can place holds on custodians and their data sources. Additionally, you can place a non-custodial hold on mailboxes and OneDrive for Business sites. You can also place a hold on the group mailbox, SharePoint site, and OneDrive for Business site for an Microsoft 365 Group. Similarly, you can place a hold on the mailbox and site that are associated with Microsoft Teams. When you place content locations on hold, content is held until you release the custodian, remove a specific data location, or delete the hold policy entirely.

Manage custodian-based holds

In some cases, you may have a set of custodians that you have identified and have decided to preserve their data during the case. In Advanced eDiscovery, when these custodians are placed on hold, the user and their selected data sources are automatically added to a custodian hold policy.

To view the custodian hold policy:

1. In the Microsoft 365 compliance center, click **eDiscovery** > **Advanced** to display the list of cases in your organization.
2. Go to the **Sources** tab to add custodians within your case. To learn how you can add and place custodians on hold within an Advanced eDiscovery case, see Add Custodians to a case. If you have already added custodians and placed them on hold, go to step 3.
3. Go to the **Holds** tab and click **CustodianHold<HoldId>**.
4. On the flyout page, you can see hold statistics for the policy. You can also perform actions like apply a query to your custodian-based hold. For more information about creating a hold query and using conditions, see Keyword queries and search conditions for Content Search.

Manage non-custodial holds

When you create a hold, you have the following options to scope the content that is held in the specified content locations:

- You create an infinite hold where all content is placed on hold. Alternatively, you can create a query-based hold where only content that matches a search query is placed on hold.
- You can specify a date range to hold only the content that was sent, received, or created within that date range. Alternatively, you can hold all content regardless of when it was sent, received, or created.

To create a non-custodial hold for an Advanced eDiscovery case:

1. In the Microsoft 365 compliance center, click **eDiscovery** > **Advanced** to display the list of cases in your organization.
2. Click **Open** next to the case that you want to create the holds in.
3. From the home page for the case, click the **Holds** tab.
4. On the **Holds** tab, click **Create**.

5. On the **Name your hold** page, give the hold a name. The name of the hold must be unique in your organization.
6. (Optional) In the **Description** box, add a description of the hold.
7. Click **Next**.
8. Choose the content locations that you want to place on hold. You can place mailboxes, sites, and public folders on hold.

    a. **Exchange email** - Click **Choose users, groups, or teams** and then click **Choose users, groups, or teams** again to specify mailboxes to place on hold. Use the search box to find user mailboxes and distribution groups (to place a hold on the mailboxes of group members) to place on hold. You can also place a hold on the associated mailbox for an Microsoft 365 Group or a Microsoft Team. Select the user, group, team check box, click **Choose**, and then click **Done**.

    > **NOTE**
    > When you click **Choose users, groups, or teams** to specify mailboxes to place on hold, the mailbox picker that's displayed is empty. This is by design to enhance performance. To add people to this list, type a name (a minimum of 3 characters) in the search box.

    b. **SharePoint Sites** - Click **Choose sites** and then click **Choose sites** again to specify SharePoint and OneDrive for Business sites to place on hold. Type the URL for each site that you want to place on hold. You can also add the URL for the SharePoint site for an Microsoft 365 Group or a Microsoft Team. Click **Choose**, and then click **Done**.

    See the **FAQ** section for tips on putting Microsoft 365 Groups and Microsoft Teams on hold.

    > **NOTE**
    > The URL for a user's OneDrive account includes their user principal name (UPN) (for example, `https://alpinehouse-my.sharepoint.com/personal/sarad_alpinehouse_onmicrosoft_com`). In the rare case that a person's UPN is changed, their OneDrive URL will also change to incorporate the new UPN. If a user's OneDrive account is part of a non-custodial hold and their UPN is changed, you need to update the hold and point to the new OneDrive URL. For more information, see How UPN changes affect the OneDrive URL.

    c. **Exchange public folders** - Move the toggle switch to the All position to put all public folders in your Exchange Online organization on hold. Note that you can't choose specific public folders to put on hold. Leave the toggle switch set to **None** if you don't want to put a hold on public folders.
9. When you're done adding content locations to the hold, click **Next**.
10. To create a query-based hold with conditions, complete the following. Otherwise, just click **Next**.

    - In the box under **Keywords**, type a search query in the box so that only the content that meets the search criteria is placed on hold. You can specify keywords, message properties, or document properties, such as file names. You can also use more complex queries that use a Boolean operator, such as AND, OR, or NOT. If you leave the keyword box empty, then all content located in the specified content locations will be placed on hold.
    - Click **Add** conditions to add one or more conditions to narrow the search query for the hold. Each condition adds a clause to the KQL search query that is created and run when you create the hold. For example you can specify a date range so that email or site documents that were created within the date ranged are placed on hold. A condition is logically connected to the keyword query

(specified in the keyword box) by the AND operator. That means that items have to satisfy both the keyword query and the condition to be placed on hold.

For more information about creating a search query and using conditions, see Keyword queries and search conditions for Content Search.

11. After configuring a query-based hold, click **Next**.
12. Review your settings, and then click **Create this hold**.

View hold statistics

After some time, information about the new hold is displayed in the details pane on the **Holds** tab for the selected hold. This information includes the number of mailboxes and sites on hold and statistics about the content that was placed on hold, such as the total number and size of items placed on hold and the last time the hold statistics were calculated. These hold statistics help you identify how much content that's related to the eDiscovery case is being held.

Keep the following things in mind about hold statistics:

- The total number of items on hold indicates the number of items from all content sources that are placed on hold. If you've created a query-based hold, this statistic indicates the number of items that match the query.
- The number of items on hold also includes unindexed items found in the content locations. Note that if you create a query-based hold, all unindexed items in the content locations are placed on hold. This includes unindexed items that don't match the search criteria of a query-based hold and unindexed items that might fall outside of a date range condition. This is different than what happens when you run a Content Search, in which unindexed items that don't match the search query or are excluded by a date range condition aren't included in the search results. For more information about unindexed items, see Partially indexed items in Content Search in Office 365.
- You can get the latest hold statistics by clicking Update statistics to re-run a search estimate that calculates the current number of items on hold.
- If necessary, click Refresh in the toolbar to update the hold statistics in the details pane.
- It's normal for the number of items on hold to increase over time because users whose mailbox or site is on hold are typically sending or receiving new email message and creating new SharePoint and OneDrive for Business documents.
- If a SharePoint site or OneDrive account is moved to a different region in a multi-geo environment, the statistics for that site won't be included in the hold statistics. However, the content in the site will still be on hold. Also, if a site is moved to a different region the URL that's displayed in the hold will not be updated. You'll have to edit the hold and update the URL.

Place a hold on Microsoft Teams and Office 365 Groups

Microsoft Teams are built on Office 365 Groups. Therefore, placing them on hold in Advanced eDiscovery is very similar.

- **How do I map an additional Microsoft 365 Groups or Microsoft Teams site to a custodian? And what about placing a non-Custodial hold on Microsoft 365 Groups and Microsoft Teams?** Microsoft Teams are built on Microsoft 365 Groups. Therefore, placing them on hold in an eDiscovery case is very similar. Keep the following things in mind when placing Microsoft 365 Groups and Microsoft Teams on hold.
  - To place content located in Microsoft 365 Groups and Microsoft Teams on hold, you have to

specify the mailbox and SharePoint site that associated with a group or team.

- Run the **Get-UnifiedGroup** cmdlet in Exchange Online to view properties for an Microsoft 365 Group or Microsoft Team. This is a good way to get the URL for the site that's associated with an Microsoft 365 Group or a Microsoft Team. For example, the following command displays selected properties for an Microsoft 365 Group named Senior Leadership Team:

```
Get-UnifiedGroup "Senior Leadership Team" | FL
DisplayName,Alias,PrimarySmtpAddress,SharePointSiteUrl
DisplayName            : Senior Leadership Team
Alias                  : seniorleadershipteam
PrimarySmtpAddress     : seniorleadershipteam@contoso.onmicrosoft.com
SharePointSiteUrl      : https://contoso.sharepoint.com/sites/seniorleadershipteam
```

> **NOTE**
> To run the Get-UnifiedGroup cmdlet, you have to be assigned the View-Only Recipients role in Exchange Online or be a member of a role group that's assigned the View-Only Recipients role.

- When a user's mailbox is searched, any Microsoft 365 Group or Microsoft Team that the user is a member of won't be searched. Similarly, when you place an Microsoft 365 Group or Microsoft Team hold, only the group mailbox and group site are placed on hold; the mailboxes and OneDrive for Business sites of group members aren't placed on hold unless you explicitly add them as custodians or place their data sources hold. Therefore, if you the need to place an Microsoft 365 Group or Microsoft Team on hold for a specific custodian, consider mapping the group site and group mailbox to the custodian (See Managing Custodians in Advanced eDiscovery). If the Microsoft 365 Group or Microsoft Team is not attributable to a single custodian, consider adding the source to a non-custodial hold.

- To get a list of the members of a Microsoft 365 Group or Microsoft Team, you can view the properties on the Home > Groups page in the Microsoft 365 admin center. Alternatively, you can run the following command in Exchange Online PowerShell:

```
Get-UnifiedGroupLinks <group or team name> -LinkType Members | FL DisplayName,PrimarySmtpAddress
```

> **NOTE**
> To run the **Get-UnifiedGroupLinks** cmdlet, you have to be assigned the View-Only Recipients role in Exchange Online or be a member of a role group that's assigned the View-Only Recipients role.

- Channel conversations that are part of a Microsoft Teams channel are stored in the mailbox that's associated with the Team. Similarly, files that team members share in a channel are stored on the team's SharePoint site. Therefore, you have to place the Microsoft Team mailbox and SharePoint site on hold to retain conversations and files in a channel.

- Alternatively, conversations that are part of the Chat list in Microsoft Teams are stored in the mailbox of the user's who participate in the chat. Files that a user shares in Chat conversations are stored in the OneDrive for Business site of the user who shares the file. Therefore, you have to place the individual user mailboxes and OneDrive for Business sites on hold to retain conversations and files in the Chat list.

- Every Microsoft Team or team channel contains a Wiki for note-taking and collaboration. The Wiki content is automatically saved to a file with a .mht format. This file is stored in the Teams Wiki Data document library on the team's SharePoint site. You can place the content in the Wiki on hold by placing the team's SharePoint site on hold.

NOTE

The capability to retain Wiki content for a Microsoft Team or team channel (when you place the team's SharePoint site on hold) was released on June 22, 2017. If a team site is on hold, the Wiki content will be retained starting on that date. However, if a team site is on hold and the Wiki content was deleted before June 22, 2017, the Wiki content was not retained.

Work with processing errors in Advanced eDiscovery

2/18/2021 • 2 minutes to read • Edit Online

Processing is the process of file identification, expansion of embedded documents and attachments, text extraction, and Optical Character Recognition (OCR)) of image files and the subsequent indexing of that content.

When you add custodians and non-custodian data sources to a case on the **Sources** tab, all partially indexed items from Microsoft 365 are processed to make them fully searchable. Likewise, when content is added to a review set from both Microsoft 365 and non-Microsoft 365 data sources, this content is also processed.

The **Processing** tab in Advanced eDiscovery provides insight into the status of advanced indexing for different processing scenarios.

For more information, see the following articles:

- Advanced indexing of custodian data
- Error remediation when processing data
- Single item error remediation

Advanced indexing of custodian data

2/18/2021 • 2 minutes to read • Edit Online

When a custodian is added to an Advanced eDiscovery case, any content that was deemed as partially indexed is reprocessed to make it fully searchable. This process is called *Advanced indexing*. Content can be partially indexed for a number of reasons including the existence of images, unsupported file types or when indexing file size limits are encountered.

To learn more about processing support and partially indexed items, see:

- Supported file types in Advanced eDiscovery
- Partially indexed items in Content Search in Office 365
- File formats indexed by Exchange Search
- Default crawled file name extensions and parsed file types in SharePoint Server

Viewing Advanced indexing results

After the Advanced indexing process is completed, you can get an understanding of the effectiveness of reprocessing. In the Advanced indexing results view on the **Processing** tab for a case, the graph lists the number of items added to the *hybrid index*. The hybrid index is where Advanced eDiscovery stores the reprocessed content.

This view also includes the number of items that require remediation and another graph of errors by file type. For more information, see:

- Error remediation when processing data
- Single item error remediation

Updating the Advanced index for custodians

When a custodian is added to an Advanced eDiscovery case, all partially indexed items are reprocessed. However, as time passes, more partially indexed items may be added to a user's mailbox or OneDrive account. If necessary, you can update the index for specific custodian. For more information, see Manage custodians in an Advanced eDiscovery case. You can also update the index for all custodians in a case by clicking the **Update index** on the **Processing** tab.

> **NOTE**
>
> Updating custodian indexes is a long running process. It's recommended that you don't update indexes more than once a day in a case.

Error remediation when processing data

2/18/2021 • 5 minutes to read • Edit Online

Error remediation allows eDiscovery administrators the ability to rectify data issues that prevent Advanced eDiscovery from properly processing the content. For example, files that are password protected can't be processed since the files are locked or encrypted. Using error remediation, eDiscovery administrators can download files with such errors, remove the password protection, and then upload the remediated files.

Use the following workflow to remediate files with errors in Advanced eDiscovery cases.

Create an error remediation session to remediate files with processing errors

> **NOTE**
> If the the error remediation wizard is closed at any time during the following procedure, you can return to the error remediation session from the **Processing** tab by selecting **Remediations** in the **View** drop-down menu.

1. On the **Processing** tab in the Advanced eDiscovery case, select **Errors** in the **View** drop-down menu and then select a review set or the entire case in the **Scope** drop-down menu. This section displays all errors from the case or error from a specific review set.

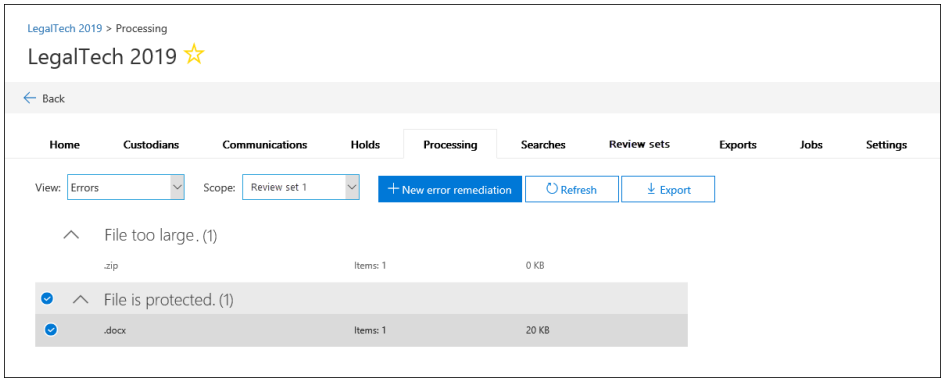

2. Select the errors you want to remediate by clicking the radio button next to either the error type or file type. In the following example, we're remediating a password protected file.

3. Click **New error remediation**.

   The error remediation workflow starts with a preparation stage where the files with errors are copied to a Microsoft-provided Azure Storage location so that you can download them to your local computer to remediate.

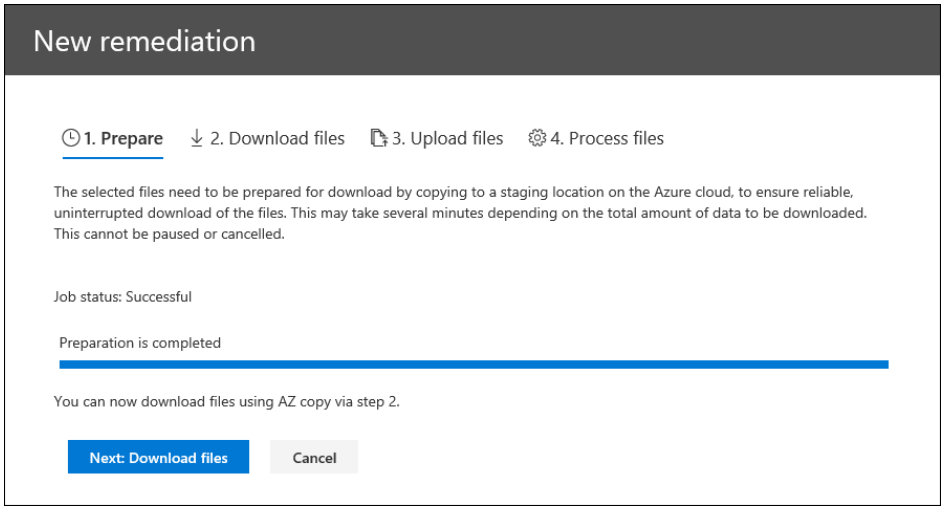

4. After the preparation is complete, click **Next: Download files** to proceed with download.

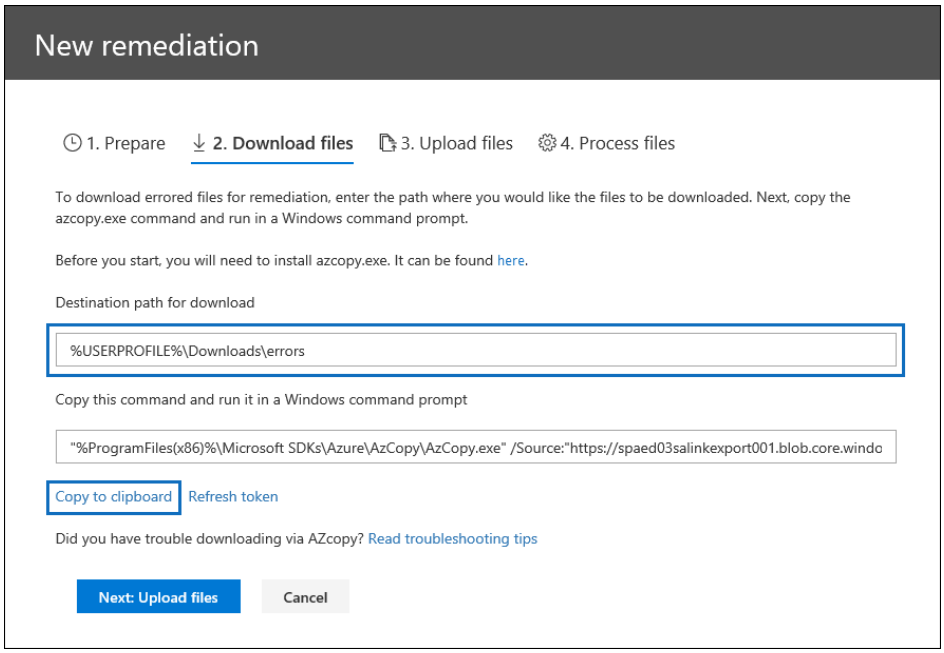

5. To download files, specify the **Destination path for download**. This is a path to the parent folder on your local computer where the file will be downloaded. The default path, %USERPROFILE%\Downloads\errors, points to the logged-in user's downloads folder. You can change this path if desired. If you do change it, we recommend that you use a local file path for the best performance. Don't use a remote network path. For example, you could use the path **C:\Remediation**.

   The path to the parent folder is automatically added to AzCopy command (as the value of the **/Dest** parameter).

6. Copy the predefined command by clicking **Copy to clipboard**. Open a Windows Command Prompt, paste the AzCopy command, and then press **Enter**.

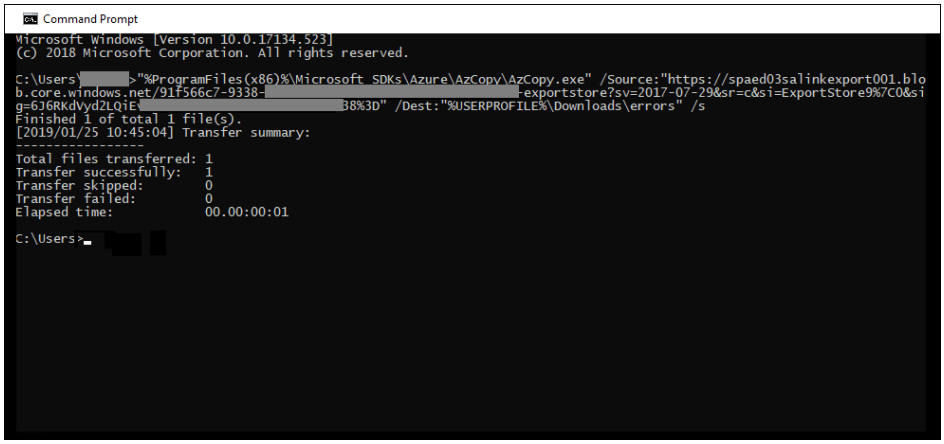

> **NOTE**
>
> You must use AzCopy v8.1 to successfully use the command that's provided on the **Download files** page. You also must use AzCopy v8.1 to upload the files in step 10. To install this version of AzCopy, see Transfer data with the AzCopy v8.1 on Windows. If the supplied AzCopy command fails, please see Troubleshoot AzCopy in Advanced eDiscovery.

The files that you selected are downloaded to the location that you specified in step 5. In the parent folder (for example, **C:\Remediation**), the following subfolder structure is automatically created:

```
<Parent folder>\Subfolder 1\Subfolder 2\<file>
```

- *Subfolder 1* is named with the ID for the case or the review set, depending on the scope that you selected in step 1.
- *Subfolder 2* is named with the file ID of the downloaded file
- The downloaded file is located in *Subfolder 2* and is also named with the file ID.

Here's an example of the folder path and error file name that's created when items are downloaded to the **C:\Remediation** parent folder:

```
C:\Remediation\232f8b7e-089c-4781-88c6-
210da0615d32\d1459499146268a096ea20202cd029857d64087706e6d6ca2a224970ae3b8938\d1459499146268a096ea20202cd029857d64087706e6d6ca2a2249
```

If multiple files are downloaded, each one is downloaded to a subfolder that's named with the file ID.

> **IMPORTANT**
>
> When you upload files in step 9 and step 10, the remediated files must have that same filename and be located in the same subfolder structure. The subfolder and file names are used to associated the remediated file with the original error file. If the folder structure or file names are changed, you'll receive the following error: `Cannot apply Error Remediation to the current Workingset`. To prevent any issues, we recommend that keep the remediated files in the same parent folder and subfolder structure.

7. After downloading the files, you can remediate them with an appropriate tool. For password-protected files, there are several password cracking tools you can use. If you know the passwords for the files, you can open them and remove the password protection.

8. Return to Advanced eDiscovery and the error remediation wizard and then click **Next: Upload files**. This moves to the next page where you can now upload the files.

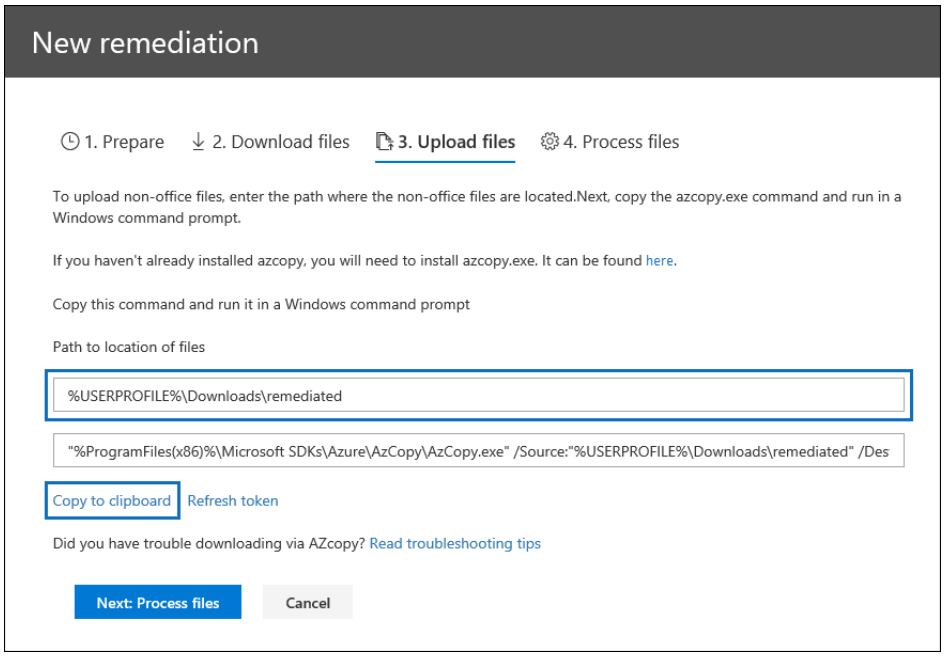

9. Specify the parent folder where the remediated files are located in the **Path to location of files** text box. Again, the parent folder must have the same subfolder structure that was created when you downloaded the files.

   The path to the parent folder is automatically added to AzCopy command (as the value of the **/Source** parameter).

10. Copy the predefined command by clicking **Copy to clipboard**. Open a Windows Command Prompt, paste the AzCopy command, and then press **Enter**. upload the files.

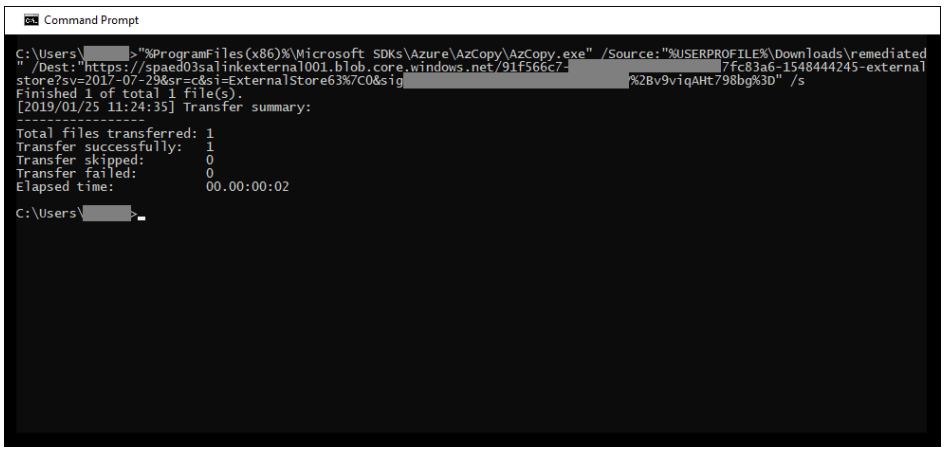

11. After you run the AzCopy command, click **Next: Process files**.

    When processing is complete, you can go to review set and view the remediated files.

Remediating errors in container files

In situations when the contents of a container file (such as a .zip file) can't be extracted by Advanced eDiscovery, the containers can be downloaded and the contents expanded into the same folder in which the original container resides. The expanded files will be attributed to the parent container as if it was originally expanded by Advanced eDiscovery. The process works as described as above except for uploading a single file as the replacement file. When you upload remediated files, don't include the original container file.

Remediating errors by uploading the extracted text

Sometimes it's not possible to remediate a file to native format that Advanced eDiscovery can interpret. But you can replace the original file with a text file that contains the original text of the native file (in a process called *text overlay*). To do this, follow the steps described in this article but instead of remediating the original file in the native format, you would create a text file that contains the extracted text from the original file, and then upload the text file using the original filename appended with a .txt suffix. For example, you download a file during error remediation with the filename 335850cc-6602-4af0-acfa-1d14d9128ca2.abc. You open the file in the native application, copy the text, and then paste it into a new file named 335850cc-6602-4af0-acfa-1d14d9128ca2.abc.txt. When you do this, be sure to remove the original file in the native format from the remediated file location on your local computer before uploading the remediated text file to Advanced eDiscovery.

What happens when files are remediated

When remediated files are uploaded, the original metadata is preserved except for the following fields:

- ExtractedTextSize
- HasText
- IsErrorRemediate
- LoadId
- ProcessingErrorMessage
- ProcessingStatus
- Text
- WordCount
- WorkingsetId

For a definition of all metadata fields in Advanced eDiscovery, see Document metadata fields.

Single item error remediation in Advanced eDiscovery

2/18/2021 • 2 minutes to read • Edit Online

Error remediation gives Advanced eDiscovery users the ability to rectify data issues that prevent Advanced eDiscovery from properly processing the content. For example, files that are password protected can't be processed because those files are locked or encrypted. Previously, you could only remediate errors in bulk by using this workflow. But sometimes, it doesn't make sense to remediate errors in multiple files when you're unsure if any of those files are responsive to the case you're investigating. It also might not make sense to remediate errors before you've had a chance to review the file metadata (such as file location or who had access) to help you make up-front decisions about responsiveness. A new feature called *single item error remediation* gives eDiscovery managers the ability to view the metadata of files with a processing error and if necessary remediate the error directly in the review set. The article discusses how to identify, ignore, and remediate files with processing errors in a review set.

Identify documents with errors

Documents with processing errors in a review set are now identified (with a banner). You can remediate or ignore the error. The following screenshot shows the processing error banner for a Word document in a review set that is password-protected. Also notice that you can view the file metadata of documents with processing errors.

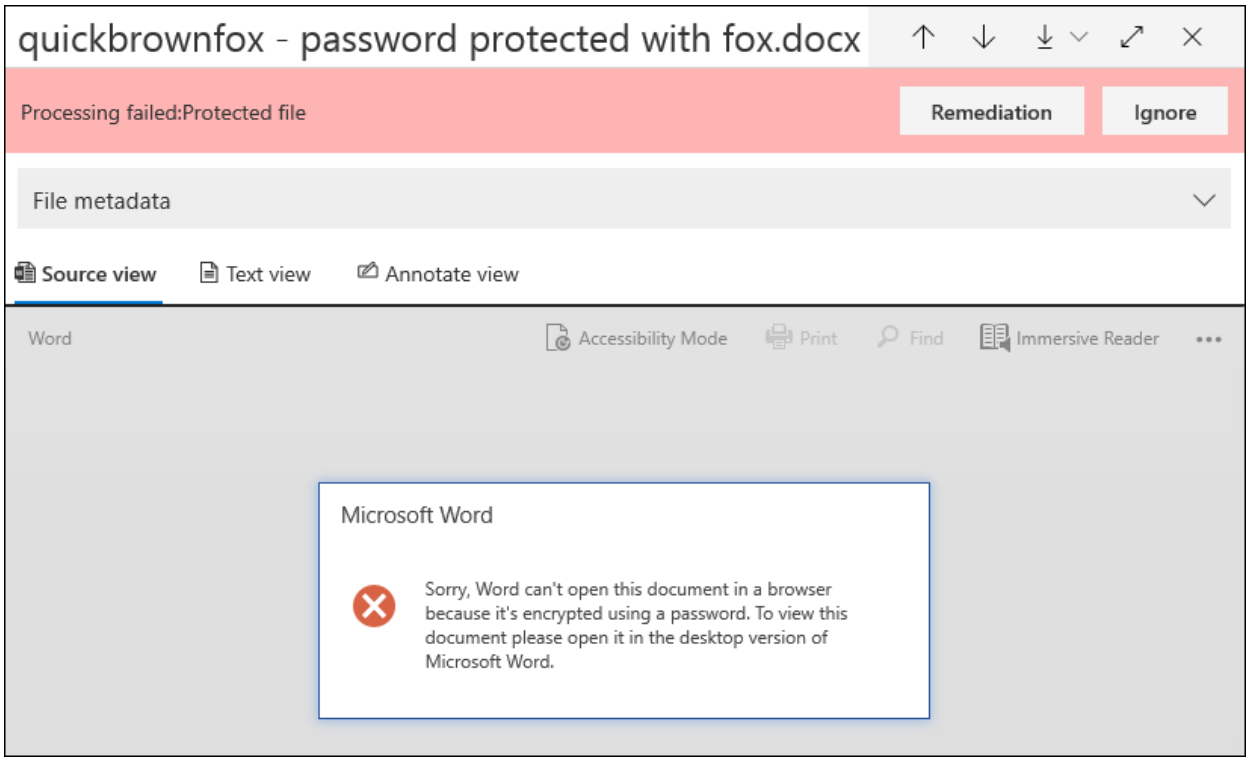

You can also search for documents that have processing errors by using the **Processing status** condition when querying the documents in a review set.

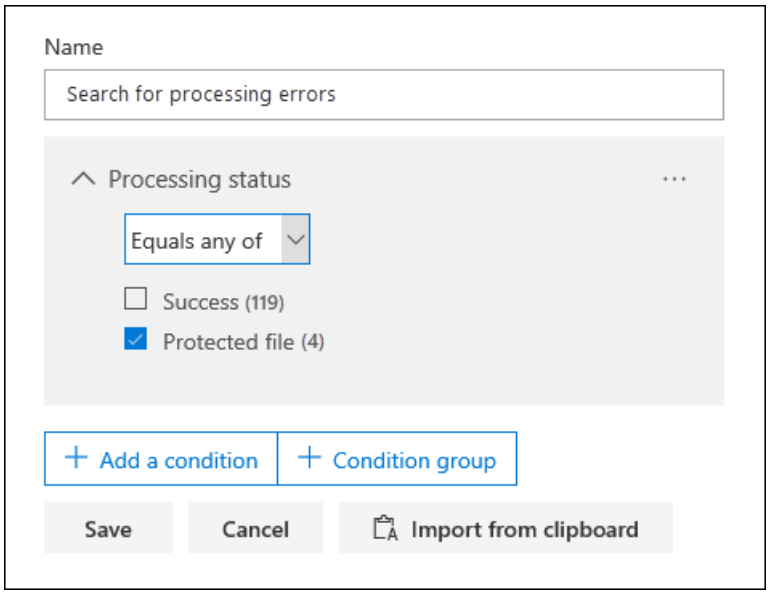

Ignore errors

You can ignore a processing error by clicking **Ignore** in the processing error banner. When you ignore an error, the document is removed from the bulk error remediation workflow. After an error is ignored, the document banner changes color and indicates that the processing error was ignored. At any time, you can revert the decision to ignore the error by clicking **Revert**.

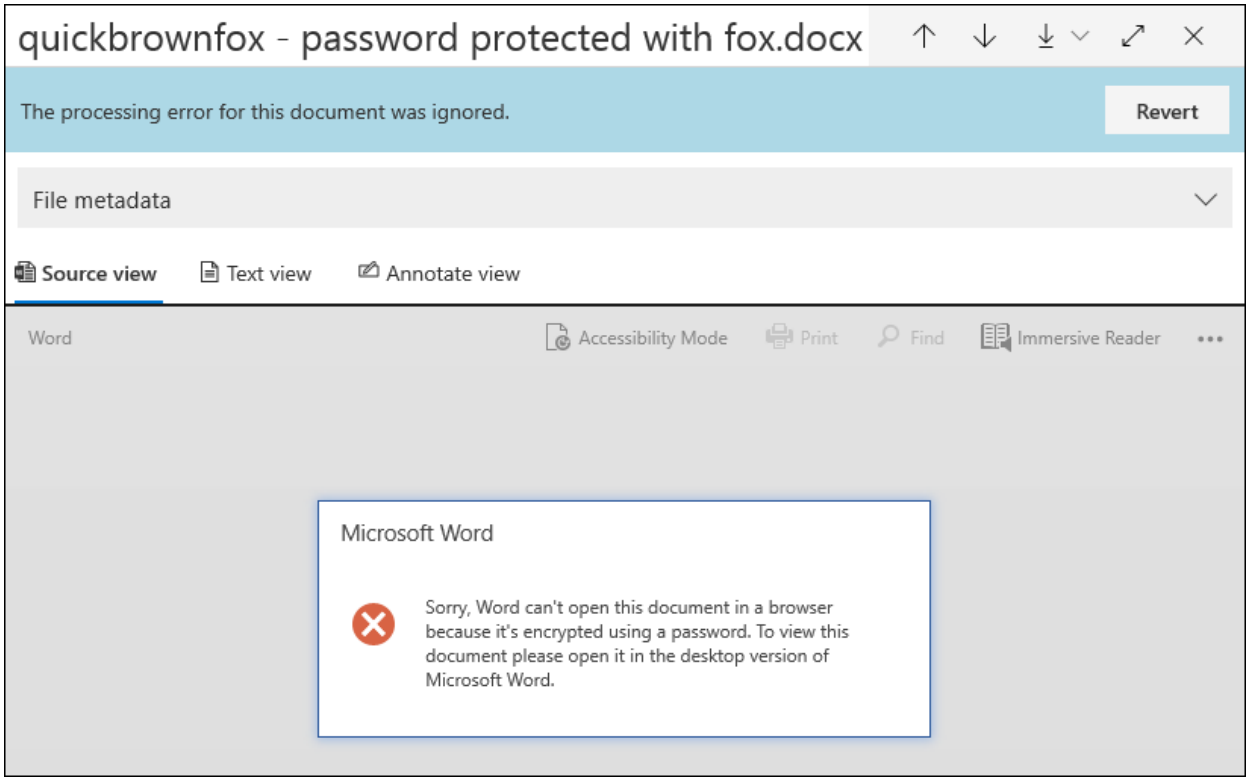

You can also search for all documents that had a processing error that was ignored by using the *Ignored processing errors* condition when querying documents in a review set.

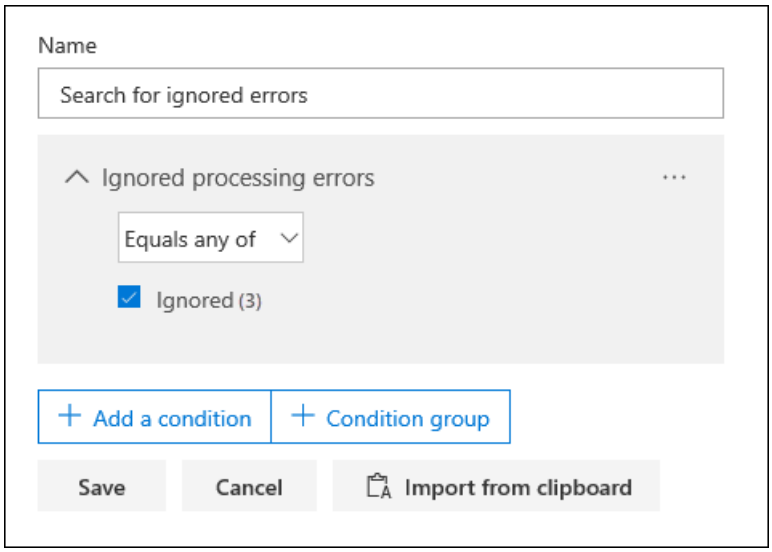

Remediate a document with errors

Sometimes you may be required to remediate a processing error in documents (by removing a password, decrypting an encrypted file, or recovering a corrupted document) and then add the remediated document to the review set. This allows you to review and export the error document together with the other documents in the review set.

To remediate a single document, follow these steps:

1. Click **Download** > **Download original** to download a copy of the file to a local computer.

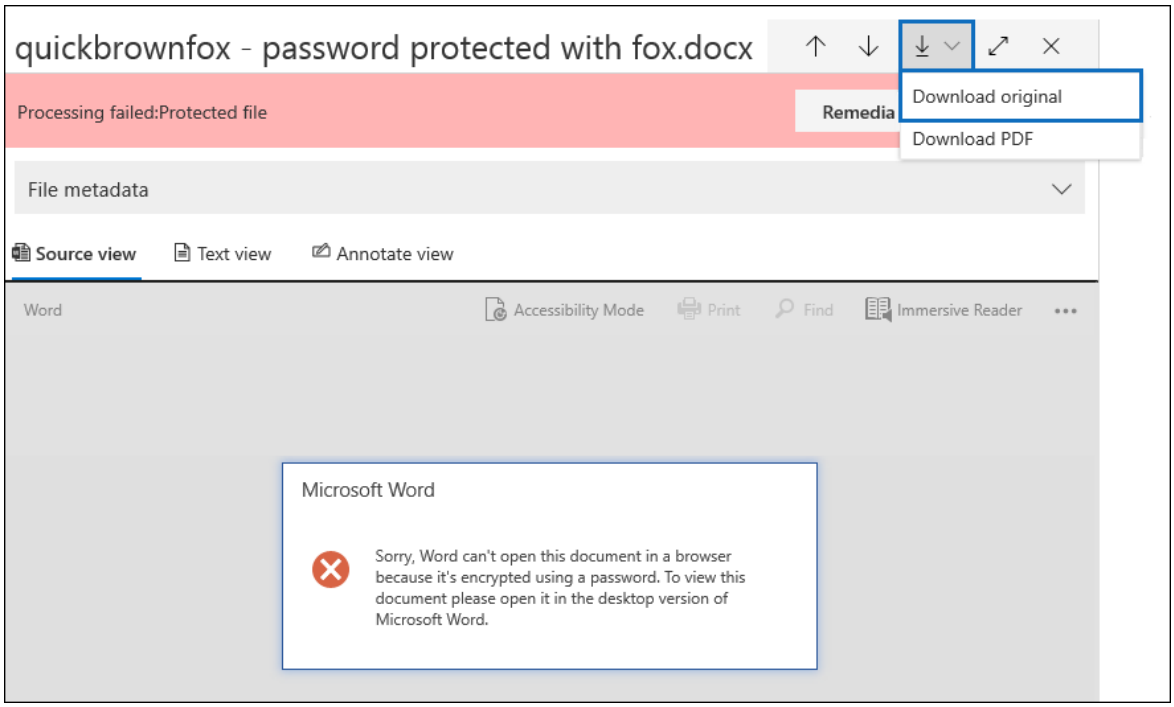

2. Remediate the error in the file offline. For encrypted files, that would require decryption software, to remove password protection, either provide the password and save the file or use a password cracker. After you remediate the file, go to the next step.

3. In the review set, select the file with the processing error that you remediated, and then click **Remediation**.

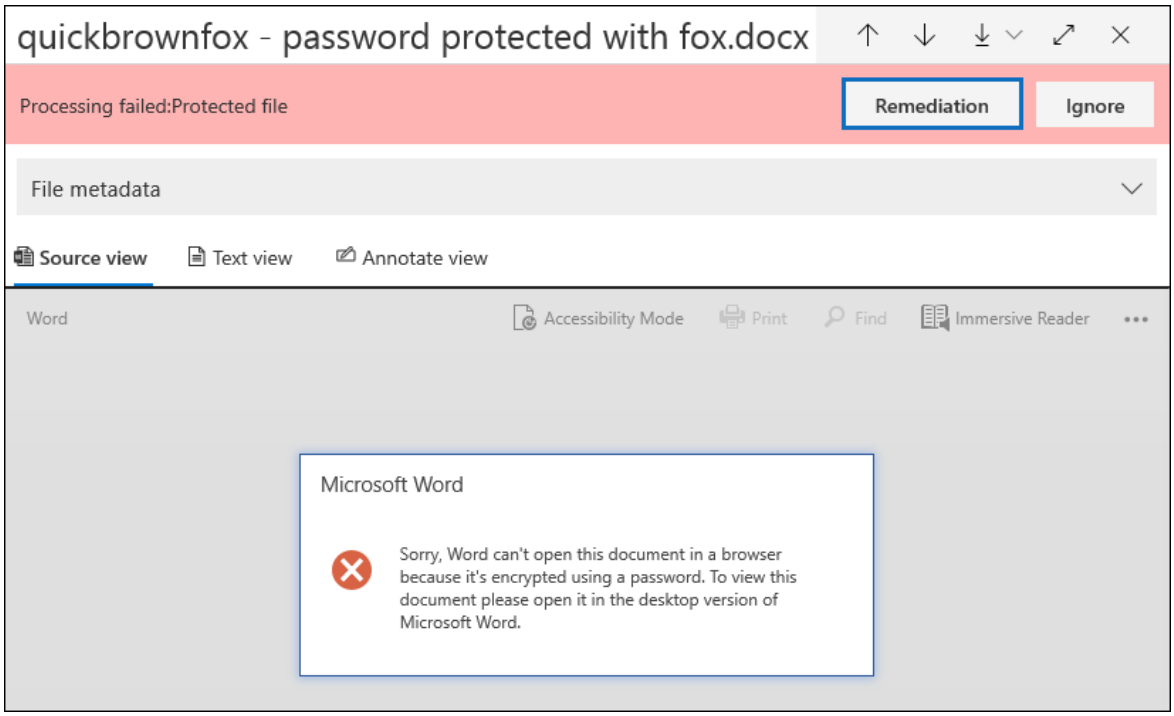

4. Click **Browse**, go to the location of the remediated file on your local computer, and then select the file.

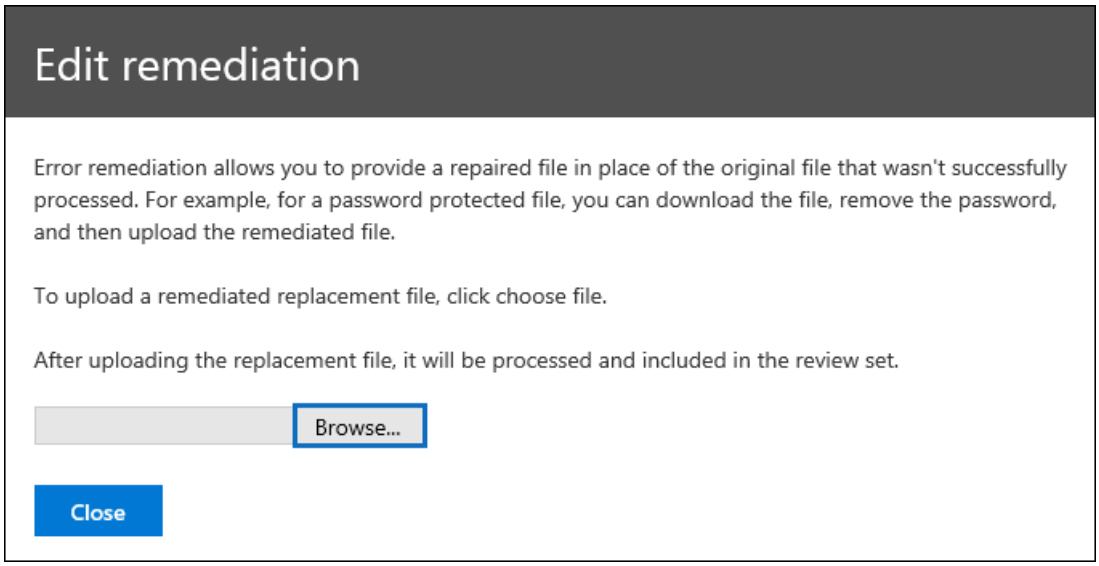

After selecting the remediated file, it is automatically uploaded to the review set. You can track the processing status of the file.

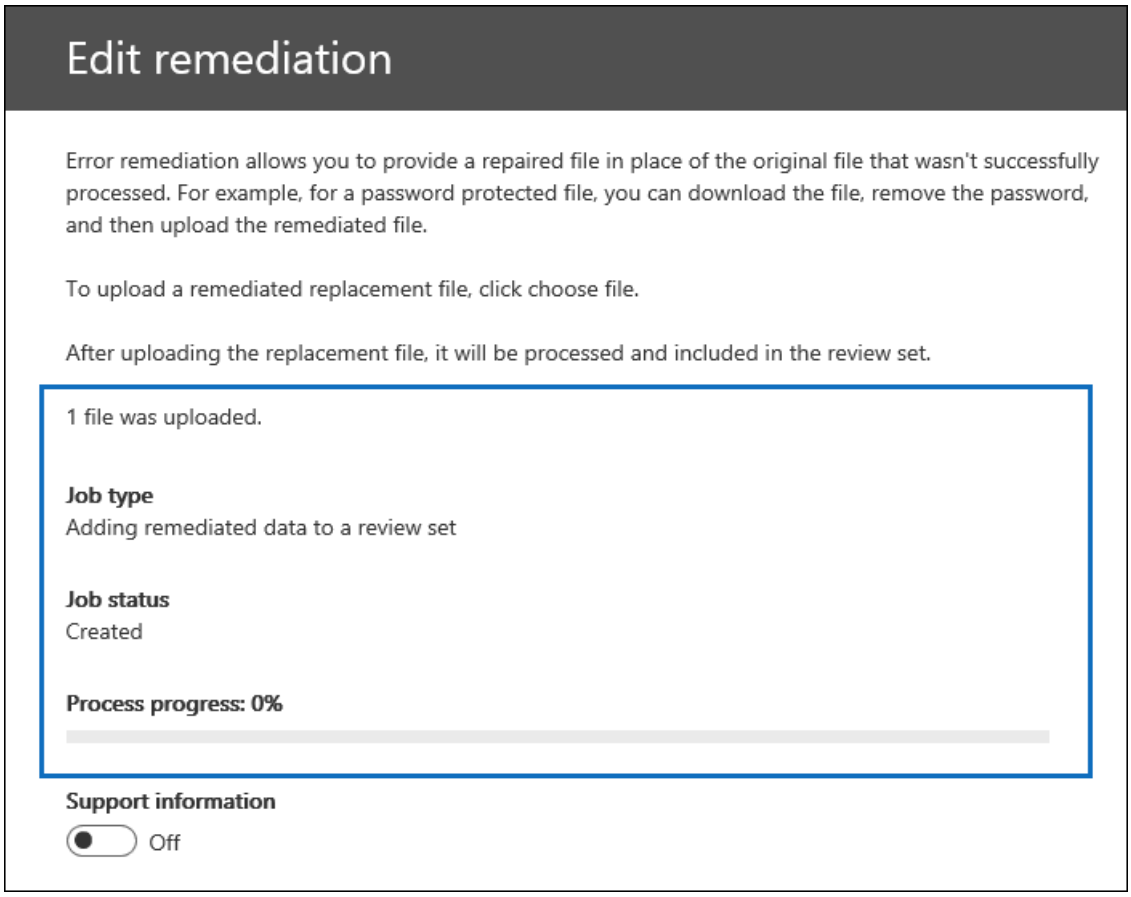

After processing is completed, you can view the remediated document.

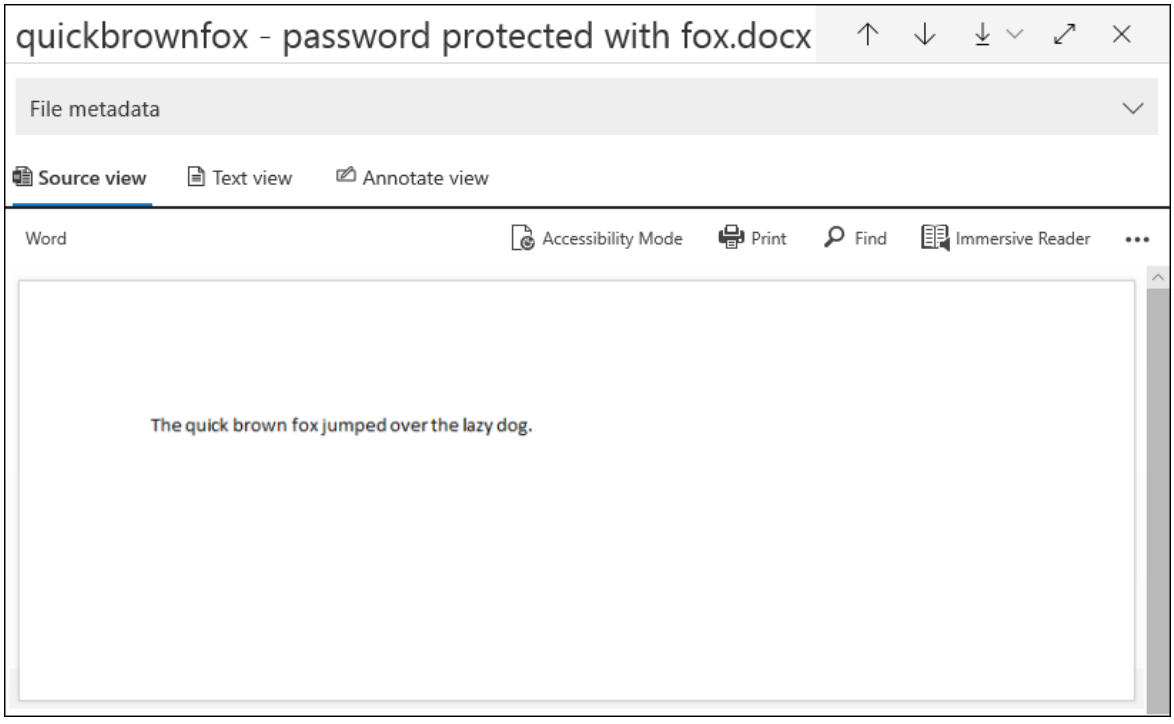

For more information about what happens when a document is remediated, see What happens when files are remediated.

Search for remediated documents

You can search for all documents in a review set that were remediated by using the **Keywords** condition and specifying the following property:value pair: **IsFromErrorRemediation:true**. This property is also available in the export load file when you export documents from a review set.

Collect data for a case in Advanced eDiscovery

2/18/2021 • 2 minutes to read • Edit Online

Once you've identified custodians and data sources that are of interest for your case, it's time to identify the set of documents to delve into. You can use the Search tool in Advanced eDiscovery to identify relevant documents from custodial and non-custodial locations in Microsoft 365.

After you run a search, you can view statistics on the retrieved items, such as the locations that had the most items that matched the search query. You can also preview a subset of the results. When you've identified the set of documents you want to further examine, you can add the search results to a review set to collect and process.

Create a search

Selecting **New search** on the **Searches** tab will start a wizard that guides you through creating a search. For detailed information on how to create a search, see Create a search to collect data.

After a search is created, a flyout page with details is displayed. The **Statistics** and **Preview** buttons are initially unavailable because the search hasn't completed yet. You can keep track of the progress of the search on the **Searches** tab.

View search results and statistics

There are two components of a content search: Statistics (Estimates) and Preview. As each of these components complete, you'll see the status displayed in the corresponding columns on the **Searches** tab change from **Submitted** to **In progress** to **Completed**.

Once the search estimate is completed, select the search to display the flyout page, which will display some high-level statistics about the results of the search. At this point, the **Statistics** button will be active. You can select it to see search statistics, such as:

- Summary
- Top locations
- Queries

For more information about search statistics, see Search statistics.

Once preview is completed, the **Preview** button will be active. Select it to preview a sampled subset of the results.

Add search results to a review set

When you're ready to collect and process the entire results of a search, you can do so by adding it to a review set. For details, see Add data to a review set.

Add non-Microsoft 365 data to a review set

As part of the collection process for a case, you can also add non-Office 365 data to a review set and review and analyze together with the Office 365 data that you collected by using the search tool. When you add non-Office 365, you have to associate it with a specific custodian in the case. For more information, see Load non-Microsoft 365 data into a review set.

Create a search

5/5/2020 • 2 minutes to read • Edit Online

On the **Searches** tab in your case, you can create a new search by clicking **New search** and following the wizard.

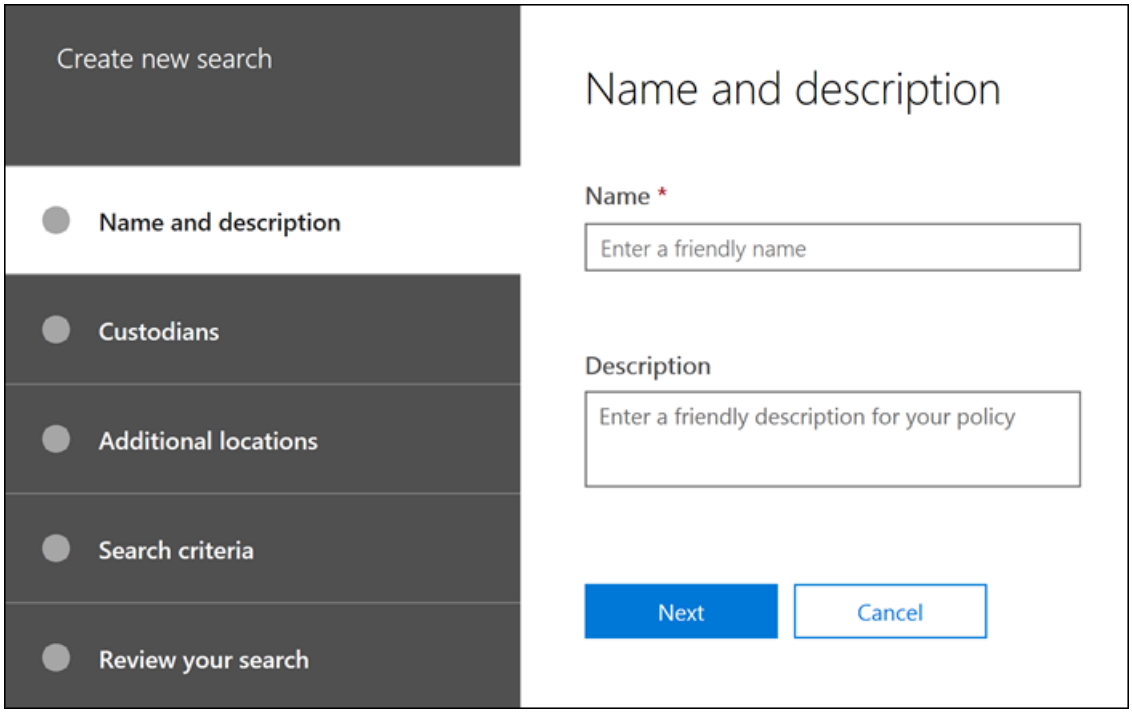

Name the search and give it a description

Each search with a case should have a unique name. You can optionally provide a description for your search.

Choose the custodians and custodial locations to search

Choose custodian content locations to search by specifying that custodians you have added to the case. By selecting a custodian, you will run the search against all data sources mapped to the custodian. You also have the option to narrow the search to selected data sources for each custodian. For more information about how to add custodians and manage their data sources, see Work with custodians.

Choose non-custodial locations

In some cases, you may want to search data sources that are not associated with a custodian. In this case, you can specify the locations you want to search, or choose to search all content locations for a specific Microsoft service (such as searching all Exchange mailboxes or all SharePoint sites and OneDrive accounts).

Define the search query and conditions

You can define the keywords query and any conditions for the search by using the pre-built condition cards or using Keyword Query Language (KQL). For more information, see Build search queries.

Build search collection queries in Advanced eDiscovery

2/18/2021 • 2 minutes to read • Edit Online

When building search queries to collect data in an Advanced eDiscovery case, you can use keywords to find specific content and conditions to narrow the scope of the search to return items that are most relevant to your legal investigation.

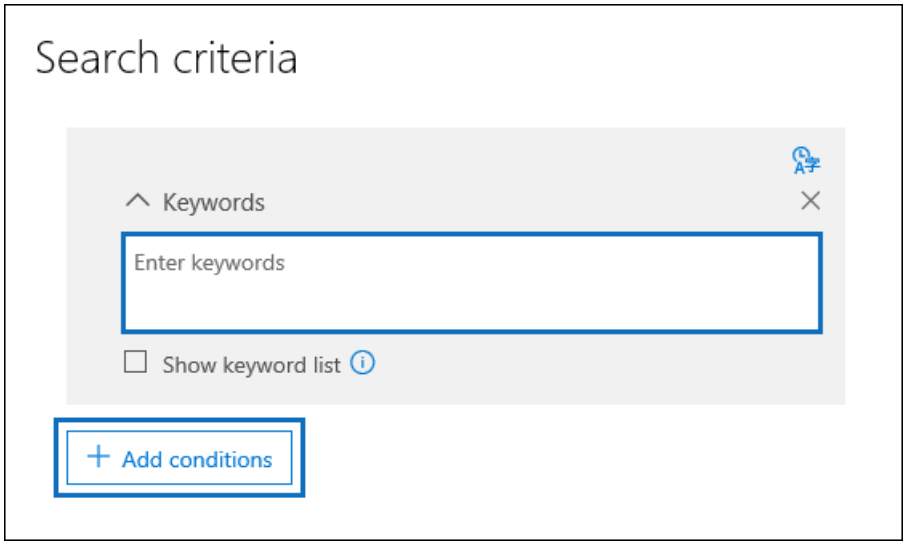

Keyword searches

Type a keyword query in the **Keywords** box in the search query. You can specify keywords, email message properties, such as sent and received dates, or document properties, such as file names or the date that a document was last changed. You can use more complex queries that use a Boolean operator, such as **AND**, **OR**, **NOT**, and **NEAR**. You can also search for sensitive information (such as social security numbers) in documents in SharePoint and OneDrive (not in email messages), or search for documents that have been shared externally. If you leave the **Keywords** box empty, all content located in the specified content locations is in the search results.

Keyword list

Alternatively, you can select the **Show keyword list** check box and the type a keyword or keyword phrase in each row. The keywords in each row are connected by a logical operator (which is represented as *c:s* in the search query syntax) that is similar in functionality to the **OR** operator in the search query that's created. This means items that contain any keyword in any row are in the search results. You can add up to 180 rows in the keyword list in Advanced eDiscovery search queries.

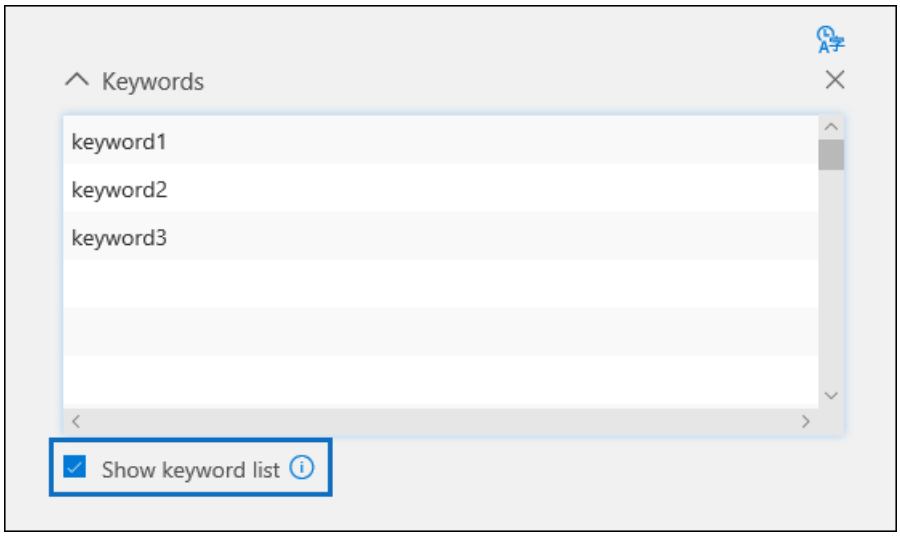

Why use the keyword list? You can get statistics that show how many items match each keyword in the keyword list. This can help you quickly identify the keywords that are the most (and least) effective. You can also use a keyword phrase (surrounded by parentheses) in a row in the keywords list. For more information about search statistics, see Search statistics.

Conditions

You can add search conditions to narrow the scope of a search and return a more refined set of results. Each condition adds a clause to the search query that is created and run when you start the search. A condition is logically connected to the keyword query specified in the keyword box by a logical operator (which is represented as *c:c* in the search query syntax) that is similar in functionality to the **AND** operator. That means items have to satisfy both the keyword query and one or more conditions to be included in the search results. This is how conditions help to narrow your results. For a list and description of conditions that you can use in a search query, see the "Search conditions" section in Keyword queries and search conditions.

Keyword queries and search conditions for Content Search and eDiscovery

2/18/2021 • 34 minutes to read • Edit Online

This topic describes the email and document properties that you can search for in email items in Exchange Online and documents stored on SharePoint and OneDrive for Business sites by using the Content Search feature in the Microsoft 365 compliance center. You can also use the ***-ComplianceSearch** cmdlets in Security & Compliance Center PowerShell to search for these properties. The topic also describes:

- Using Boolean search operators, search conditions, and other search query techniques to refine your search results.
- Searching for sensitive data types and custom sensitive data types in SharePoint and OneDrive for Business.
- Searching for site content that's shared with users outside of your organization

For step-by-step instructions on how to create a Content Search, see Content Search.

> **NOTE**
> Content Search in the Microsoft 365 compliance center and the corresponding ***-ComplianceSearch** cmdlets in Security & Compliance Center PowerShell use the Keyword Query Language (KQL). For more detailed information, see Keyword Query Language syntax reference.

Searchable email properties

The following table lists email message properties that can be searched by using the Content Search feature in the Microsoft 365 compliance center or by using the **New-ComplianceSearch** or the **Set-ComplianceSearch** cmdlet. The table includes an example of the *property:value* syntax for each property and a description of the search results returned by the examples. You can type these `property:value` pairs in the keywords box for a Content Search.

> **NOTE**
> When searching email properties, it's not possible to search for items in which the specified property is empty or blank. For example, using the *property:value* pair of **subject:""** to search for email messages with an empty subject line will return zero results. This also applies when searching site and contact properties.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLES	SEARCH RESULTS RETURNED BY THE EXAMPLES

PROPERTY	PROPERTY DESCRIPTION	EXAMPLES	SEARCH RESULTS RETURNED BY THE EXAMPLES

AttachmentNames	The names of files attached to an email message.	`attachmentnames:annualreport.ppt`
		`attachmentnames:annual*`
		`attachmentnames:.pptx`
			Messages that have an attached file named annualreport.ppt. In the second example, using the wildcard returns messages with the word "annual" in the file name of an attachment. The third example returns all attachments with the pptx file extension.

Bcc	The Bcc field of an email message.[1]	`bcc:pilarp@contoso.com`
		`bcc:pilarp`
		`bcc:"Pilar Pinilla"`
			All examples return messages with Pilar Pinilla included in the Bcc field.

Category	The categories to search. Categories can be defined by users by using Outlook or Outlook on the web (formerly known as Outlook Web App). The possible values are: blue green orange purple red yellow	`category:"Red Category"`	Messages that have been assigned the red category in the source mailboxes.
Cc	The Cc field of an email message.[1]	`cc:pilarp@contoso.com` `cc:"Pilar Pinilla"`	In both examples, messages with Pilar Pinilla specified in the Cc field.
Folderid	The folder ID (GUID) of a specific mailbox folder. If you use this property, be sure to search the mailbox that the specified folder is located in. Only the specified folder will be searched. Any subfolders in the folder won't be searched. To search subfolders, you need to use the Folderid property for the subfolder you want to search. For more information about searching for the Folderid property and using a script to obtain the folder IDs for a specific mailbox, see Use Content Search for targeted collections.	`folderid:4D6DD7F943C29041A65[illegible]A0000` `folderid:2370FB455F82FC44BE3[illegible]00001160000 AND participants:garthf@conto[illegible]`	The first example returns all items in the specified mailbox folder. The second example returns all items in the specified mailbox folder that were sent or received by garthf@contoso.com.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLES	SEARCH RESULTS RETURNED BY THE EXAMPLES
From	The sender of an email message.[1]	`from:pilarp@contoso.com` `from:contoso.com`	Messages sent by the specified user or sent from a specified domain.
HasAttachment	Indicates whether a message has an attachment. Use the values **true** or **false**.	`from:pilar@contoso.com AND hasattachment:true`	Messages sent by the specified user that have attachments.

Importance	The importance of an email message, which a sender can specify when sending a message. By default, messages are sent with normal importance, unless the sender sets the importance as **high** or **low**.	`importance:high`
		`importance:medium`
		`importance:low`
			Messages that are marked as high importance, medium importance, or low importance.

IsRead	Indicates whether messages have been read. Use the values **true** or **false**.	`isread:true` `isread:false`	The first example returns messages with the IsRead property set to **True**. The second example returns messages with the IsRead property set to **False**.
ItemClass	Use this property to search specific third-party data types that your organization imported to Office 365. Use the following syntax for this property: `itemclass:ipm.externaldata.<third-party data type>*`	`itemclass:ipm.externaldata.Facebook* AND subject:contoso` `itemclass:ipm.externaldata.Twitter* AND from:"Ann Beebe" AND "Northwind Traders"`	The first example returns Facebook items that contain the word "contoso" in the Subject property. The second example returns Twitter items that were posted by Ann Beebe and that contain the keyword phrase "Northwind Traders". For a complete list of values to use for third-party data types for the ItemClass property, see Use Content Search to search third-party data that was imported to Office 365.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLES	SEARCH RESULTS RETURNED BY THE EXAMPLES

Kind	The type of email message to search for. Possible values:
	contacts
	docs
	email
	externaldata
	faxes
	im
	journals
	meetings
	microsoftteams (returns items from chats, meetings, and calls in Microsoft Teams)
	notes
	posts
	rssfeeds
	tasks
	voicemail
		`kind:email`
		`kind:email OR kind:im OR kind:voicemail`
		`kind:externaldata`
			The first example returns email messages that meet the search criteria. The second example returns email messages, instant messaging conversations (including Skype for Business conversations and chats in Microsoft Teams), and voice messages that meet the search criteria. The third example returns items that were imported to mailboxes in Microsoft 365 from third-party data sources, such as Twitter, Facebook, and Cisco Jabber, that meet the search criteria. For more information, see Archiving third-party data in Office 365.

Participants	All the people fields in an email message. These fields are From, To, Cc, and Bcc.[1]	`participants:garthf@contoso.com` `participants:contoso.com`	Messages sent by or sent to garthf@contoso.com. The second example returns all messages sent by or sent to a user in the contoso.com domain.
Received	The date that an email message was received by a recipient.	`received:04/15/2016` `received>=01/01/2016 AND received<=03/31/2016`	Messages that were received on April 15, 2016. The second example returns all messages received between January 1, 2016 and March 31, 2016.
Recipients	All recipient fields in an email message. These fields are To, Cc, and Bcc.[1]	`recipients:garthf@contoso.com` `recipients:contoso.com`	Messages sent to garthf@contoso.com. The second example returns messages sent to any recipient in the contoso.com domain.
Sent	The date that an email message was sent by the sender.	`sent:07/01/2016` `sent>=06/01/2016 AND sent<=07/01/2016`	Messages that were sent on the specified date or sent within the specified date range.
Size	The size of an item, in bytes.	`size>26214400` `size:1..1048567`	Messages larger than 25?? MB. The second example returns messages from 1 through 1,048,567 bytes (1 MB) in size.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLES	SEARCH RESULTS RETURNED BY THE EXAMPLES
Subject	The text in the subject line of an email message. **Note:** When you use the Subject property in a query, the search returns all messages in which the subject line contains the text you're searching for. In other words, the query doesn't return only those messages that have an exact match. For example, if you search for `subject:"Quarterly Financials"`, your results will include messages with the subject "Quarterly Financials 2018".	`subject:"Quarterly Financials"` `subject:northwind`	Messages that contain the phrase "Quarterly Financials" anywhere in the text of the subject line. The second example returns all messages that contain the word northwind in the subject line.

To	The To field of an email message.[1]	`to:annb@contoso.com`
		`to:annb`
		`to:"Ann Beebe"`
			All examples return messages where Ann Beebe is specified in the To: line.

NOTE

[1] For the value of a recipient property, you can use email address (also called *user principal name* or UPN), display name, or alias to specify a user. For example, you can use annb@contoso.com, annb, or "Ann Beebe" to specify the user Ann Beebe.

Recipient expansion

When searching any of the recipient properties (From, To, Cc, Bcc, Participants, and Recipients), Microsoft 365 attempts to expand the identity of each user by looking them up in Azure Active Directory (Azure AD). If the user is found in Azure AD, the query is expanded to include the user's email address (or UPN), alias, display name, and LegacyExchangeDN. For example, a query such as `participants:ronnie@contoso.com` expands to `participants:ronnie@contoso.com OR participants:ronnie OR participants:"Ronald Nelson" OR participants:"<LegacyExchangeDN>"`.

To prevent recipient expansion, add a wild card character (asterisk) to the end of the email address and use a reduced domain name; for example, `participants:"ronnie@contoso*"` Be sure to surround the email address with double quotation marks.

However, be aware that preventing recipient expansion in the search query may result in relevant items not being returned in the search results. Email messages in Exchange can be saved with different text formats in the recipient fields. Recipient expansion is intended to help mitigate this fact by returning messages that may contain different text formats. So preventing recipient expansion may result in the search query not returning all items that may be relevant to your investigation.

NOTE

If you need to review or reduce the items returned by a search query due to recipient expansion, consider using Advanced eDiscovery. You can search for messages (taking advantage of recipient expansion), add them to a review set, and then use review set queries or filters to review or narrow the results. For more information, see Collect data for a case and Query the data in a review set.

Searchable site properties

The following table lists some of the SharePoint and OneDrive for Business properties that can be searched by using the Content Search feature in the Security & Compliance Center or by using the **New-ComplianceSearch** or the **Set-ComplianceSearch** cmdlet. The table includes an example of the *property:value* syntax for each property and a description of the search results returned by the examples.

For a complete list of SharePoint properties that can be searched, see Overview of crawled and managed properties in SharePoint. Properties marked with a **Yes** in the **Queryable** column can be searched.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLE	SEARCH RESULTS RETURNED BY THE EXAMPLES
Author	The author field from Office documents, which persists if a document is copied. For example, if a user creates a document and the emails it to someone else who then uploads it to SharePoint, the document will still retain the original author. Be sure to use the user's display name for this property.	`author:"Garth Fort"`	All documents that are authored by Garth Fort.
ContentType	The SharePoint content type of an item, such as Item, Document, or Video.	`contenttype:document`	All documents would be returned.
Created	The date that an item is created.	`created>=06/01/2016`	All items created on or after June 1, 2016.
CreatedBy	The person that created or uploaded an item. Be sure to use the user's display name for this property.	`createdby:"Garth Fort"`	All items created or uploaded by Garth Fort.
DetectedLanguage	The language of an item.	`detectedlanguage:english`	All items in English.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLE	SEARCH RESULTS RETURNED BY THE EXAMPLES
DocumentLink	The path (URL) of a specific folder on a SharePoint or OneDrive for Business site. If you use this property, be sure to search the site that the specified folder is located in. To return items located in subfolders of the folder that you specify for the documentlink property, you have to add /* to the URL of the specified folder; for example, `documentlink: "https://contoso.sharepoint.com/Shared Documents/*"` For more information about searching for the documentlink property and using a script to obtain the documentlink URLs for folders on a specific site, see Use Content Search for targeted collections.	`documentlink:"https://contoso-my.sharepoint.com/personal/garthf_contoso_com/Documents/Private` `documentlink:"https://contoso-my.sharepoint.com/personal/garthf_contoso_com/Documents/Shared with Everyone/*" AND filename:confidential`	The first example returns all items in the specified OneDrive for Business folder. The second example returns documents in the specified site folder (and all subfolders) that contain the word "confidential" in the file name.
FileExtension	The extension of a file; for example, docx, one, pptx, or xlsx.	`fileextension:xlsx`	All Excel files (Excel 2007 and later)
FileName	The name of a file.	`filename:"marketing plan"` `filename:estimate`	The first example returns files with the exact phrase "marketing plan" in the title. The second example returns files with the word "estimate" in the file name.
LastModifiedTime	The date that an item was last changed.	`lastmodifiedtime>=05/01/2016` `lastmodifiedtime>=05/10/2016 AND lastmodifiedtime<=06/1/2016`	The first example returns items that were changed on or after May 1, 2016. The second example returns items changed between May 1, 2016 and June 1, 2016.
ModifiedBy	The person who last changed an item. Be sure to use the user's display name for this property.	`modifiedby:"Garth Fort"`	All items that were last changed by Garth Fort.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLE	SEARCH RESULTS RETURNED BY THE EXAMPLES
Path	The path (URL) of a specific site in a SharePoint or OneDrive for Business site. To return items located in folders in the site that you specify for the path property, you have to add /* to the URL of the specified site; for example, `path: "https://contoso.sharepoint.com/Shared Documents/*"` **Note:** Using the `Path` property to search OneDrive locations won't return media files, such as .png, .tiff, or .wav files, in the search results. Use a different site property in your search query to search for media files in OneDrive folders.	`path:"https://contoso-my.sharepoint.com/personal/garthf_contoso_com/"` `path:"https://contoso-my.sharepoint.com/personal/garthf_contoso_com/*" AND filename:confidential`	The first example returns all items in the specified OneDrive for Business site. The second example returns documents in the specified site (and folders in the site) that contain the word "confidential" in the file name.
SharedWithUsersOWSUser	Documents that have been shared with the specified user and displayed on the **Shared with me** page in the user's OneDrive for Business site. These are documents that have been explicitly shared with the specified user by other people in your organization. When you export documents that match a search query that uses the SharedWithUsersOWSUser property, the documents are exported from the original content location of the person who shared the document with the specified user. For more information, see Searching for site content shared within your organization.	`sharedwithusersowsuser:garthf` `sharedwithusersowsuser:"garthf@contoso.com"`	Both examples return all internal documents that have been explicitly shared with Garth Fort and that appear on the **Shared with me** page in Garth Fort's OneDrive for Business account.
Site	The URL of a site or group of sites in your organization.	`site:"https://contoso-my.sharepoint.com"` `site:"https://contoso.sharepoint.com/sites/teamA"`	The first example returns items from the OneDrive for Business sites for all users in the organization. The second example returns items from all team sites.
Size	The size of an item, in bytes.	`size>=1` `size:1..10000`	The first example returns items larger than 1 byte. The second example returns items from 1 through 10,000 bytes in size.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLE	SEARCH RESULTS RETURNED BY THE EXAMPLES
Title	The title of the document. The Title property is metadata that's specified in Microsoft Office documents. It's different from the file name of the document.	`title:"communication plan"`	Any document that contains the phrase "communication plan" in the Title metadata property of an Office document.

Searchable contact properties

The following table lists the contact properties that are indexed and that you can search for using Content Search. These are the properties that are available for users to configure for the contacts (also called personal contacts) that are located in the personal address book of a user's mailbox. To search for contacts, you can select the mailboxes to search and then use one or more contact properties in the keyword query.

> **TIP**
> To search for values that contain spaces or special characters, use double quotation marks (" ") to contain the phrase; for example, `businessaddress:"123 Main Street"`.

PROPERTY	PROPERTY DESCRIPTION		
BusinessAddress	The address in the **Business Address** property. The property is also called the **Work** address on the contact properties page.		
BusinessPhone	The phone number in any of the **Business Phone** number properties.		
CompanyName	The name in the **Company** property.		
Department	The name in the **Department** property.		
DisplayName	The display name of the contact. This is the name in the **Full Name** property of the contact.		
EmailAddress	The address for any email address property for the contact. Users can add multiple email addresses for a contact. Using this property would return contacts that match any of the contact's email addresses.		

PROPERTY	PROPERTY DESCRIPTION		
FileAs	The **File as** property. This property is used to specify how the contact is listed in the user's contact list. For example, a contact could be listed as *FirstName,LastName* or *LastName,FirstName*.		
GivenName	The name in the **First Name** property.		
HomeAddress	The address in any of the **Home** address properties.		
HomePhone	The phone number in any of the **Home** phone number properties.		
IMAddress	The IM address property, which is typically an email address used for instant messaging.		
MiddleName	The name in the **Middle** name property.		
MobilePhone	The phone number in the **Mobile** phone number property.		
Nickname	The name in the **Nickname** property.		
OfficeLocation	The value in **Office** or **Office location** property.		
OtherAddress	The value for the **Other** address property.		
Surname	The name in the **Last** name property.		
Title	The title in the **Job title** property.		

Searchable sensitive data types

You can use eDiscovery search tools in the Microsoft 365 compliance center to search for sensitive data, such as credit card numbers or social security numbers, that is stored in documents on SharePoint and OneDrive for Business sites. You can do this by using the `SensitiveType` property and the name (or ID) of a sensitive information type in a keyword query. For example, the query `SensitiveType:"Credit Card Number"` returns documents that contain a credit card number. The query `SensitiveType:"U.S. Social Security Number (SSN)"` returns documents that contain a U.S. social security number.

To see a list of the sensitive information types that you can search for, go to **Data classifications** > **Sensitive info types** in the Microsoft 365 compliance center. Or you can use the **Get-DlpSensitiveInformationType** cmdlet in Security & Compliance Center PowerShell to display a list of sensitive information types.

For more information about creating queries using the `SensitiveType` property, see Form a query to find sensitive data stored on sites.

Limitations for searching sensitive data types

- To search for custom sensitive information types, you have to specify the ID of the sensitive information type in the `SensitiveType` property. Using the name of a custom sensitive information type (as shown in the example for built-in sensitive information types in the previous section) will return no results. Use the **Publisher** column on the **Sensitive info types** page in the compliance center (or the **Publisher** property in PowerShell) to differentiate between built-in and custom sensitive information types. Built-in sensitive data types have a value of `Microsoft Corporation` for the **Publisher** property.

  To display the name and ID for the custom sensitive data types in your organization, run the following command in Security & Compliance Center PowerShell:

  ```
  Get-DlpSensitiveInformationType | Where-Object {$_.Publisher -ne "Microsoft Corporation"} | FT
  Name,Id
  ```

  Then you can use the ID in the `SensitiveType` search property to return documents that contain the custom sensitive data type; for example, `SensitiveType:7e13277e-6b04-3b68-94ed-1aeb9d47de37`

- You can't use sensitive information types and the `SensitiveType` search property to search for sensitive data at-rest in Exchange Online mailboxes. However, you can use data loss prevention (DLP) policies to protect sensitive email data in transit. For more information, see Overview of data loss prevention policies and Search for and find personal data.

Search operators

Boolean search operators, such as **AND**, **OR**, and **NOT**, help you define more-precise searches by including or excluding specific words in the search query. Other techniques, such as using property operators (such as `>=` or `..`), quotation marks, parentheses, and wildcards, help you refine a search query. The following table lists the operators that you can use to narrow or broaden search results.

OPERATOR	USAGE	DESCRIPTION	
AND	keyword1 AND keyword2	Returns items that include all of the specified keywords or `property:value` expressions. For example, `from:"Ann Beebe" AND subject:northwind` would return all messages sent by Ann Beebe that contained the word northwind in the subject line. [2]	

OPERATOR	USAGE	DESCRIPTION	
+	keyword1 + keyword2 + keyword3	Returns items that contain *either* `keyword2` or `keyword3` *and* that also contain `keyword1`. Therefore, this example is equivalent to the query `(keyword2 OR keyword3) AND keyword1`. The query `keyword1 + keyword2` (with a space after the **+** symbol) isn't the same as using the **AND** operator. This query would be equivalent to `"keyword1 + keyword2"` and return items with the exact phase `"keyword1 + keyword2"`.	
OR	keyword1 OR keyword2	Returns items that include one or more of the specified keywords or `property:value` expressions. [2]	
NOT	keyword1 NOT keyword2 NOT from:"Ann Beebe" NOT kind:im	Excludes items specified by a keyword or a `property:value` expression. In the second example excludes messages sent by Ann Beebe. The third example excludes any instant messaging conversations, such as Skype for Business conversations that are saved to the Conversation History mailbox folder. [2]	
-	keyword1 -keyword2	The same as the **NOT** operator. So this query returns items that contain `keyword1` and would exclude items that contain `keyword2`.	
NEAR	keyword1 NEAR(n) keyword2	Returns items with words that are near each other, where n equals the number of words apart. For example, `best NEAR(5) worst` returns any item where the word "worst" is within five words of "best". If no number is specified, the default distance is eight words. [2]	

OPERATOR	USAGE	DESCRIPTION	
:	property:value	The colon (:) in the `property:value` syntax specifies that the value of the property being searched for contains the specified value. For example, `recipients:garthf@contoso.com` returns any message sent to garthf@contoso.com.	
=	property=value	The same as the : operator.	
<	property<value	Denotes that the property being searched is less than the specified value. [1]	
>	property>value	Denotes that the property being searched is greater than the specified value.[1]	
<=	property<=value	Denotes that the property being searched is less than or equal to a specific value.[1]	
>=	property>=value	Denotes that the property being searched is greater than or equal to a specific value.[1]	
..	property:value1..value2	Denotes that the property being searched is greater than or equal to value1 and less than or equal to value2.[1]	
" "	"fair value" subject:"Quarterly Financials"	Use double quotation marks (" ") to search for an exact phrase or term in keyword and `property:value` search queries.	

OPERATOR	USAGE	DESCRIPTION	
*	cat* subject:set*	Prefix wildcard searches (where the asterisk is placed at the end of a word) match for zero or more characters in keywords or `property:value` queries. For example, `title:set*` returns documents that contain the word set, setup, and setting (and other words that start with "set") in the document title. **Note:** You can use only prefix wildcard searches; for example, **cat*** or **set***. Suffix searches (***cat**), infix searches (**c*t**), and substring searches (***cat***) are not supported.	
( )	(fair OR free) AND (from:contoso.com) (IPO OR initial) AND (stock OR shares) (quarterly financials)	Parentheses group together Boolean phrases, `property:value` items, and keywords. For example, `(quarterly financials)` returns items that contain the words quarterly and financials.	

NOTE

[1] Use this operator for properties that have date or numeric values.

[2] Boolean search operators must be uppercase; for example, **AND**. If you use a lowercase operator, such as **and**, it will be treated as a keyword in the search query.

Search conditions

You can add conditions to a search query to narrow a search and return a more refined set of results. Each condition adds a clause to the KQL search query that is created and run when you start the search.

Conditions for common properties

Conditions for mail properties

Conditions for document properties

Operators used with conditions

Guidelines for using conditions

Examples of using conditions in search queries

Conditions for common properties

Create a condition using common properties when searching mailboxes and sites in the same search. The following table lists the available properties to use when adding a condition.

CONDITION	DESCRIPTION
Date	For email, the date a message was received by a recipient or sent by the sender. For documents, the date a document was last modified.
Sender/Author	For email, the person who sent a message. For documents, the person cited in the author field from Office documents. You can type more than one name, separated by commas. Two or more values are logically connected by the **OR** operator.
Size (in bytes)	For both email and documents, the size of the item (in bytes).
Subject/Title	For email, the text in the subject line of a message. For documents, the title of the document. As previously explained, the Title property is metadata specified in Microsoft Office documents. You can type the name of more than one subject/title, separated by commas. Two or more values are logically connected by the **OR** operator.
Compliance label	For both email and documents, retention labels that have been assigned to messages and documents automatically by autolabel policies or retention labels that have been manually assigned by users. Retention labels are used to classify email and documents for information governance and enforce retention rules based on the settings defined by the label. You can type part of the retention label name and use a wildcard or type the complete label name. For more information about retention labels, see Learn about retention policies and retention labels.

Conditions for mail properties

Create a condition using mail properties when searching mailboxes or public folders. The following table lists the email properties that you can use for a condition. These properties are a subset of the email properties that were previously described. These descriptions are repeated for your convenience.

CONDITION	DESCRIPTION
Message kind	The message type to search. This is the same property as the Kind email property. Possible values: contacts docs email externaldata faxes im journals meetings microsoftteams notes posts rssfeeds tasks voicemail

CONDITION	DESCRIPTION
Participants	All the people fields in an email message. These fields are From, To, Cc, and Bcc.
Type	The message class property for an email item. This is the same property as the ItemClass email property. It's also a multi-value condition. So to select multiple message classes, hold the **CTRL** key and then click two or more message classes in the drop-down list that you want to add to the condition. Each message class that you select in the list will be logically connected by the **OR** operator in the corresponding search query. For a list of the message classes (and their corresponding message class ID) that are used by Exchange and that you can select in the **Message class** list, see Item Types and Message Classes.
Received	The date that an email message was received by a recipient. This is the same property as the Received email property.
Recipients	All recipient fields in an email message. These fields are To, Cc, and Bcc.
Sender	The sender of an email message.
Sent	The date that an email message was sent by the sender. This is the same property as the Sent email property.
Subject	The text in the subject line of an email message.
To	The recipient of an email message in the To field.

Conditions for document properties

Create a condition using document properties when searching for documents on SharePoint and OneDrive for Business sites. The following table lists the document properties that you can use for a condition. These properties are a subset of the site properties that were previously described. These descriptions are repeated for your convenience.

CONDITION	DESCRIPTION
Author	The author field from Office documents, which persists if a document is copied. For example, if a user creates a document and the emails it to someone else who then uploads it to SharePoint, the document will still retain the original author.
Title	The title of the document. The Title property is metadata that's specified in Office documents. It's different than the file name of the document.
Created	The date that a document is created.
Last modified	The date that a document was last changed.

CONDITION	DESCRIPTION
File type	The extension of a file; for example, docx, one, pptx, or xlsx. This is the same property as the FileExtension site property.

Operators used with conditions

When you add a condition, you can select an operator that is relevant to type of property for the condition. The following table describes the operators that are used with conditions and lists the equivalent that is used in the search query.

OPERATOR	QUERY EQUIVALENT	DESCRIPTION
After	`property>date`	Used with date conditions. Returns items that were sent, received, or modified after the specified date.
Before	`property<date`	Used with date conditions. Returns items that were sent, received, or modified before the specified date.
Between	`date..date`	Use with date and size conditions. When used with a date condition, returns items there were sent, received, or modified within the specified date range. When used with a size condition, returns items whose size is within the specified range.
Contains any of	`(property:value) OR (property:value)`	Used with conditions for properties that specify a string value. Returns items that contain any part of one or more specified string values.
Doesn't contain any of	`-property:value` `NOT property:value`	Used with conditions for properties that specify a string value. Returns items that don't contain any part of the specified string value.
Doesn't equal any of	`-property=value` `NOT property=value`	Used with conditions for properties that specify a string value. Returns items that don't contain the specific string.
Equals	`size=value`	Returns items that are equal to the specified size.[1]
Equals any of	`(property=value) OR (property=value)`	Used with conditions for properties that specify a string value. Returns items that are an exact match of one or more specified string values.
Greater	`size>value`	Returns items where the specified property is greater than the specified value.[1]

OPERATOR	QUERY EQUIVALENT	DESCRIPTION
Greater or equal	`size>=value`	Returns items where the specified property is greater than or equal to the specified value.[1]
Less	`size<value`	Returns items that are greater than or equal to the specific value.[1]
Less or equal	`size<=value`	Returns items that are greater than or equal to the specific value.[1]
Not equal	`size<>value`	Returns items that don't equal the specified size.[1]

NOTE

[1] This operator is available only for conditions that use the Size property.

Guidelines for using conditions

Keep the following in mind when using search conditions.

- A condition is logically connected to the keyword query (specified in the keyword box) by the **AND** operator. That means that items have to satisfy both the keyword query and the condition to be included in the results. This is how conditions help to narrow your results.
- If you add two or more unique conditions to a search query (conditions that specify different properties), those conditions are logically connected by the **AND** operator. That means only items that satisfy all the conditions (in addition to any keyword query) are returned.
- If you add more than one condition for the same property, those conditions are logically connected by the **OR** operator. That means items that satisfy the keyword query and any one of the conditions are returned. So, groups of the same conditions are connected to each other by the **OR** operator and then sets of unique conditions are connected by the **AND** operator.
- If you add multiple values (separated by commas or semi-colons) to a single condition, those values are connected by the **OR** operator. That means items are returned if they contain any of the specified values for the property in the condition.
- The search query that is created by using the keywords box and conditions is displayed on the **Search** page, in the details pane for the selected search. In a query, everything to the right of the notation `(c:c)` indicates conditions that are added to the query.
- Conditions only add properties to the search query; the don't add operators. This is why the query displayed in the detail pane doesn't show operators to the right of the `(c:c)` notation. KQL adds the logical operators (according to the previously explained rules) when the executing the query.
- You can use the drag and drop control to resequence the order of conditions. Click on the control for a condition and move it up or down.
- As previously explained, some condition properties allow you to type multiple values. Each value is logically connected by the **OR** operator. This results in the same logic as having multiple instances of the same condition, where each has a single value. The following illustrations show an example of a single condition with multiple values and an example of multiple conditions (for the same property) with a single value. Both examples result in the same query:
  `(filetype:docx) OR (filetype:pptx) OR (filetype:xlsx)`

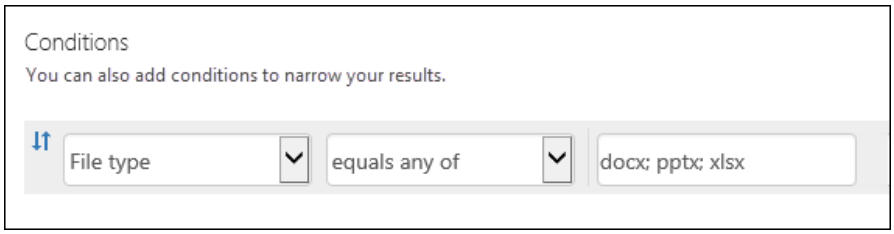

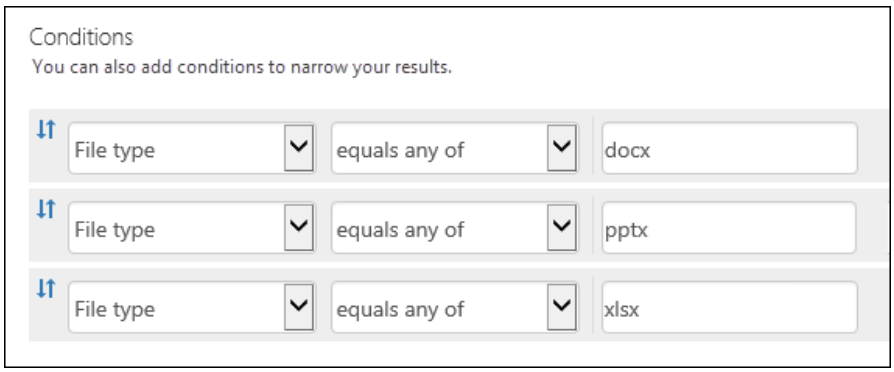

TIP

If a condition accepts multiple values, we recommend that you use a single condition and specify multiple values (separated by commas or semi-colons). This helps ensure the query logic that's applied is what you intend.

Examples of using conditions in search queries

The following examples show the GUI-based version of a search query with conditions, the search query syntax that is displayed in the details pane of the selected search (which is also returned by the **Get-ComplianceSearch** cmdlet), and the logic of the corresponding KQL query.

Example 1

This example returns documents on SharePoint and OneDrive for Business sites that contain a credit card number and were last modified before January 1, 2016.

GUI

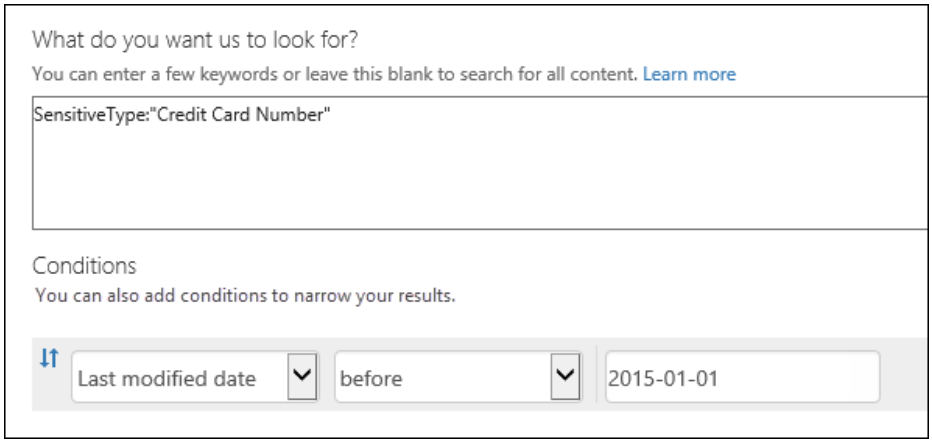

Search query syntax

```
SensitiveType:"Credit Card Number"(c:c)(lastmodifiedtime<2016-01-01)
```

Search query logic

```
SensitiveType:"Credit Card Number" AND (lastmodifiedtime<2016-01-01)
```

Example 2

This example returns email items or documents that contain the keyword "report", that were sent or created before April 1, 2105, and that contain the word "northwind" in the subject field of email messages or in the title

property of documents. The query excludes Web pages that meet the other search criteria.

GUI

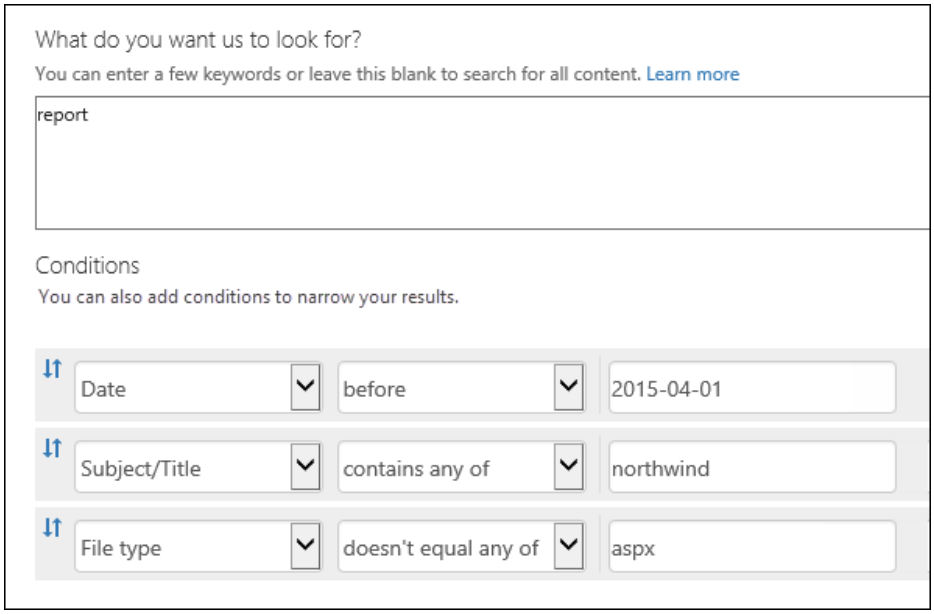

Search query syntax

```
report(c:c)(date<2016-04-01)(subjecttitle:"northwind")(-filetype:aspx)
```

Search query logic

```
report AND (date<2016-04-01) AND (subjecttitle:"northwind") NOT (filetype:aspx)
```

Example 3

This example returns email messages or calendar meetings that were sent between 12/1/2016 and 11/30/2016 and that contain words that start with "phone" or "smartphone".

GUI

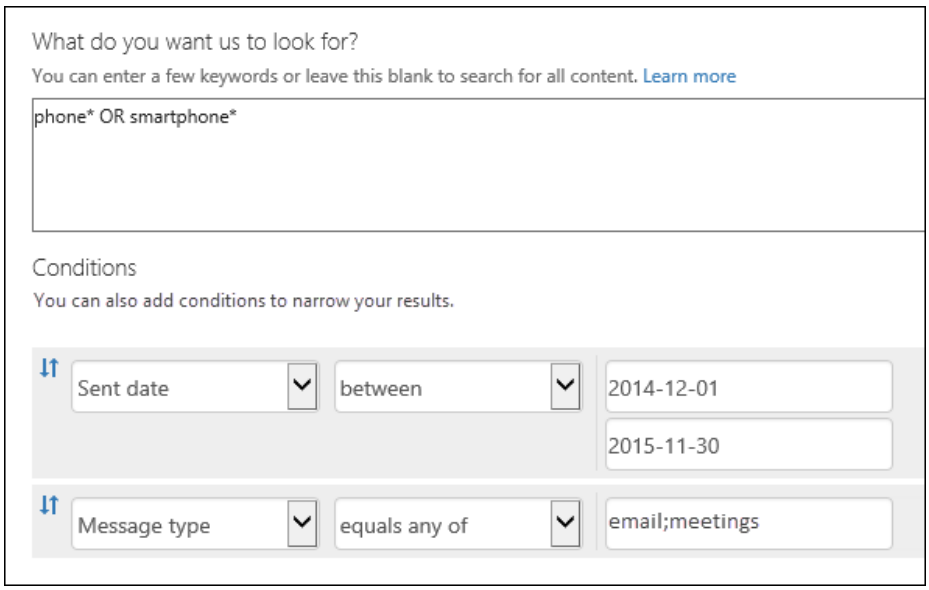

Search query syntax

```
phone* OR smartphone*(c:c)(sent=2016-12-01..2016-11-30)(kind="email")(kind="meetings")
```

Search query logic

```
phone* OR smartphone* AND (sent=2016-12-01..2016-11-30) AND ((kind="email") OR (kind="meetings"))
```

Special characters

Some special characters are not included in the search index and therefore are not searchable. This also includes the special characters that represent search operators in the search query. Here's a list of special characters that are either replaced by a blank space in the actual search query or cause a search error.

```
+ - = : ! @ # % ^ & ; _ / ? ( ) [ ] { }
```

Searching for site content shared with external users

You can also use the Content Search feature in the Security & Compliance Center to search for documents stored on SharePoint and OneDrive for Business sites that have been shared with people outside of your organization. This can help you identify sensitive or proprietary information that's being shared outside your organization. You can do this by using the `ViewableByExternalUsers` property in a keyword query. This property returns documents or sites that have been shared with external users by using one of the following sharing methods:

- A sharing invitation that requires users to sign in to your organization as an authenticated user.
- An anonymous guest link, which allows anyone with this link to access the resource without having to be authenticated.

Here are some examples:

- The query `ViewableByExternalUsers:true AND SensitiveType:"Credit Card Number"` returns all items that have been shared with people outside your organization and contain a credit card number.
- The query
  ```
  ViewableByExternalUsers:true AND ContentType:document AND
  site:"https://contoso.sharepoint.com/Sites/Teams"
  ```
  returns a list of documents on all team sites in the organization that have been shared with external users.

> **TIP**
> A search query such as `ViewableByExternalUsers:true AND ContentType:document` might return a lot of .aspx files in the search results. To eliminate these (or other types of files), you can use the `FileExtension` property to exclude specific file types; for example `ViewableByExternalUsers:true AND ContentType:document NOT FileExtension:aspx`.

What is considered content that is shared with people outside your organization? Documents in your organization's SharePoint and OneDrive for Business sites that are shared by sending a sharing invitation or that are shared in public locations. For example, the following user activities result in content that is viewable by external users:

- A user shares a file or folder with a person outside your organization.
- A user creates and sends a link to a shared file to a person outside your organization. This link allows the external user to view (or edit) the file.
- A user sends a sharing invitation or a guest link to a person outside your organization to view (or edit) a shared file.

Issues using the ViewableByExternalUsers property

While the `ViewableByExternalUsers` property represents the status of whether a document or site is shared with external users, there are some caveats to what this property does and doesn't reflect. In the following scenarios, the value of the `ViewableByExternalUsers` property won't be updated, and the results of a Content Search query that uses this property may be inaccurate.

- Changes to sharing policy, such as turning off external sharing for a site or for the organization. The property will still show previously shared documents as being externally accessible even though external access might have been revoked.
- Changes to group membership, such as adding or removing external users to Microsoft 365 Groups or Microsoft 365 security groups. The property won't automatically be updated for items the group has access to.
- Sending sharing invitations to external users where the recipient hasn't accepted the invitation, and therefore doesn't yet have access to the content.

In these scenarios, the `ViewableByExternalUsers` property won't reflect the current sharing status until the site or document library is recrawled and reindexed.

Searching for site content shared within your organization

As previously explained, you can use the `SharedWithUsersOWSUser` property so search for documents that have been shared between people in your organization. When a person shares a file (or folder) with another user inside your organization, a link to the shared file appears on the **Shared with me** page in the OneDrive for Business account of the person who the file was shared with. For example, to search for the documents that have been shared with Sara Davis, you can use the query `SharedWithUsersOWSUser:"sarad@contoso.com"`. If you export the results of this search, the original documents (located in the content location of the person who shared the documents with Sara) will be downloaded.

Documents must be explicitly shared with a specific user to be returned in search results when using the `SharedWithUsersOWSUser` property. For example, when a person shares a document in their OneDrive account, they have the option to share it with anyone (inside or outside the organization), share it only with people inside the organization, or share it with a specific person. Here's a screenshot of the **Share** window in OneDrive, that shows the three sharing options.

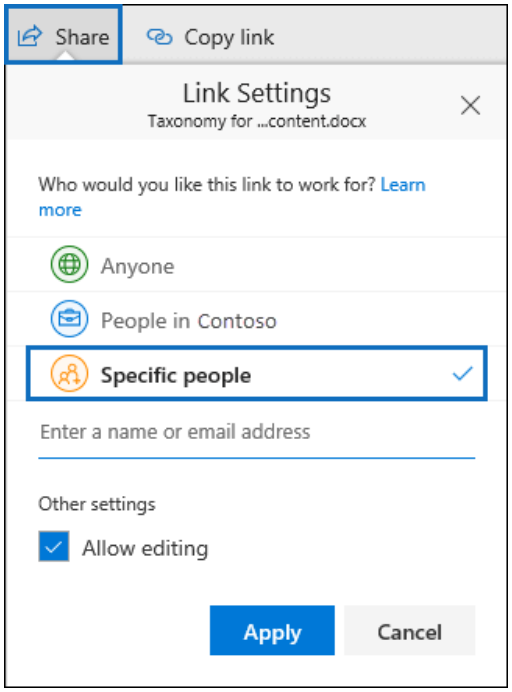

Only documents that are shared by using the third option (shared with **Specific people**) will be returned by a search query that uses the `SharedWithUsersOWSUser` property.

Searching for Skype for Business conversations

You can use the following keyword query to specifically search for content in Skype for Business conversations:

```
kind:im
```

The previous search query also returns chats from Microsoft Teams. To prevent this, you can narrow the search results to include only Skype for Business conversations by using the following keyword query:

```
kind:im AND subject:conversation
```

The previous keyword query excludes chats in Microsoft Teams because Skype for Business conversations are saved as email messages with a Subject line that starts with the word "Conversation".

To search for Skype for Business conversations that occurred within a specific date range, use the following keyword query:

```
kind:im AND subject:conversation AND (received=startdate..enddate)
```

Search tips and tricks

- Keyword searches are not case-sensitive. For example, **cat** and **CAT** return the same results.
- The Boolean operators **AND**, **OR**, **NOT**, and **NEAR** must be uppercase.
- A space between two keywords or two `property:value` expressions is the same as using **AND**. For example, `from:"Sara Davis" subject:reorganization` returns all messages sent by Sara Davis that contain the word reorganization in the subject line.
- Use syntax that matches the `property:value` format. Values are not case-sensitive, and they can't have a space after the operator. If there is a space, your intended value will be a full-text search. For example `to: pilarp` searches for "pilarp" as a keyword, rather than for messages that were sent to pilarp.
- When searching a recipient property, such as To, From, Cc, or Recipients, you can use an SMTP address, alias, or display name to denote a recipient. For example, you can use pilarp@contoso.com, pilarp, or "Pilar Pinilla".
- You can use only prefix wildcard searches; for example, **cat*** or **set***. Suffix searches (***cat**), infix searches (**c*t**), and substring searches (***cat***) are not supported.
- When searching a property, use double quotation marks (" ") if the search value consists of multiple words. For example `subject:budget Q1` returns messages that contain **budget** in the subject line and that contain **Q1** anywhere in the message or in any of the message properties. Using `subject:"budget Q1"` returns all messages that contain **budget Q1** anywhere in the subject line.
- To exclude content marked with a certain property value from your search results, place a minus sign (-) before the name of the property. For example, `-from:"Sara Davis"` excludes any messages sent by Sara Davis.
- You can export items based on message type. For example, to export Skype conversations and chats in Microsoft Teams, use the syntax `kind:im`. To return only email messages, you would use `kind:email`. To return chats, meetings, and calls in Microsoft Teams, use `kind:microsoftteams`.

Search statistics in Advanced eDiscovery

2/18/2021 • 2 minutes to read • Edit Online

One way you can validate your search results is to look at the statistics around your results to make sure they align with your expectations. When a search completes, high-level statistics are shown on the search details flyout:

- Number and volume of items retrieved by the search
- Number and volume of partially indexed or unindexed items that were found in the search locations
- Number of mailboxes and locations searched. In order to view more detailed statistics, click on "Statistics" from the search details flyout.

Summary view

In the Summary view, you can see the search results broken down by location type (e.g. Exchange). For each location type, you can see:

- Number of locations that had items that matched the search conditions
- Number of items from these locations that matched the search conditions
- Total volume of items that matched the search conditions.

Top locations view

In the Top locations view, you see the individual locations with the most matches. For each location, you will see:

- Location name (e.g. SharePoint URL)
- Location type
- Number of items that matched the search conditions
- Total volume of items that matched the search conditions.

Queries view

If you have used (c:s) keyword or keyword rows in your query, then you can see the breakdown of your query in Queries view per location type. For each location type, you will see:

- Part: this column will either have the word "Primary" or "Keyword". "Primary" means that the row presents statistics on the entire query, whereas "Keyword" means one of the query components.
- Query: the actual query component the row refers to. If Part is "Primary", this will be the entire query; if Part was "Keyword", you will see one of the query components here.
    - When you search all contentin mailboxes (by not specifying any keywords), the actual query is (size >= 0) so that all items are returned
    - When you search SharePoint Online and OneDrive for Business sites, the two following components are added:
        - NOT IsExternalContent:1 - excludes any content from an on-premises SharePoint organization

    - NOT isOneNotePage: 1 - excludes all OneNote files because these would be duplicates of any document that matches the search query.
- Number of locations that had items that matched the search conditions.
- Number of items from these locations that matched the search conditions.
- Total volume of items that matched the search conditions.

Add search results to a review set

2/18/2021 • 4 minutes to read • Edit Online

When you're satisfied with the results of a search and you're ready to review and analyze those search results, you can add them to a review set in the case. Copying the original data to the review set also facilitates the review and analysis process by providing you with advanced analytics tools such as themes detection, near-duplicate detection, and email thread identification. You can also add data from non-Microsoft 365 data sources to a review set so that you can review that data in addition to the data you collect from Microsoft 365.

When you add the results of a search to a review set (the review sets in a case are listed on the **Review sets** tab), the following things occur:

- The search is run again. This means the actual search results copied to the review set may be different than the estimated results that were returned when the search was last run.
- All items in the search results are copied from the original data source in the live services, and copied to a secure Azure Storage location in the Microsoft cloud.
- All items (including the content and metadata) are reindexed so that all data in the review set is fully searchable during the review of the case data. Reindexing the data results in thorough and fast searches when you search the data in the review set during the case investigation.
- A file encrypted with a Microsoft encryption technology and is attached to an email message that's returned in the search results is decrypted when the email message and attached file are added to the review set. You can review and query the decrypted file in the review set. You have to be assigned the RMS Decrypt role to add decrypted email attachments to a review set. For more information, see Decryption in Microsoft 365 eDiscovery tools.

To add data to a review set, click a search on the **Searches** tab, and then click **Add results to review set** on the flyout page.

You can add to an existing review set or create a new review set. If adding to a new review set, specify the name and then click **Add** to display the flyout page.

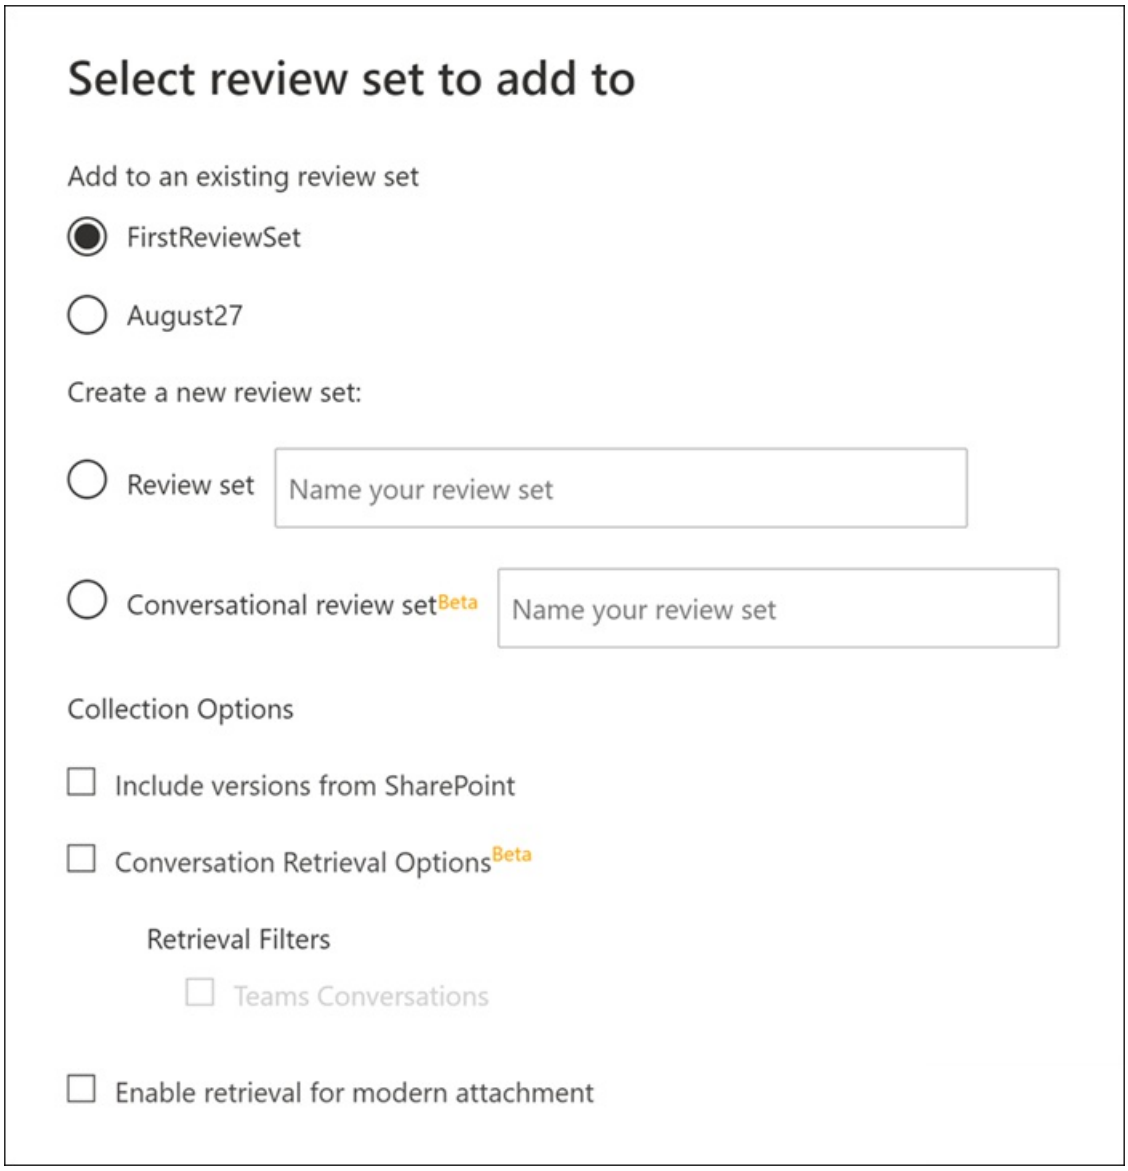

Adding data to a review set is a long-running process. This process includes gathering items from the original data sources in Microsoft 365 (for example, from mailboxes and sites), copying them to the Azure Storage location (this copying process is also called *ingestion*), and then reindexing the items. You can track the progress on the **Jobs** tab or on the **Searches** tab by monitoring the status in the **Added data to review set** column. After the review set processing is completed, click the **Review sets** tab in the case, and then click the review set to start the process of filtering, reviewing, tagging, and exporting data in the review set.

Define options to scope your collection for review

When you add the content of a search to an existing or new review set, you have the following options for how to collect the content for review:

- **Include versions from SharePoint (beta)**: Use this option to enable the collection of all version of a SharePoint document per the version limits and search parameters of the collection. Selecting this option will significantly increase the size of items that are added to the review set.
- **Conversation retrieval options**: Items added to the review set are enabled for threaded conversations to help review content in context of the back and forth conversation. For more information, see Review conversations in Advanced eDiscovery.
- **Enable retrieval for modern attachments**: Use this option to include modern attachments or linked files in the collection for further review. For more information about the searchable properties related to modern attachments, see Document metadata fields in Advanced eDiscovery.

Add a sample to a review set

If you want to validate the results of a search more thoroughly before adding all of them to a review set, you can add a sample of the search results to a review set instead of adding everything.

To add a sample to a review set, click a search on the **Searches** tab and click **Sample** on the flyout page. On the **Sampling parameters** page, choose one of the following options:

- **Confidence level %** and **Confidence interval %** - The items added to the review set will be determined by the statistical parameters that you set. If you typically use a confidence level and interval when sampling results, specify them in the drop-down boxes. Otherwise, use the default settings.
- **Random sample %** - The items added to the review set is based on a random selection of the specified percentage of the total number of items returned by the search.

After selecting and configuring one of the previous options, choose a review set to add the sample to and then click **Send**. Again, you can track the progress on the **Jobs** tab or on the **Searches** tab by monitoring the status in the **Added data to review set** column.

Optical character recognition

When you add search results to a review set, optical character recognition (OCR) functionality in Advanced eDiscovery automatically extracts text from images, and includes the image text with the data that's added to a review set. You can view the extracted text in the Text viewer of the selected image file in the review set. This lets you conduct further review and analysis on text in images. OCR is supported for loose files, email attachments, and embedded images. For a list of image file formats that are supported for OCR, see Supported file types in Advanced eDiscovery.

You have to enable OCR functionality for each case that you create in Advanced eDiscovery. For more information, see Configure search and analytics settings.

Manage review sets in Advanced eDiscovery

5/5/2020 • 2 minutes to read • Edit Online

Review sets are a static set of documents where you can analyze, query, view, tag, and export data in a case. For more information about performing these tasks, see:

- Analyze data in a review set
- Query the data in a review set
- View documents in a review set
- Tag documents in a review set
- Export case data

Create a review set

Review sets can be created on the **Review sets** tab by clicking **+ Add review set**.

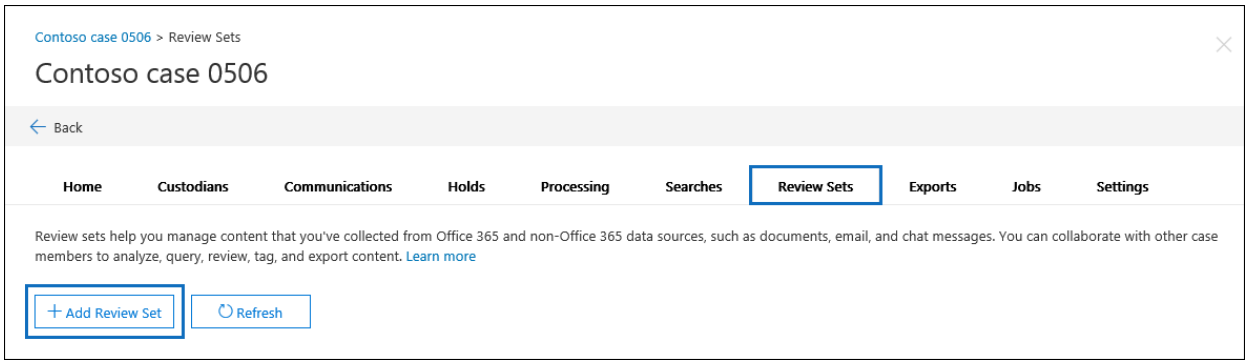

On the **Add review set** flyout page, type a name for the review set and then click **Save**. The new review set is displayed in the list on the **Review sets** tab.

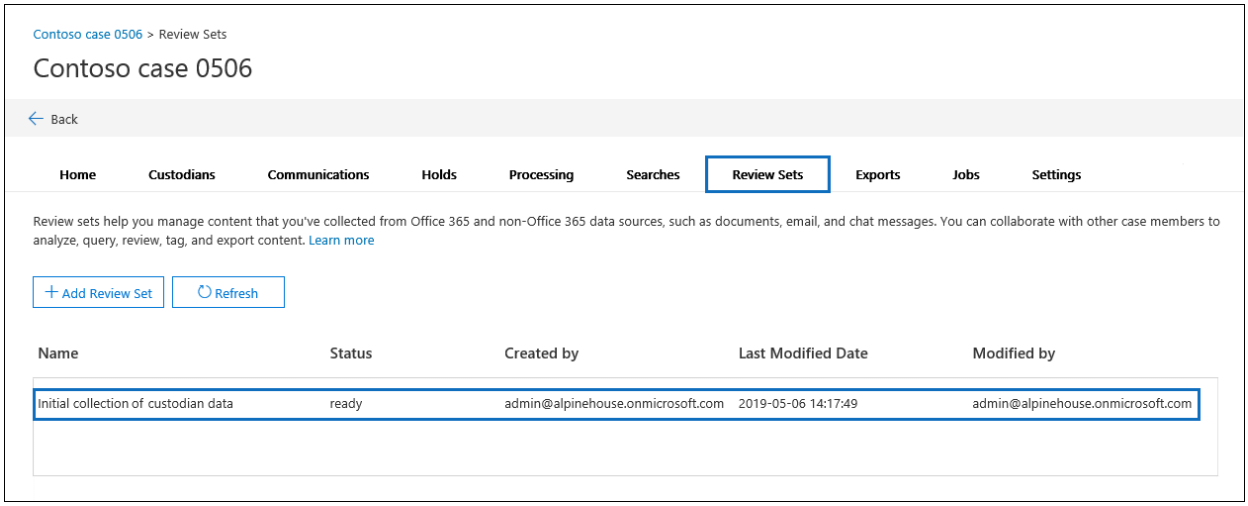

There are three different ways to add data to a review set in an Advanced eDiscovery case.

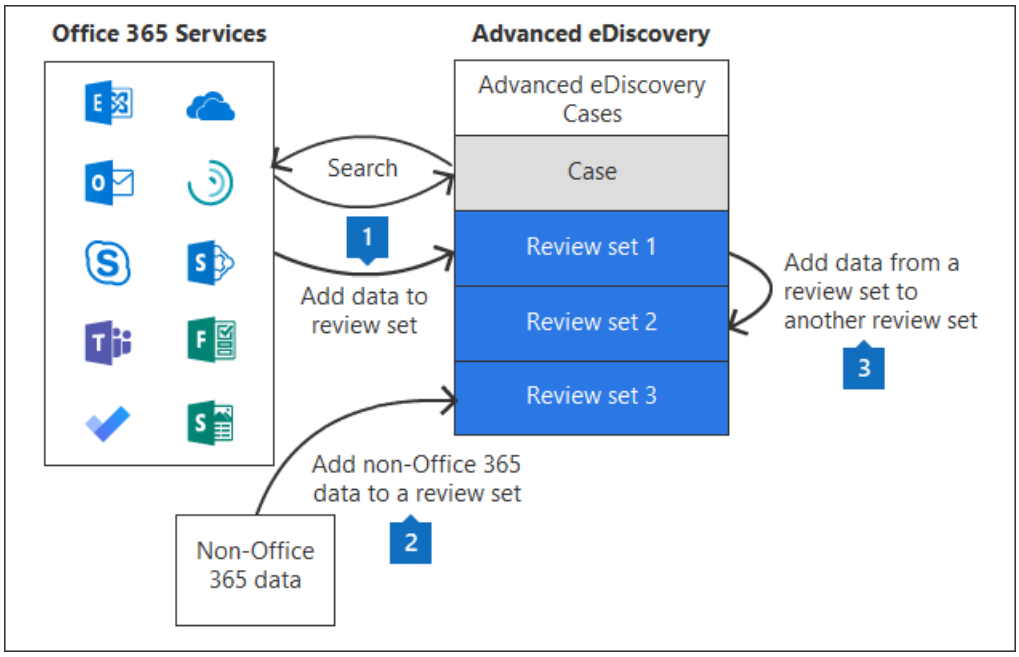

1. Add search results to a review set
2. Load non-Microsoft 365 data into a review set
3. Add data to a review set from another review set

Load non-Microsoft 365 data into a review set

11/2/2020 • 3 minutes to read • Edit Online

Not all documents that you need to analyze in Advanced eDiscovery are located in Microsoft 365. With the non-Microsoft 365 data import feature in Advanced eDiscovery, you can upload documents that aren't located in Microsoft 365 to a review set. This article shows you how to bring your non-Microsoft 365 documents into Advanced eDiscovery for analysis.

Requirements to upload non-Office 365 content

Using the upload non-Microsoft 365 feature described in this article requires that you have the following:

- All custodians that you want to associate non-Microsoft 365 content to must be assigned the appropriate license. For more information, see Get started with Advanced eDiscovery.
- An existing Advanced eDiscovery case.
- Custodians must be added to the case before you can upload and associate the non-Microsoft 365 data to them.
- Non-Microsoft 365 data must be a file type that's supported by Advanced eDiscovery. For more information, see Supported file types in Advanced eDiscovery.
- All files that are uploaded to a review set must be located in folders, where each folder is associated with a specific custodian. The names for these folders must use the following naming format: *alias@domainname*. The alias@domainname must be the user's Microsoft 365 alias and domain. You can collect all the alias@domainname folders in a root folder. The root folder can only contain the alias@domainname folders. Loose files in the root folder aren't supported.

  The folder structure for the non-Microsoft 365 data that you want to upload would be similar to the following example:

  - c:\nonO365\abraham.mcmahon@contoso.com
  - c:\nonO365\jewell.gordon@contoso.com
  - c:\nonO365\staci.gonzalez@contoso.com

  Where abraham.mcmahon@contoso.com, jewell.gordon@contoso.com, and staci.gonzalez@contoso.com are the SMTP addresses of custodians in the case.

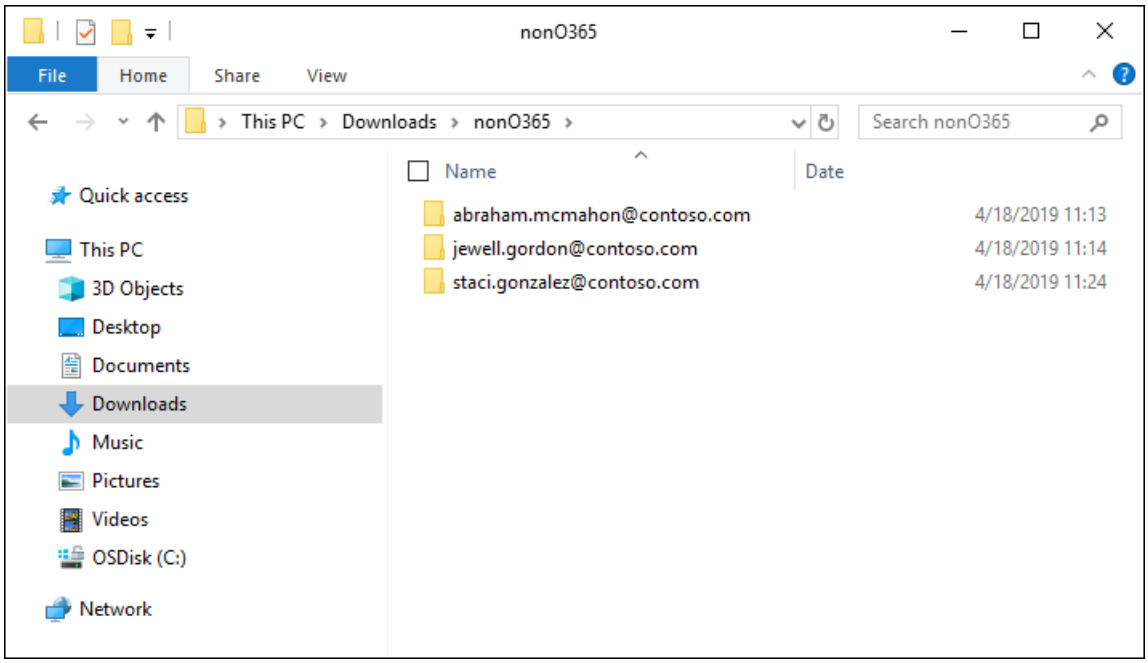

- An account that is assigned to the eDiscovery Manager role group (and added as eDiscovery Administrator).
- The AzCopy v8.1 tool installed on a computer that has access to the non-Microsoft 365 content folder structure. To install AzCopy, see Transfer data with the AzCopy v8.1 on Windows. Be sure to install AzCopy in the default location, which is **%ProgramFiles(x86)%\Microsoft SDKs\Azure\AzCopy**. You must use AzCopy v8.1. Other versions of AzCopy may not work when loading non-Microsoft 365 data in Advanced eDiscovery.

Upload non-Microsoft 365 content into Advanced eDiscovery

1. As an eDiscovery Manager or eDiscovery Administrator, open Advanced eDiscovery, and go to the case that the non-Microsoft 365 data will be uploaded to.
2. Click **Review sets**, and then select the review set to upload the non-Microsoft 365 data to. If you don't have a review set, you can create one.
3. In the review set, click **Manage review set**, and then click **View uploads** on the **Non-Microsoft 365 data** tile.
4. Click **Upload files** to start the data import wizard.

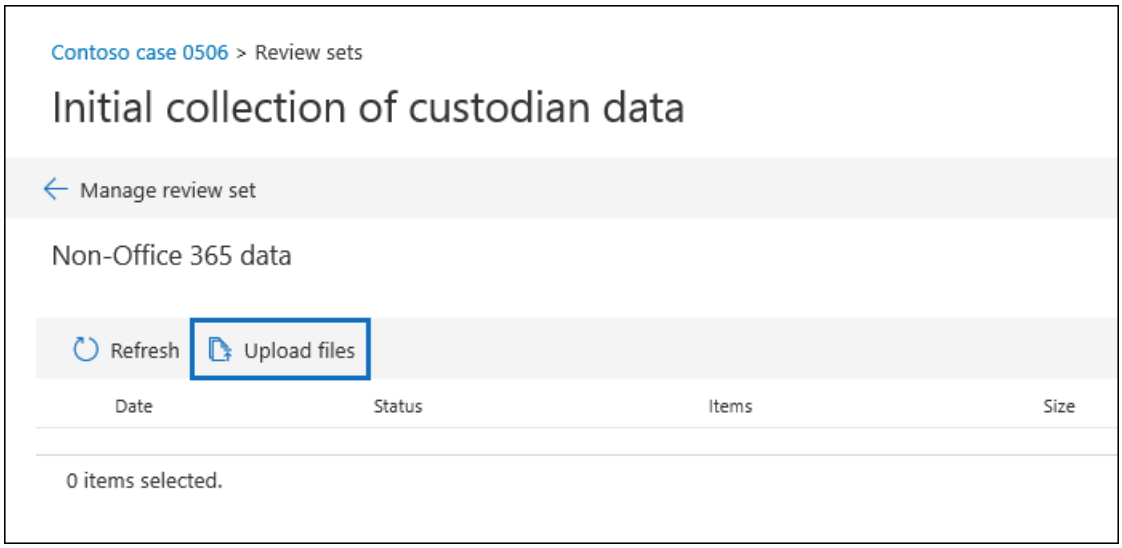

The first step in the wizard prepares a secure Microsoft-provided Azure Storage location to upload the

files to. When the preparation is completed, the **Next: Upload files** button becomes active.

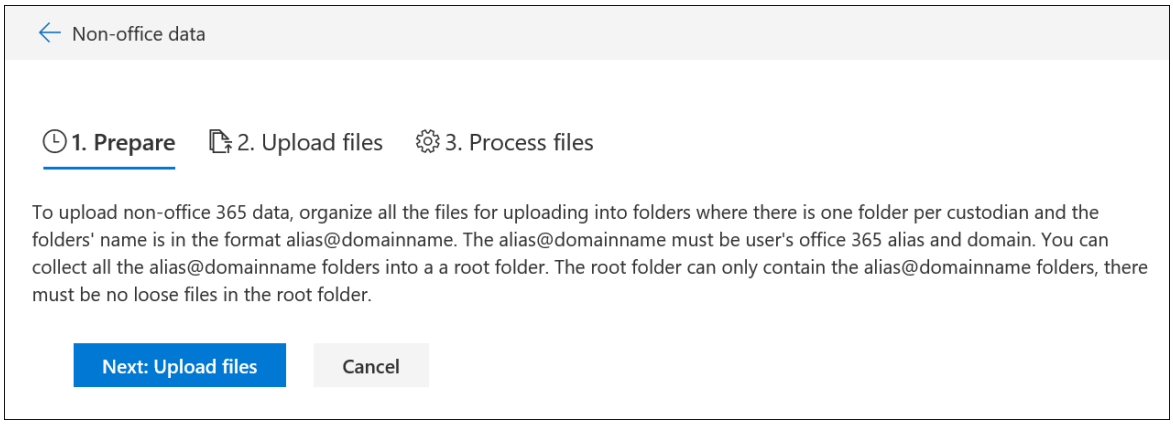

5. Click **Next: Upload files**.

6. On the **Upload files** page, do the following:

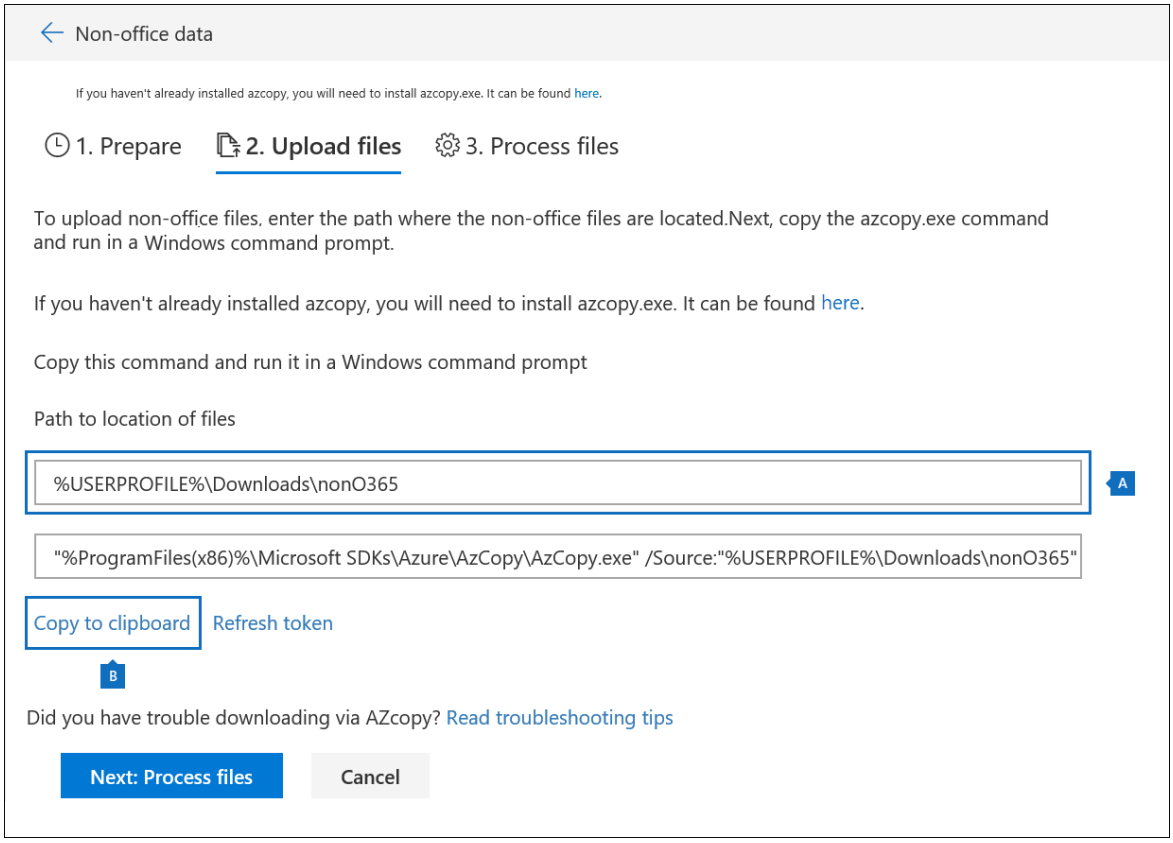

a. In the **Path to location of files** box, verify or type the location of the root folder where you've stored the non-Microsoft 365 data you want to upload. For example, for the location of the example files shown in the **Before you begin section**, you would type **%USERPROFILE\Downloads\nonO365**. Providing the correct location ensures the AzCopy command displayed in box under the path is properly updated.

b. Click **Copy to clipboard** to copy the command that is displayed in the box.

7. Start a Windows command prompt, paste the command that you copied in the previous step, and then press **Enter** to start the AzCopy command. After you start the command, the non-Microsoft 365 files will be uploaded to the Azure Storage location that was prepared in step 4.

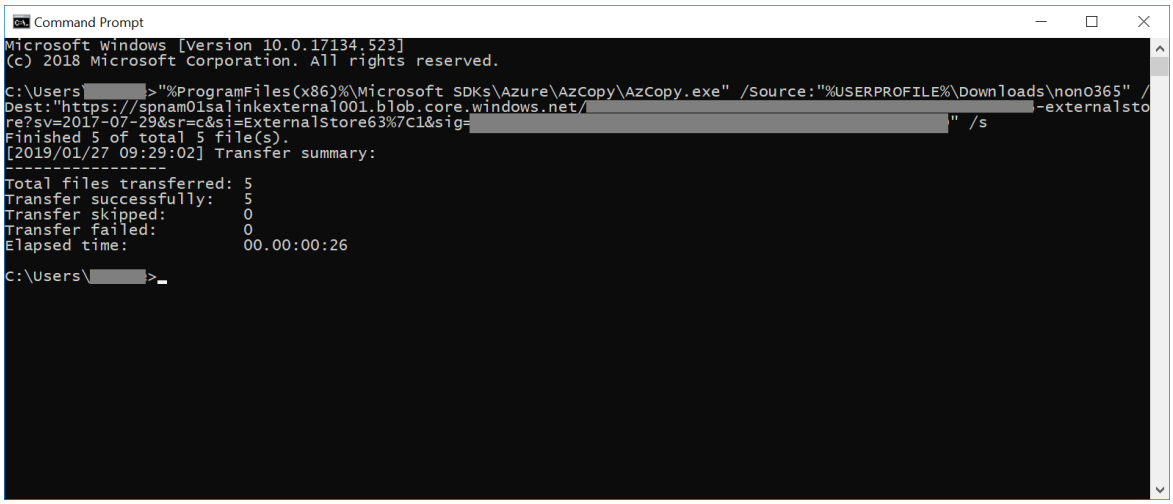

> **NOTE**
>
> As previously stated, you must use AzCopy v8.1 to successfully use the command that's provided on the **Upload files** page. If the supplied AzCopy command fails, please see Troubleshoot AzCopy in Advanced eDiscovery.

8. Go back to the Security & Compliance Center, and click **Next: Process files** in the wizard. This initiates processing, text extraction, and indexing of the non-Microsoft 365 files that were uploaded to the Azure Storage location.

9. Track the progress of processing the files on the **Process files** page or on the **Jobs** tab by viewing a job named **Adding non-Microsoft 365 data to a review set**. After the job is finished, the new files will be available in the review set.

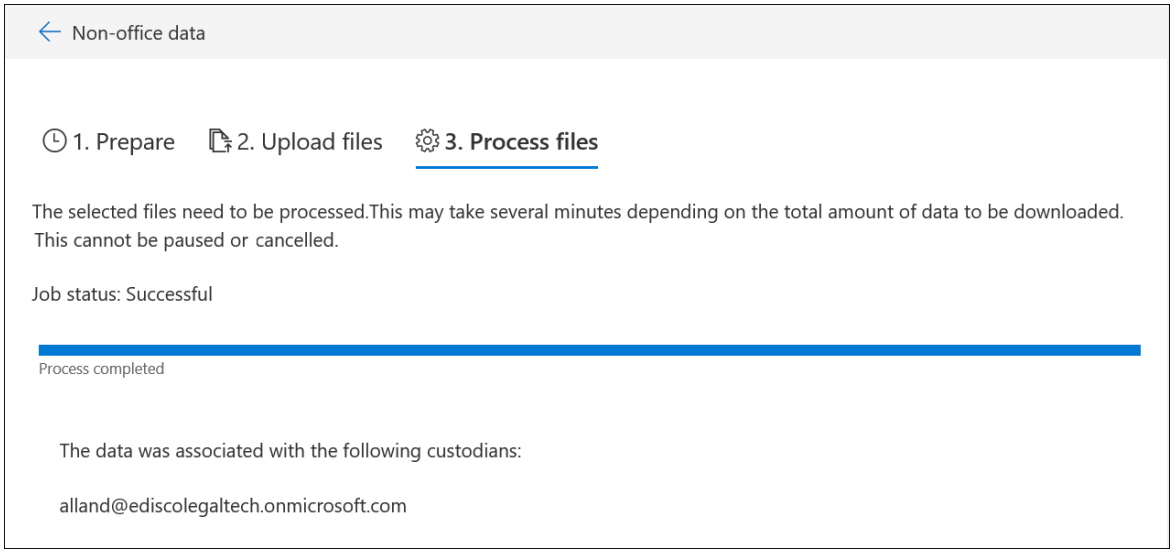

10. After the processing is finished, you can close the wizard.

Add data to a review set from another review set

11/2/2020 • 2 minutes to read • Edit Online

In some cases, it may be necessary to select documents from one review set and work with them individually in another review set. This is especially useful if you've culled content in a review set and want to run analytics on the subset of data.

Follow the workflow in this article to add content from one review set to another.

Create a review set

Before you start, you'll need to create a review set to add the data to. A new review set can be added on the **Review sets** tab of the case. For more information, see Create a review set.

Step 1: Identify content to add to another review set

You can add content from one review set to another one by selecting specific documents in the source review set or by selecting all items returned by review set query. If you're adding selected items, select the items, select **Action**, and then select **Add to another review set**.

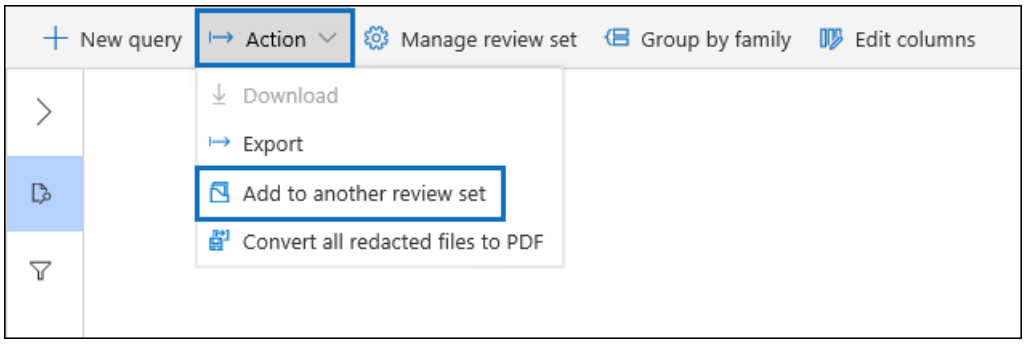

Step 2: Specify options for adding to another review set

In the **Add to another review set options** flyout page, choose the review set you want to add the items to. Choose whether to add **All search results** or **Selected items**. **Additional information** provides options to include all metadata from the items and whether to include the tags (by selecting the **Labels** check box) from the source review set when the documents are added to the new review set.

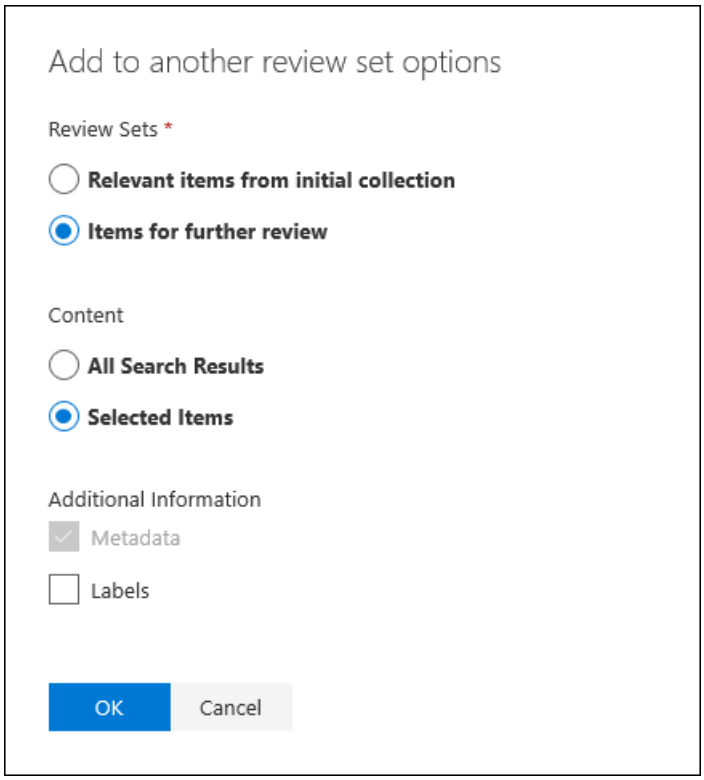

After you click **Ok**, a new job (named **Adding data to another review set**) is created to add the content to another review set. You can go to the **Jobs** tab and monitor the progress of this job. For more information, see Manage jobs.

Advanced eDiscovery dashboard for review sets

11/2/2020 • 2 minutes to read • Edit Online

For some cases in Advanced eDiscovery, you may have a large volume of documents and email messages that need to be reviewed. Before you start the review process, you may want to quickly analyze your corpus to identify trends or key statistics that will help you develop your review strategy. To do this, you can use the Advanced eDiscovery dashboard for review sets to quickly analyze your corpus.

Step 1: Create a widget on the review set dashboard

1. In the Security & Compliance Center, go to **eDiscovery > Advanced eDiscovery** to display the list of cases in your organization.
2. Select an existing case.
3. Click the **Review Set** tab, and then select a review set.
4. In the **Individual results** dropdown list, click **Search profile view**.

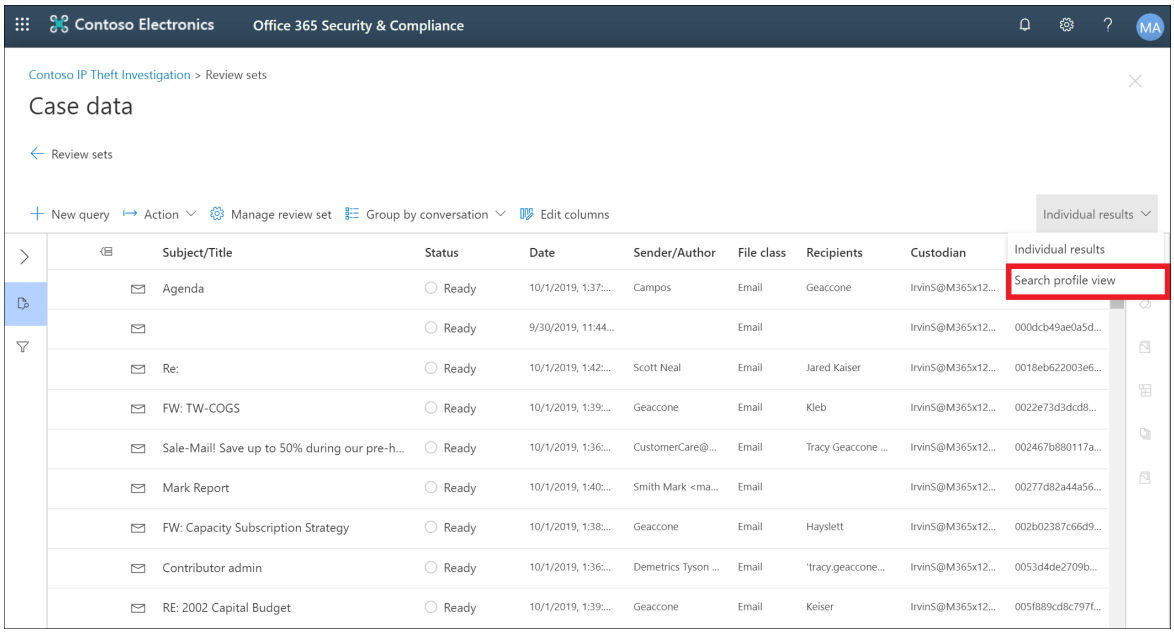

The **Search profile view** page is displayed; the first time you display this page, three default widgets are displayed.

5. Click the **New widget** and then select one of the following items:

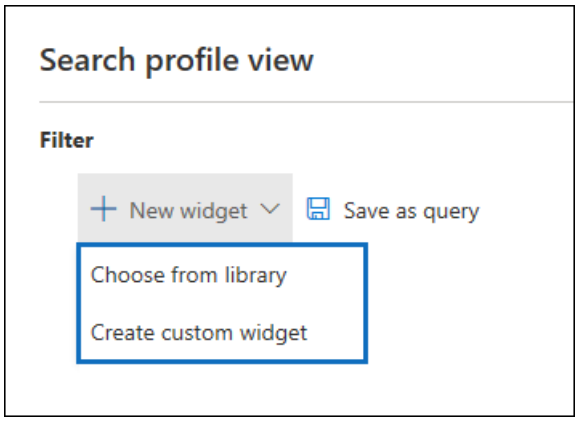

- **Choose from library:** Displays a default library of widgets. You click a widget and then click **Add** to add it to the widgets on the **Search profile view** page.
- **Create custom widget:** Displays a flyout page that you can use to set up a custom widget.

6. To create a custom widget, do the following on the **Add widget** flyout page:

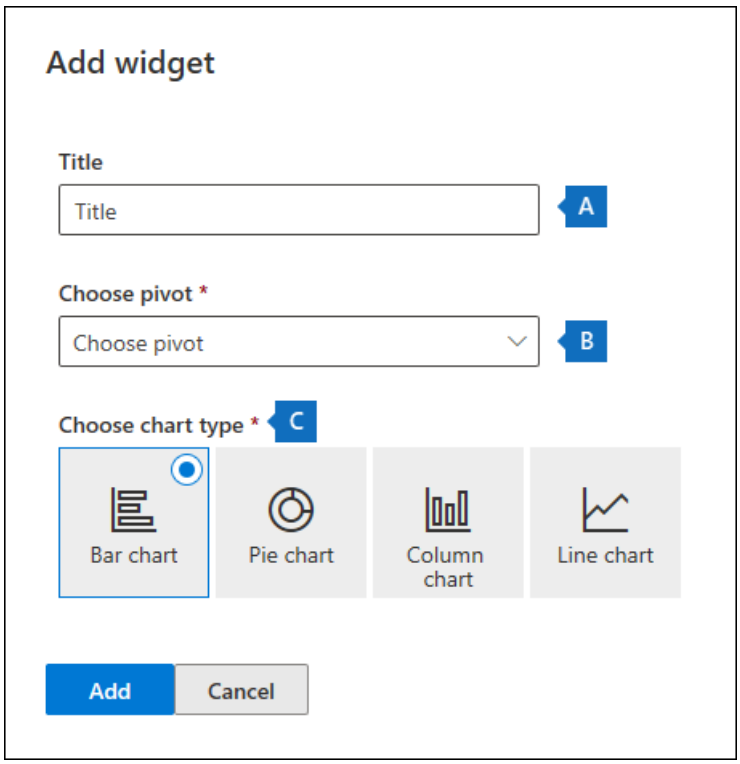

a. Type a name for the widget, which is displayed in the widget title bar. Naming a widget is required, but it's helpful to identify the widget data.

b. Select a property in the **Choose pivot** dropdown list that will be used for the widget data. The items in this list are the searchable properties for the items in the review set. For a description of these properties, see Document metadata fields in Advanced eDiscovery. The pivot options for the widget are listed in the **Searchable field name** column in this topic.

c. Select a chart type to display the data from the selected pivot property.

7. Click **Add** to create the custom widget and display it on the **Search profile view** page.

Step 2: Create a review set search query

1. Click ... in the widget title bar, and then click **Apply condition**.

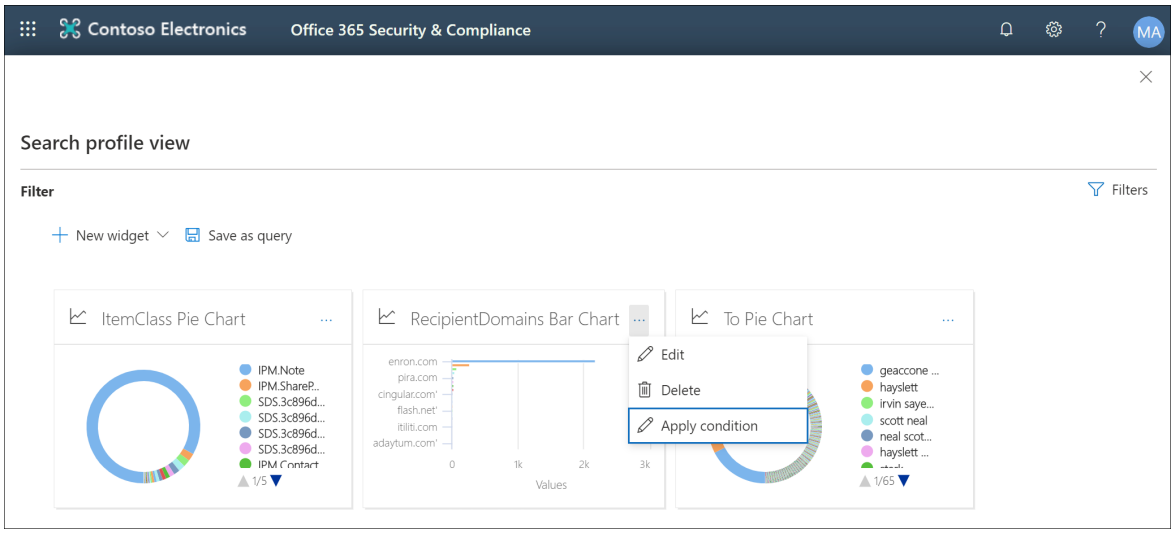

2. On the flyout page, click an element on the widget key or widget chart to create a filter.

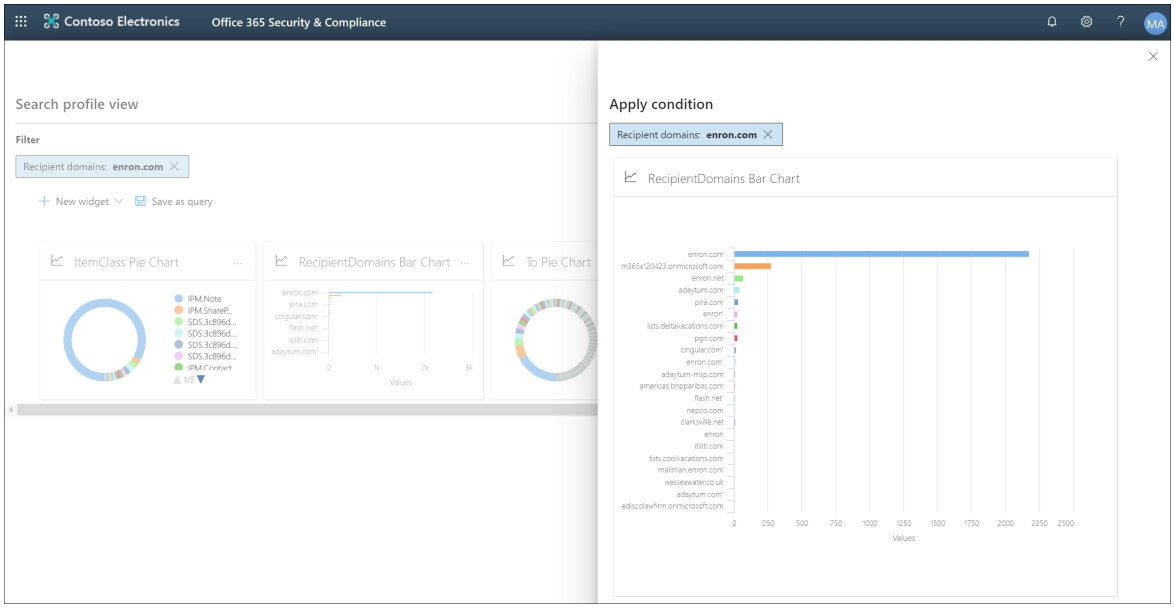

3. Repeat steps 1-2 for other widgets multiple widgets.

4. When you're done, click **Save as query** to save your conditions as a new search query for the review set.

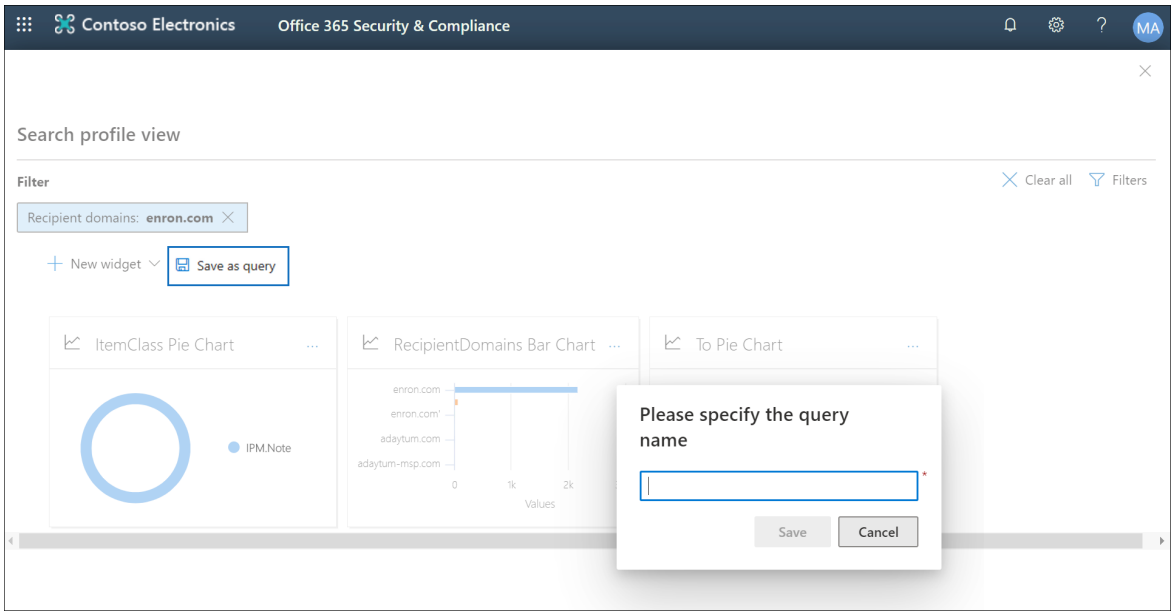

5. Close the **Search profile view** to return to the search results view.

If you have created any visual filters, the resulting query is applied to the search results that are displayed, and the search query that you saved in step 4 is displayed under **Saved queries**. For more information about review set queries, see Query the data in a review set.

View documents in a review set in Advanced eDiscovery

11/2/2020 • 2 minutes to read • Edit Online

Advanced eDiscovery displays content via several viewers each with different purposes. The various viewers can be used by clicking on any document within a review set. The viewers currently provided are:

- File metadata
- Native view
- Text view
- Annotate view
- Converted view

File metadata

This panel can be toggled on/off to display various metadata associated with the document. Although the search results grid can be customized to display specific metadata, there are instances where scrolling horizontally can be difficult while reviewing data. The File metadata panel allows a user to toggle on a view within the viewer.

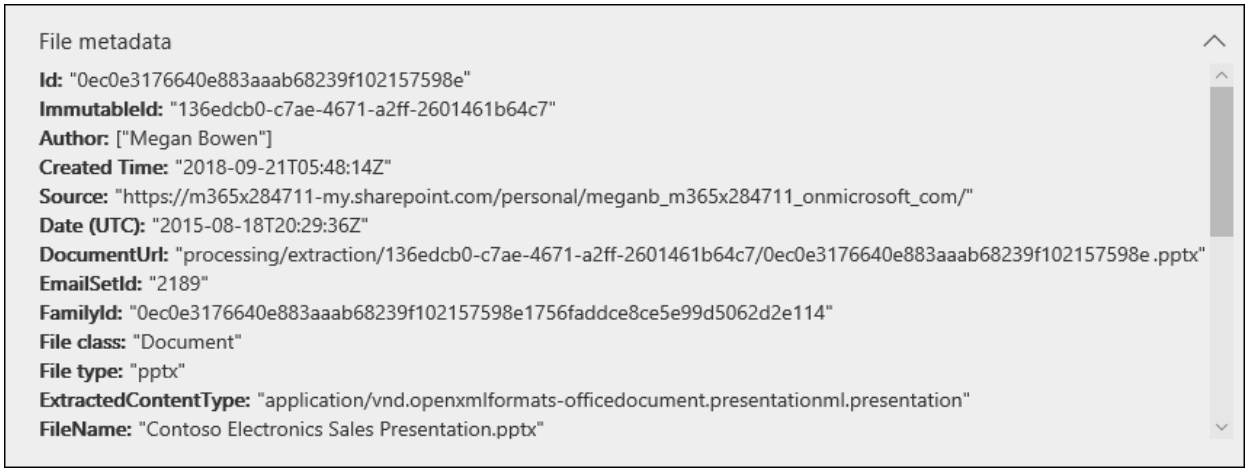

Native view

The Native viewer displays the richest view of a document. It supports hundreds of file types and is meant to display the truest to native experience possible. For Microsoft Office files, the viewer uses the web version of Office apps to display content such as document comments, Excel formulas, hidden rows/columns, and PowerPoint notes.

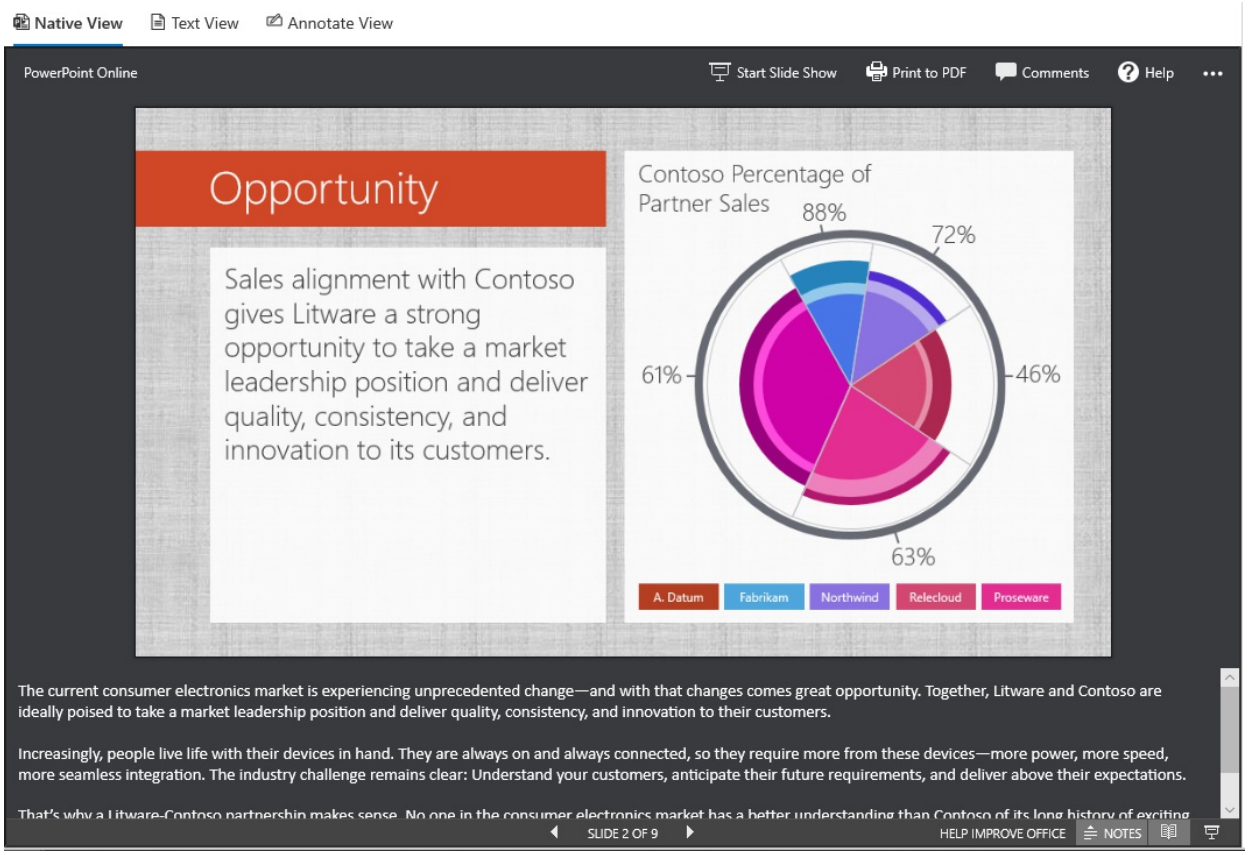

Text view

The Text viewer provides a view of the extracted text of a file. It ignores any embedded images and formatting but is very effective if you are trying to understand the content quickly. Text view also includes these features:

- Line counter makes it easier to reference specific portions of a document
- Search hit highlighting that will highlight terms within the document as well as the scrollbar
- Diff view provides a comparison view that highlights textual differences when viewing Near Duplicate documents

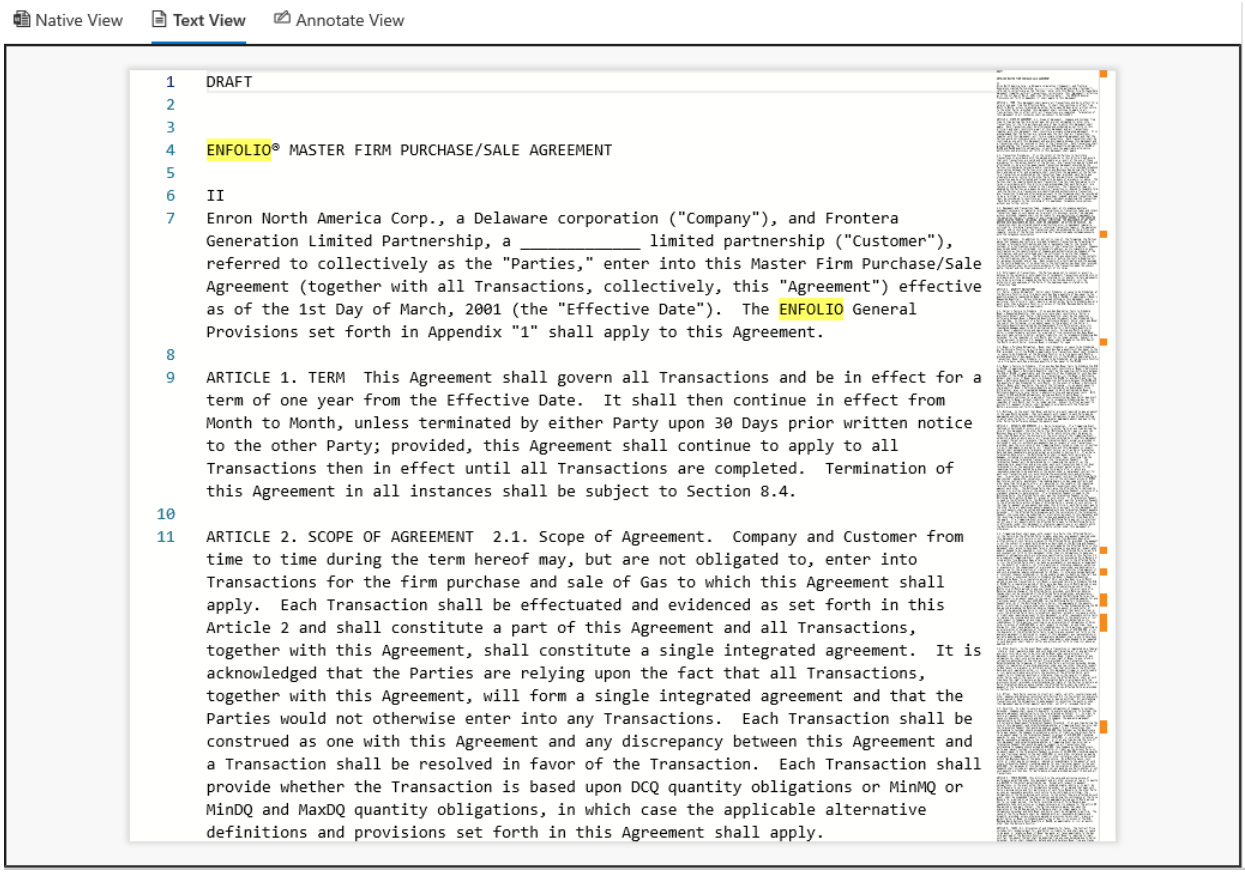

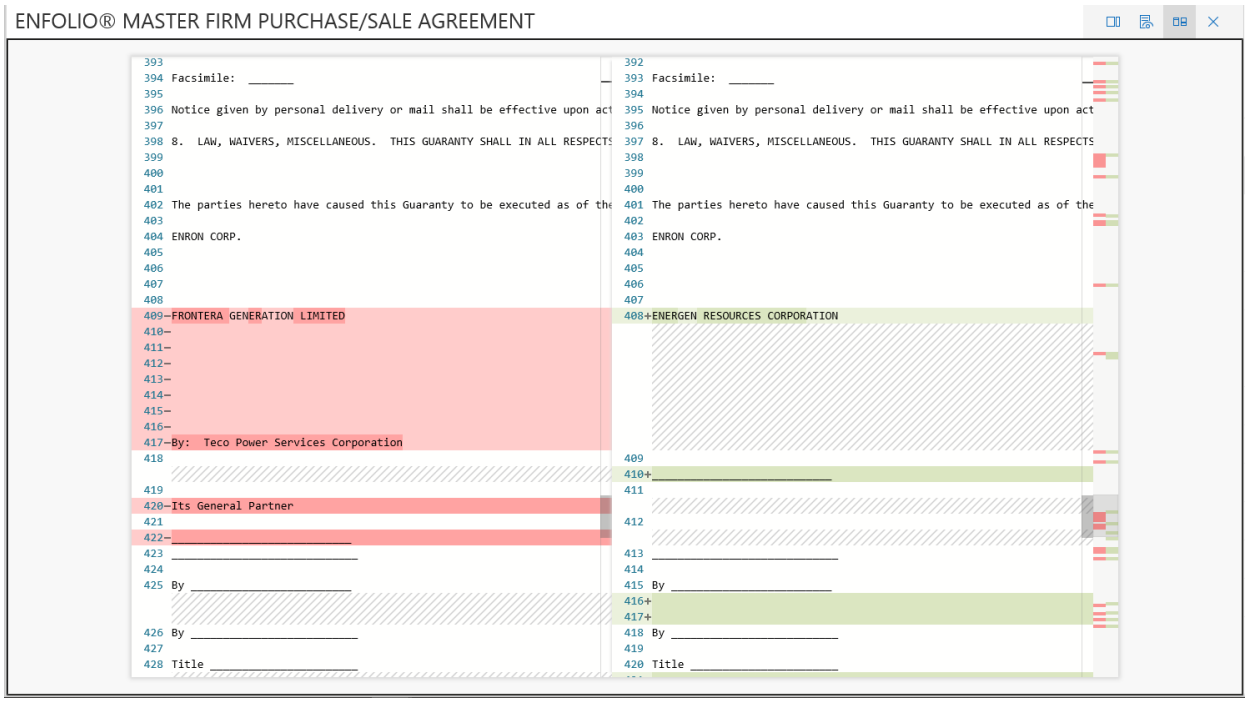

Annotate view

The Annotate view provides features that allow users to apply markup on a document including:

- Area redactions – users can draw a box on the document in order to hide sensitive content
- Pencil – users can free-hand draw on a document in order to bring attention to certain portions of a document
- Select annotations - users can select annotations on a document in order to delete
- Toggle annotation transparency – makes annotations semi-transparent in order to view the content behind the annotation

- Previous page – navigates to previous page
- Next page – navigates to the next page
- Go to page – user can enter a specific page number to navigate to
- Zoom – set zoom level for annotate view
- Rotate – user can rotate document clockwise
- Search – user can search within a document and navigate to the various hits within the document

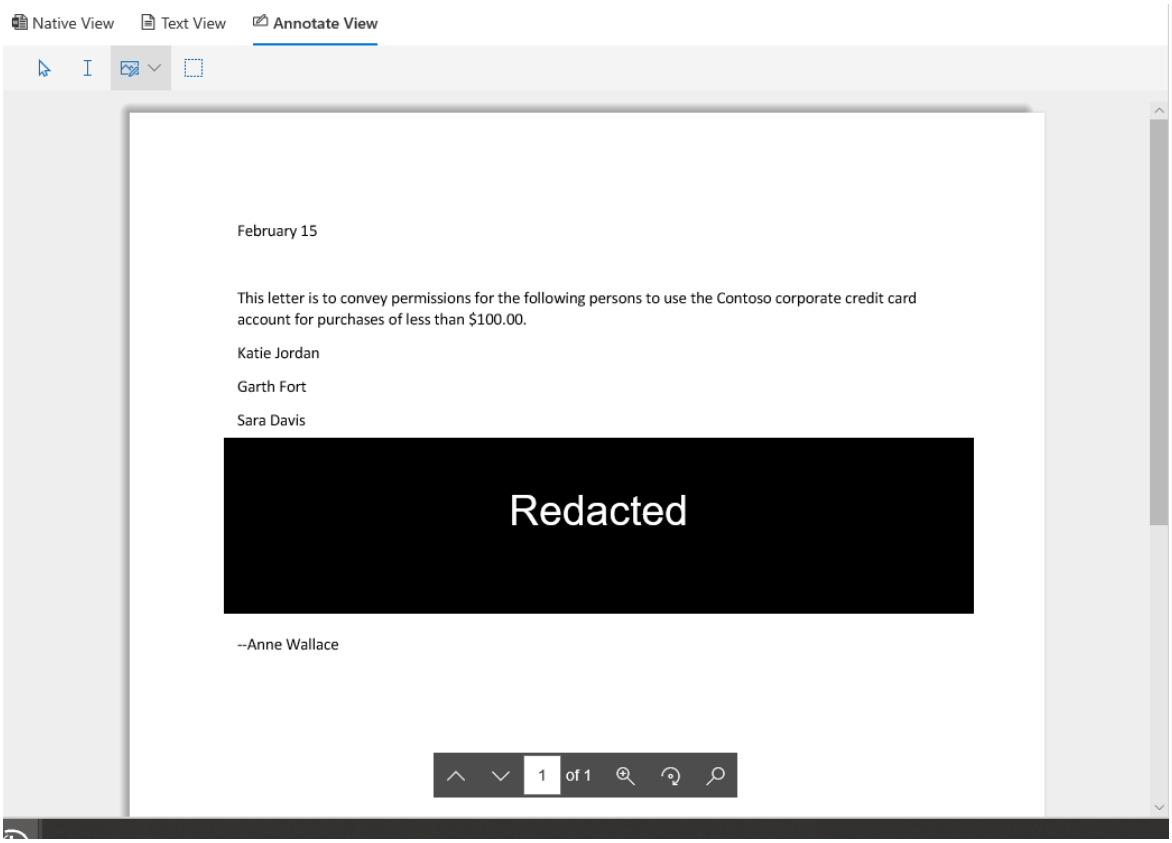

Dashboard View

The dashboard view allows you to visualize and summarize the data in your search results grid. In this view, you can create custom widgets to make analyzing and reporting on your review set intuitive and easy. Once you have created your widgets, you can interact with them to get item counts or to create a search.

Query the data in a review set

11/2/2020 • 2 minutes to read • Edit Online

In most cases, it will be useful to be able to dig deeper into the data in a review set and organize that data to facilitate a more efficient review. Using Queries in a review set helps you focus on a subset of documents that meet the criteria of your review.

Creating and running a query in a review set

To create and run a query on the documents in a review set, select **New query** in the review set. After you name your query and define the conditions, select **Save** to save and run the query. To run a query that has been previously saved, select a saved query.

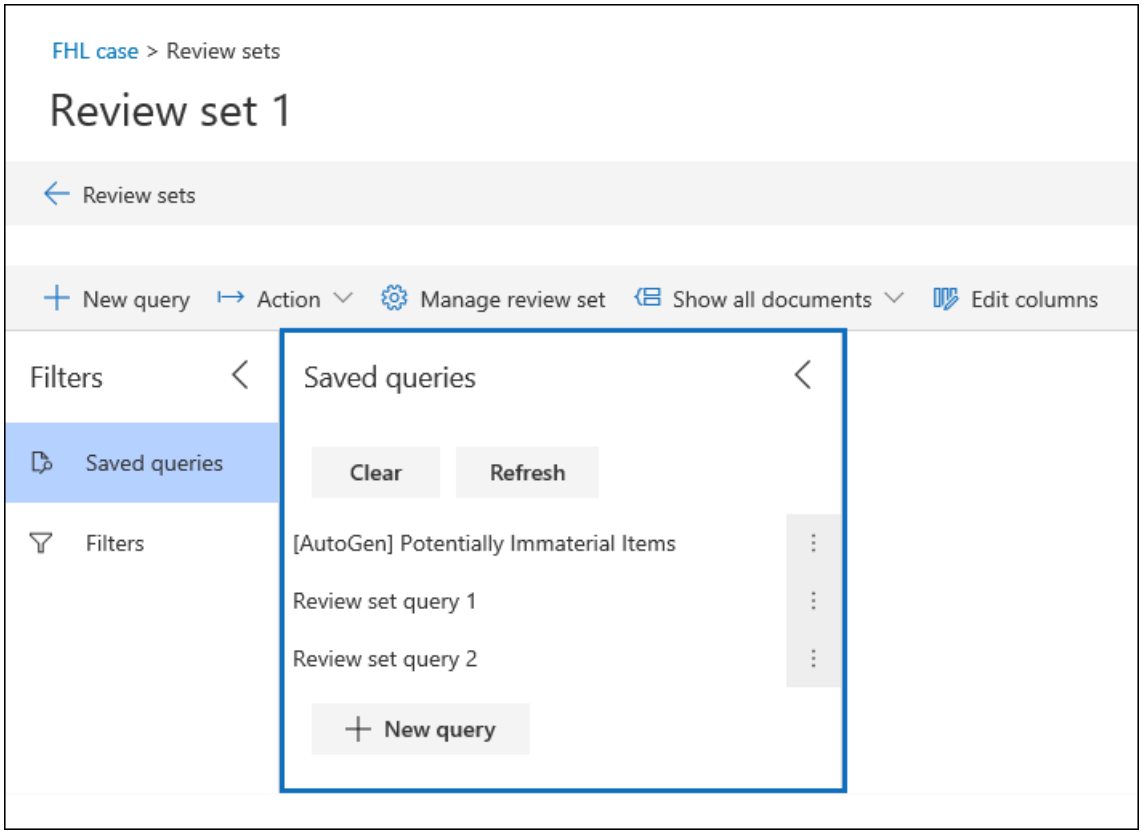

Building a review set query

You can build a query by using a combination of keywords, properties, and conditions in the Keywords condition. You can also group conditions as a block (called a *condition group*) to build a more complex query. For a list and description of metadata properties that you can search, see Document metadata fields in Advanced eDiscovery.

Conditions

Every searchable field in a review set has a corresponding condition that you can use to build your query.

There are multiple types of conditions:

- Freetext: A freetext condition is used for text fields such as subject. You can list multiple search terms by separating them out with a comma.
- Date: A date condition is used for date fields such as last modified date.

- Search options: A search options condition will provide a list of possible values for the particular field in your review set. This is used for fields, such as sender, where there is a finite number of possible values in your review set.
- Keyword: A keyword condition is a specific instance of freetext condition that you can use to search for terms, or use KQL-like query language in. See below for more detail.

Query language

In addition to conditions, you can use a KQL-like query language in the Keywords condition to build your query. The query language for review set queries supports standard Boolean operators, such as **AND**, **OR**, **NOT**, and **NEAR**. It also supports a single-character wildcard (?) and a multi-character wildcard (*).

Filters

In addition to queries that you can save, you can use review set filters to quickly apply additional conditions to a review set query. Using filters help you further refine the results displayed by a review set query.

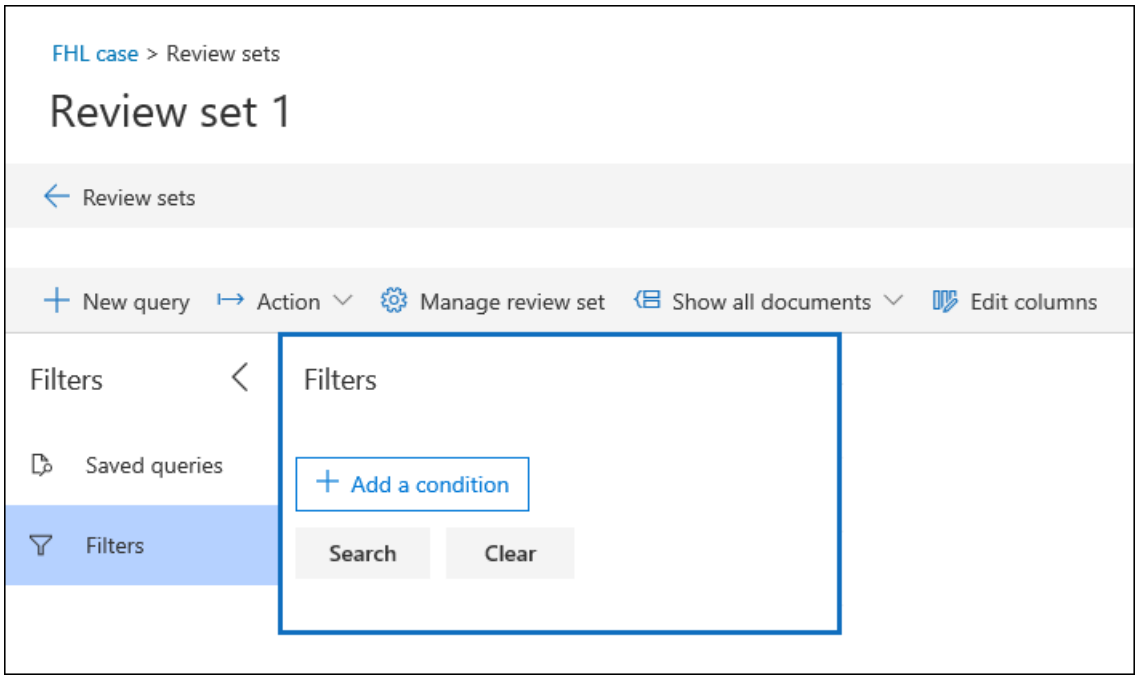

Filters differ from queries in two significant ways:

- Filters are transient. They don't persist beyond the existing session. In other words, you can't save a filter. Queries are saved to the review set, and access them whenever open the review set.
- Filters are always additive. Filters are applied in addition to the current review set query. Applying a different query will replace the results returned by the current query.

Tag documents in a review set in Advanced eDiscovery

11/2/2020 • 3 minutes to read • Edit Online

Organizing content in a review set is important to complete various workflows in the eDiscovery process. This includes:

- Culling unnecessary content
- Identifying relevant content
- Identifying content that must be reviewed by an expert or an attorney

When experts, attorneys, or other users review content in a review set, their opinions related to the content can be captured by using tags. For example, if the intent is to cull unnecessary content, a user can tag documents with a tag such as "non-responsive". After content has been reviewed and tagged, a review set search can be created to exclude any content tagged as "non-responsive", which eliminates this content from the next steps in the eDiscovery workflow. The tag panel can be customized for every case so that the tags can support the intended review workflow.

Tag types

Advanced eDiscovery provides two types of tags:

- **Single choice tags** - Restricts users to select a single tag within a group. This can be useful to ensure users don't select conflicting tags such as "responsive" and "non-responsive". These will appear as radio buttons.
- **Multiple choice tags** - Allow users to select multiple tags within a group. These will appear as checkboxes.

Tag structure

In addition to the tag types, the structure of how tags are organized in the tag panel can be used to make tagging documents more intuitive. Tags are grouped by sections. Review set search supports the ability to search by tag and by tag section. This means you can create a review set search to retrieve documents tagged with any tag in a section.

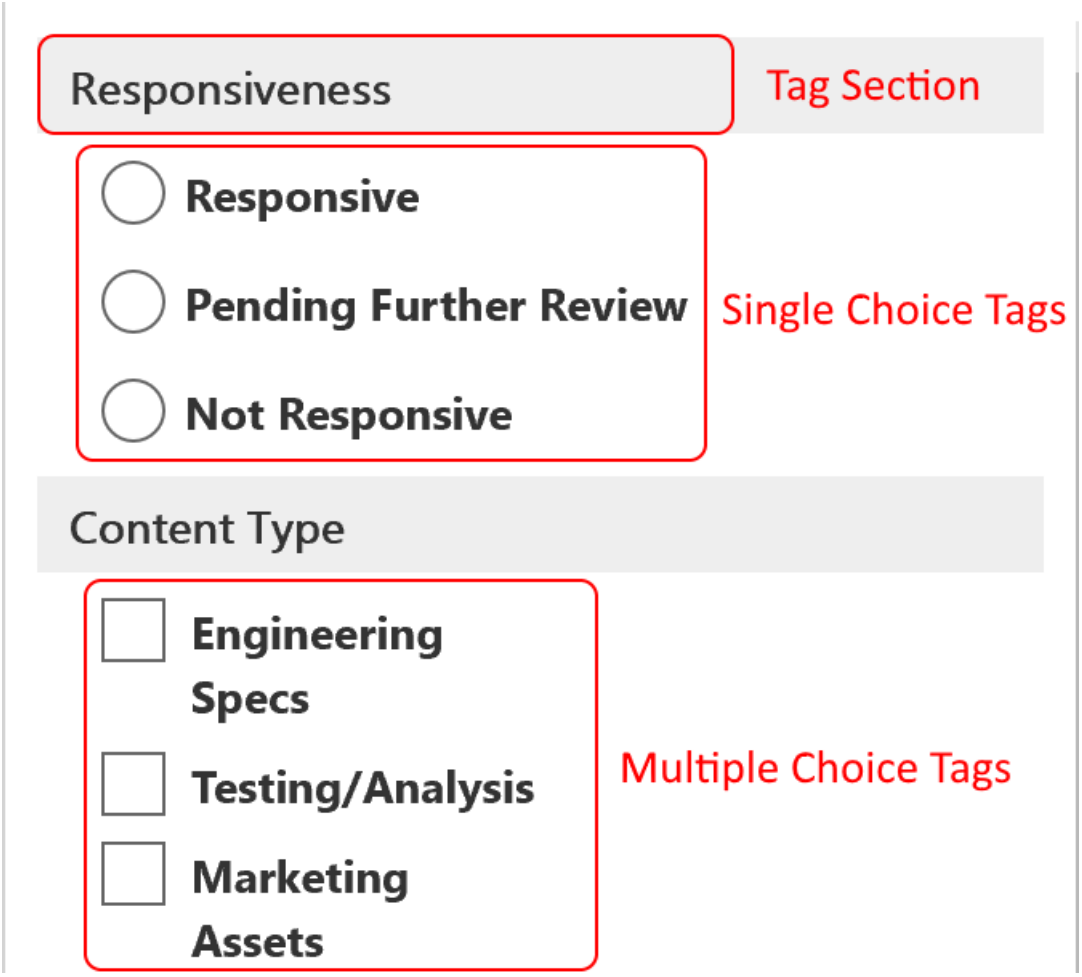

Tags can be further organized by nesting them within a section. For example, if the intent is to identify and tag privileged content, nesting can be used to make it clear that a user can tag a document as "Privileged" and select the type of privilege by checking the appropriate nested tag.

Applying tags

There are several ways to apply a tag to content.

Tagging a single document

When viewing a document in a review set, you can display the tags that a review can use by clicking **Tagging panel**.

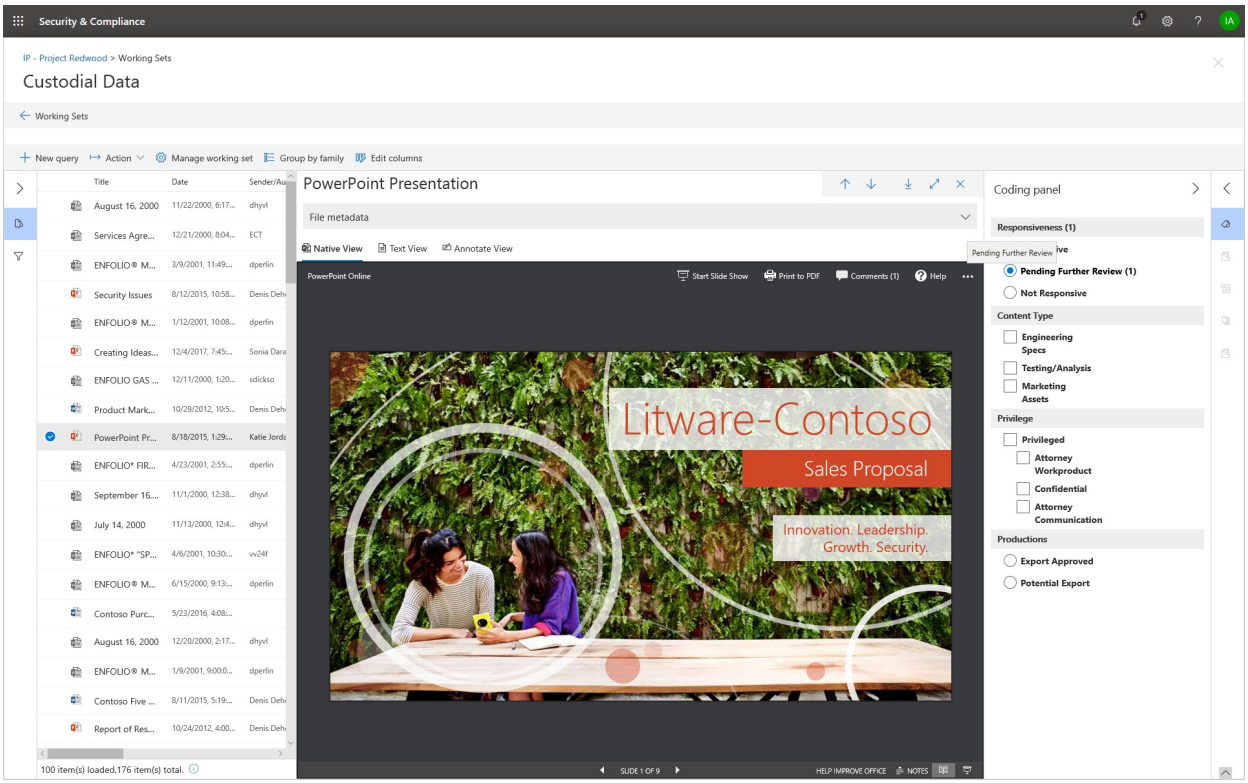

This will enable you to apply tags to the document displayed in the viewer.

Bulk tagging

Bulk tagging can be done by selecting multiple files in the results grid and then using the tags in the **Tagging panel** similar to tagging single documents. Bulk un-tagging can be done by selecting tags twice; the first click will apply the tag, and the second selection will ensure that tag is cleared for all selected files.

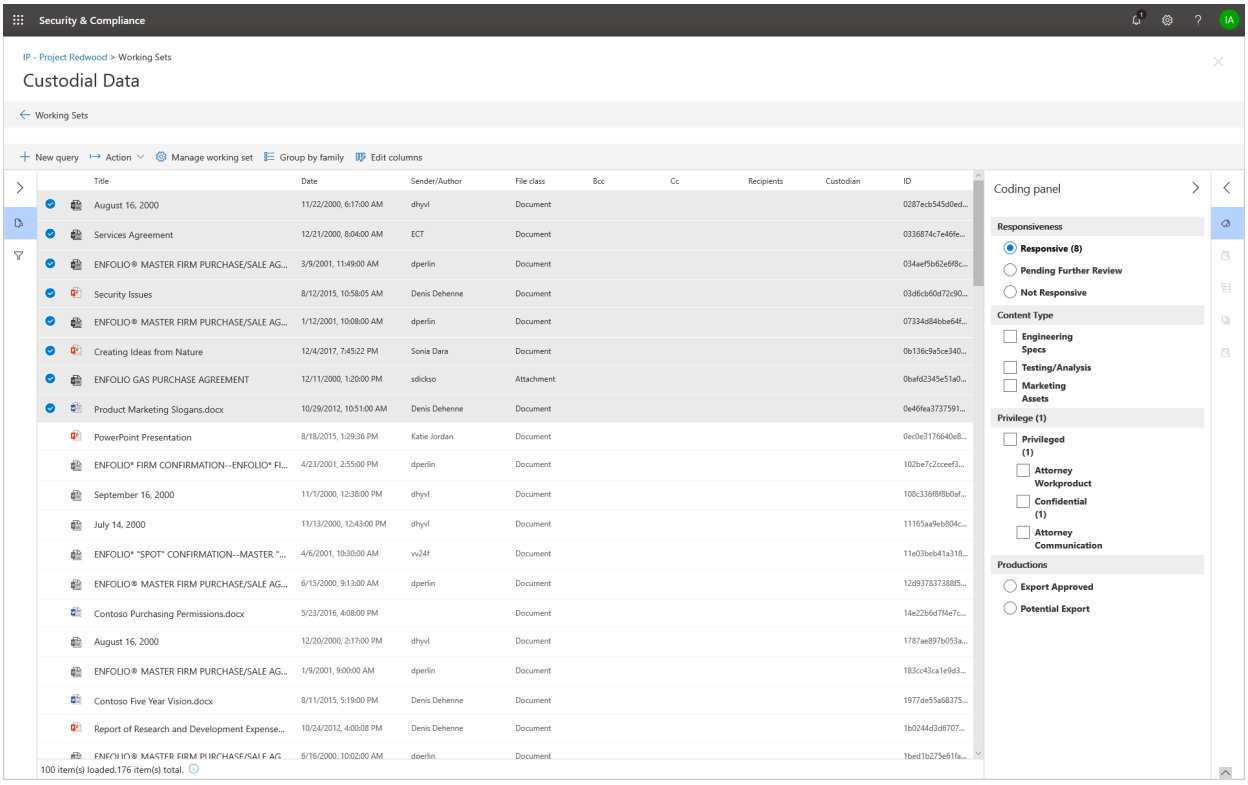

> **NOTE**
>
> When bulk tagging, the tagging panel will display a count of files that are tagged for each tag in the panel.

Tagging in other review panels

When reviewing documents, you can use the other review panels to review other characteristics of documents in the results grid. This includes reviewing other related documents, email threads, near duplicates, and hash duplicates. For example, when you're reviewing related documents (by using the **Document family** review panel), you can significantly reduce review time by bulk tagging related documents. For example, if an email message has several attachments and you want to ensure that the entire family is tagged consistently.

For example, here's how to display the **Tagging panel** when using the **Document family** review panel:

1. With the review panel open for a selected document (for example, displaying the list of related content in the **Document family** review panel, click **Tag documents** under the document family review panel.

   The tagging panel is displayed as a pop-up window.

2. Choose one or more tags to apply the selected document.

3. To tag all documents, select all documents in the **Document family** panel, click **Tag documents**, and then choose the tags to apply to the entire family of documents.

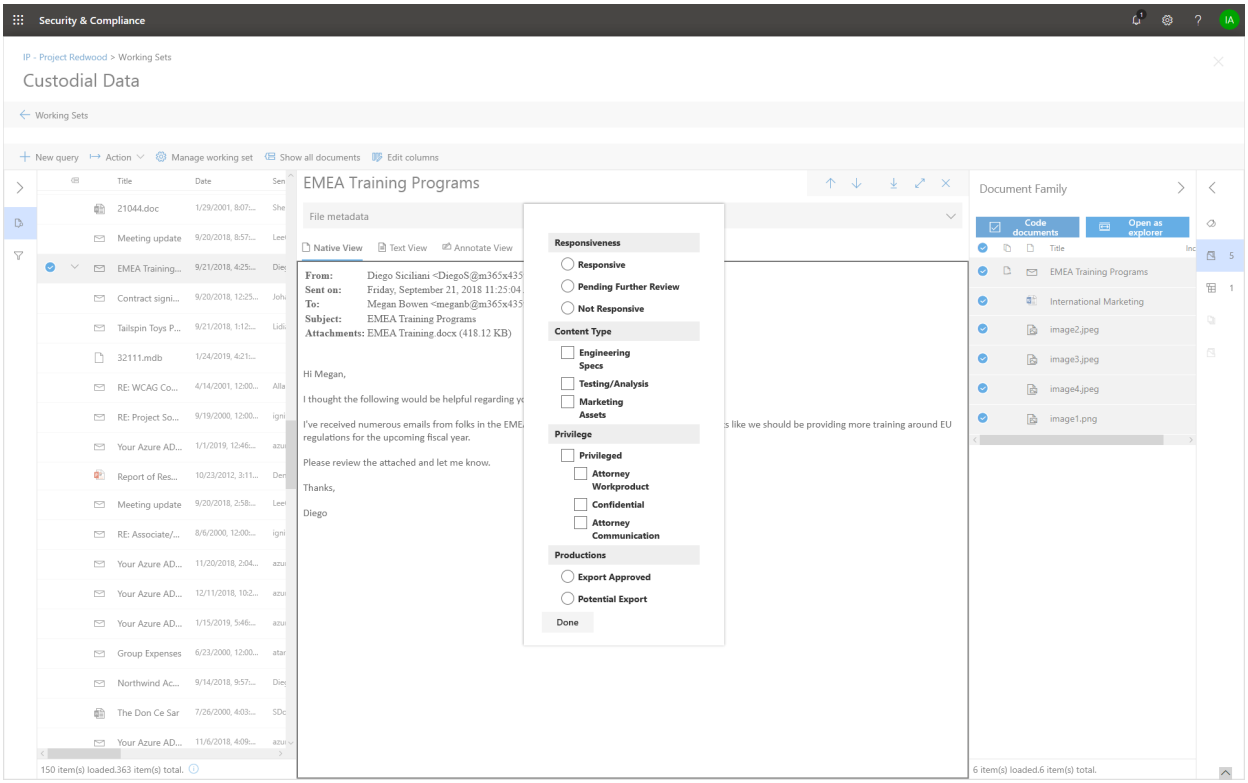

Analyze data in a review set in Advanced eDiscovery

2/18/2021 • 2 minutes to read • Edit Online

When the number of collected documents is large, it can be difficult to review them all. Advanced eDiscovery provides a number of tools to analyze the documents to reduce the volume of documents to be reviewed without any loss in information, and to help you organize the documents in a coherent manner. To learn more about these capabilities, see:

- Near duplicate detection
- Email threading
- Themes

To analyze data in a review set:

1. Configure analytics settings for your case. For more information, see Configure search and analytics settings.
2. Open the review set you want to analyze.
3. Click **Manage review set**.
4. Click **Run analytics for the review set**.

You can check the progress of analysis on the **Jobs** tab of the case.

After analysis is completed, you can view the analytics report, run queries within your review set on outputs of the analysis (see Query within your review set), and see related documents of a given document (see Reviewing data in review set).

Analytics report

To view an analytics report for a review set:

1. Open the review set.
2. Click **Manage review set**.
3. Click **View report**.

The report has seven components from analysis:

- **Target population:** The number of email messages, attachments, and loose documents found in the review set.
- **Documents (excluding attachments):** The number of loose documents that are pivots, unique near duplicates of a pivot, or an exact duplicate of another document.
- **Emails:** The number of email messages that are inclusives, inclusive copies, inclusive minuses, or none of the above.
- **Attachments:** The number of email attachments that are unique or duplicates of another email attachment in the review set.

- **Number of files by type:** The number of files, identified by file extension.
- **Documents by source:** A summary of content by its original data source.
- **Documents aggregated by process:** A summary of content by review set processes.

Near duplicate detection in Advanced eDiscovery

11/2/2020 • 2 minutes to read • Edit Online

Consider a set of documents to be reviewed in which a subset is based on the same template and has mostly the same boilerplate language, with a few differences here and there. If a reviewer could identify this subset, review one of them thoroughly, and review the differences for the rest, they would not have missed any unique information while taking only a fraction of time that would have taken them to read all documents cover to cover. Near duplicate detection groups textually similar documents together to help you make your review process more efficient.

How does it work?

When near duplicate detection is run, the system parses every document with text. Then, it compares every document against each other to determine whether their similarity is greater than the set threshold. If it is, the documents are grouped together. Once all documents have been compared and grouped, a document from each group is marked as the "pivot"; in reviewing your documents, you can review a pivot first and review the other documents in the same near duplicate set, focusing on the difference between the pivot and the document that is in review.

Email threading in Advanced eDiscovery

11/2/2020 • 2 minutes to read • Edit Online

Consider an email conversation that has been going on for a while. In most cases, the last email on the thread will include the contents of all the preceding emails; reviewing the last email will give a complete context of the conversation that happened in the thread. Email threading identifies such emails so that reviewers can review a fraction of collected documents without losing any context.

What does email threading do?

Email threading parses each email and deconstructs it to individual messages; each email is a chain of individual messages. Then, it analyzes all emails in the review set to determine whether an email has unique content or if the chain is wholly contained in a different email. In the end emails are divided into four categories:

- **Inclusive**: the last message in the email has unique content, and the email has all of the attachments that were included in other emails of which the content is wholly contained in this email.
- **Inclusive minus**: the last message in the email has unique content, but the email does not contain some of the attachments that were included in other emails of which the content is wholly contained in this email.
- **Inclusive copy**: an exact copy of an inclusive/inclusive minus email
- **None**: The content of this email is wholly contained in at least one email that is marked as inclusive/inclusive minus.

How is it different from conversations in Outlook?

At a glance, this sounds similar to conversation groupings in Outlook. However, there are some important distinctions. Consider an email conversation that got forked into two conversations; for instance, someone responded to an email that is not the latest in the conversation so the last two emails in the conversation both have unique content.

Outlook would still group the emails into a single conversation; reading only the last email would mean missing the context of the second-to-last email, which also contains unique content. Because email threading parses out each email into individual components and compares them, email threading would mark both of the last two emails as inclusive, ensuring that you won't miss any context as long as you read all emails marked as inclusive.

Themes in Advanced eDiscovery

11/2/2020 • 2 minutes to read • Edit Online

How does a person write a document? They generally start with one or more ideas they want to convey in the document, and compose using words that align with the ideas. The more prevalent an idea is, the more frequent the words that are related to that idea tend to be. This informs how people consume documents as well. The important thing to understand from reading a document is the ideas that the document is trying to convey, which ideas appear where, and what the relationships between the ideas are.

This can be extended to how a person wants to consume a set of documents. They want to see which ideas are present in the sets, and which documents are talking about those ideas. Also, if they find a particular document of interest, they want to be able to see documents that discuss similar ideas.

The Themes functionality in Advanced eDiscovery attempts to mimic how humans reason about documents, by analyzing the *themes* that are discussed in a review set and assigning a theme to documents in the review set. In Advanced eDiscovery, Themes goes one step further and identifies the *dominant theme* in each document. The dominant theme is the one that appears the most often in a document.

How does Themes work?

The Themes functionality analyzes documents with text in a review set to parse out common themes that appear across all the documents in the review set. Advanced eDiscovery assigns those themes to the documents in which they appear. It also labels each theme with the words used in the documents that are representative of the theme. Because a document can contain various types of subject matter, Advanced eDiscovery often assigns multiple themes to documents. The theme that appears most prominently in a document is designated as its dominant theme.

Use the Relevance module to analyze data in Advanced eDiscovery

11/2/2020 • 6 minutes to read • Edit Online

In Advanced eDiscovery, the Relevance module includes the Relevance training and review of files related to a case. In order to use the Relevance workflow, go to Manage review set within a review set and click on Show Relevance. There are a couple of steps you need to complete before you can start the workflow:

- Process: each load set added to the review set will show up as a "container" here. You need to process these documents before you can add them to Relevance module; this is also where you can mark them as seed or pre-tagged for a specific issue.
- Add to Relevance: Under Relevance > Loads, you can add documents that have been processed to Relevance to make them available for training.

The Relevance workflow is shown and described as follows:

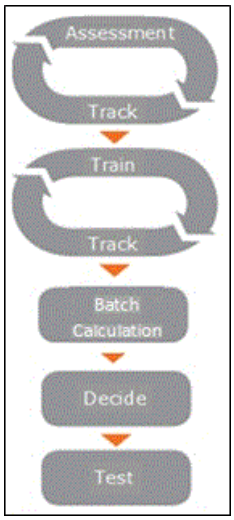

- **Cycles of assessment and tracking**:
  - **Assessment**: Enables early assessment based on a random sample of files and uses this assessment to apply decisions to determine the performance of the predictive coding process.
  - **Track**: Calculate and display interim results of the assessment while monitoring statistical validity of the process.
- **Cycles of training and tracking**
  - **Tag**: Advanced eDiscovery learns Relevance criteria specific to each issue based on the expert's iterative review and tagging of individual files.
  - **Track**: Calculate and display interim results of the Relevance training while monitoring statistical validity of the process.
- **Batch calculation**: The accumulated and learned Relevance criteria is applied to the entire file collection, and a Relevance score is generated for each file.
- **Decide**: The results of the analysis applied to the entire case is displayed after Batch calculation, and data used to make document review decisions is displayed.

- **Test**: Results can be tested to verify the validity and effectiveness of the Advanced eDiscovery processing.
- **Search**: Once the Relevance workflow is complete, you can use the output such as read percentile of a document for your issue when you run a query within your review set.

Guidelines for Relevance training and review

Following is an overview of guidelines for Relevance training and review:

- **Errors and inconsistencies**: If tagging errors are made during training, return to previous file samples to correct them. If there are too many errors to correct or there is a new perspective of the case or issue, the Relevance criteria should be redefined by the Administrator, and the Relevance training restarted.
- **Tagging and training**:
  - Files should be tagged based on content only. Do not consider metadata, such as custodian, date, or file path.
  - Do not consider date range indications in the text when tagging files.
  - Do not consider embedded graphical images when tagging files.
  - Ignore text applied to Relevance will be removed in the displayed file content in the text view in Relevance. If the values for Ignore text were defined after Relevance training already started, the new ignored text will be applied to sample files created from the point in which it was defined. The Ignore Text feature should be used cautiously, as its use may reduce the performance of file analysis
  - Use the **Skip tagging** option only when necessary. Advanced eDiscovery does not train based on skipped files. In assessment, if it's hard to tell whether a file is relevant, it is better to tag as Relevant (R) or Not relevant (NR) whenever possible rather than selecting **Skip**. When Advanced eDiscovery evaluates training, it can then be seen how well these types of files were processed.
  - Even files with a very small amount of extracted text should be tagged in training as R/NR, rather than as "Skip", when possible.
  - Tagging can impact the classifier as long as the file is readable and can be tagged as R/NR.
  - The file sequence number on the displayed Sample files list on the **Tag** tab allows the user to return to the original displayed order of the files.
  - You can go back to any sample and change the tagging of the assessment and training set files. The changes will be applied when creating the next sample.
  - Scanned Excel files in PDF format should be treated the same as native Excel files when tagging files.
  - When in doubt regarding the Relevance tagging of a file, consult an expert. Incorrect tagging during the Relevance training can lead to lost time later in the process and may also have a negative impact on the quality of the overall results.
  - Keywords that were defined in Keyword lists will be displayed in colors to help the user identify relevant files while tagging.
- **Batch calculation**: Files that were tagged as R/NR by the expert will receive a score of either 0 or 100. This applies to tagging made before Batch calculation. If the expert switched the issue to Idle after Batch Calculation and continued tagging this issue, the newly tagged scores will not be 100/0 but rather the original score.
- **Issues and sampling mode**: Issues are usually turned Off when work on them is completed (Relevance

training is stabilized and Batch calculation was performed), when the issues are canceled, or when another user is working on the issues.

Steps in Relevance training

In the **Relevance** > **Track** tab, Advanced eDiscovery provides recommendations on how to proceed in the processing, with the following next steps. The implications are described below when each of the following steps is recommended in the Relevance training process.

- Tagging / Continue tagging: File review and Relevance tagging performed by an expert for each file and issue within a sample.
    - Implication: An existing sample needs to be tagged.
- Assessment / Continue assessment: Enables early validation of case issue relevance and a preliminary view of the relevance of the file population imported for the current case.
    - Implication: More assessment is required or recommended.
- Training / Continue training: Process during which Advanced eDiscovery learns from the expert who is tagging the file samples and acquires the ability to identify Relevance criteria pertinent to each issue within the context of each case.
    - Implication: The issue needs more training; the next sample should be created and tagged.
- Batch calculation: Relevance process in which Advanced eDiscovery takes the knowledge acquired during the training stage and applies it to the entire file population. All files in the pertinent file group are assessed for relevance and assigned a Relevance score.
    - Implication: The issue has stabilized, and Batch calculation can be performed.
- Catch-up: Relevance indicates when an expert reviews and tags a sample of files selected from an additional file load during a Rolling Loads scenario.
    - Implication: A new load has been added, and Catch-up is required to continue working.
- Tag inconsistencies: Process identifies, via an Advanced eDiscovery algorithm, inconsistencies in the file tagging process that may negatively impact the analysis.
    - Implication: The next sample will include files that have been tagged in previous samples, and their tagging must be redone.
- Update classifier: Allows the user to apply tagging or seeding changes.
    - Implication: Tagging and seeding changes can be applied without needing to manually run another Relevance sample.
- On hold: The Relevance training process is completed.
    - Implication: No Relevance training is required at this point.

Although Advanced eDiscovery guides you through the process, with recommended Next steps at different stages, it also allows you to navigate between tabs and pages, and to make choices to address situations that may be pertinent to your individual case, issue, or document review process.

It is possible to accept or override Advanced eDiscovery Next step processing choices. If you want to perform a step other than the recommended Next step, click the **Next step** listed in the expanded issue display in the dialog, click the **Modify** button next to the Next step, and select another Next step option.

> **NOTE**
> Some options may remain disabled after unlocking as they are not supported for use at that point in the process.

More information

Understanding Assessment in Relevance

Tagging and Assessment

Tagging and Relevance training

Tracking Relevance analysis

Deciding based on the results

Testing Relevance analysis

Query the data in a review set

Retirement of the Relevance module in Advanced eDiscovery

2/18/2021 • 2 minutes to read • Edit Online

On March 10, 2021, we are retiring the Relevance module in Advanced eDiscovery. This retirement means that organizations will no longer have access to the Relevance module (by going to **Manage review set** > **Relevance** in an Advanced eDiscovery case) or be able to access any existing Relevance models. The current Relevance module that is being retired will be replaced with a new predictive coding solution in Q2 CY 2021. This new functionality will let organizations build their own predictive coding models in an easier and more intuitive workflow.

To prepare for this upcoming retirement, we recommend that organizations who use the Relevance module export their model's output before the retirement date by running a Batch calculation for all existing models. All Relevance scores from your model will be permanently stored in the corresponding review set and accessible when documents are exported. Relevance scores are also retained as metadata in the load file. Also, you will still be able to filter content in the review set based on relevance score and have access to all metadata produced by your Relevance models.

Complete unfinished models

For any unfinished Relevance models, please complete assessment, training, and Batch calculation so that you can apply the model to the documents in a review set. Completing the Batch calculation will preserve the information after the retirement date of the Relevance module.

Here are the steps to complete any unfinished models:

1. Train your model until it is stabilized and ready for Batch calculation. See Tagging and Relevance training.

   The following screenshot shows a module that is ready for a Batch calculation. Notice that the Assessment and Training is complete, and the next step is to run Batch calculation.

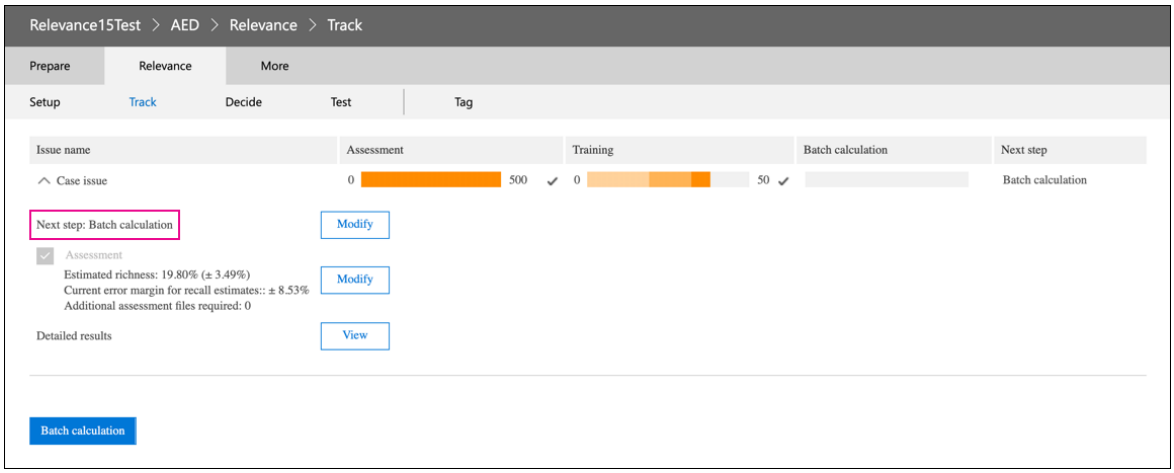

2. Run the Batch calculation. See Performing Batch calculation.

3. Verify that Batch calculation was successful. See Batch calculation results.

For help with completing unfinished Relevance models, contact Microsoft Support.

Export case data in Advanced eDiscovery

11/2/2020 • 2 minutes to read • Edit Online

There are three ways to export data from a review set:

Download: Download (by using a browser) a small set of native files. This is the quickest way to export a small set of data. This method preserves the native file names. For more information, see Download documents from a review set.

Export: Customize what data is exported. This includes exporting file metadata, native files, text files, and redacted documents that have been saved to a PDF file. Once the exported data is prepared, you can download the data from the **Export** tab to a local computer. For more information, see Export documents from a review set

Add to another review set: Copy data from one review set to a different review set. For more information, see Add data from one review set to another review set.

> **NOTE**
> In Microsoft 365, data is hashed and those hashes are provided in the output file for verification purposes. This is supplemented by audit logs and reporting functionality, such as collection statistics, load set reports, and export reports (including the export load file).

Download selected documents from a review set

11/2/2020 • 2 minutes to read • Edit Online

Download offers a simple way to download content from a review set in native format. The download tool in Advanced eDiscovery leverages the browser's data transfer features. A browser prompt will appear when a download is ready. Files downloaded using this method are zipped in a container file and will contain item-level files. This means that if you select to download an attachment, you will receive the email message with the attachment included. Similarly, if you export an Excel spreadsheet that is embedded in a Word document, the Word document with the embedded Excel spreadsheet are included in the download. When you downloaded items, the Last Modified Data property is preserved and can be viewed as a file property.

To download content from a review set in an Advanced eDiscovery case, start by selecting the files you want to download then select **Action items** > **Download**.

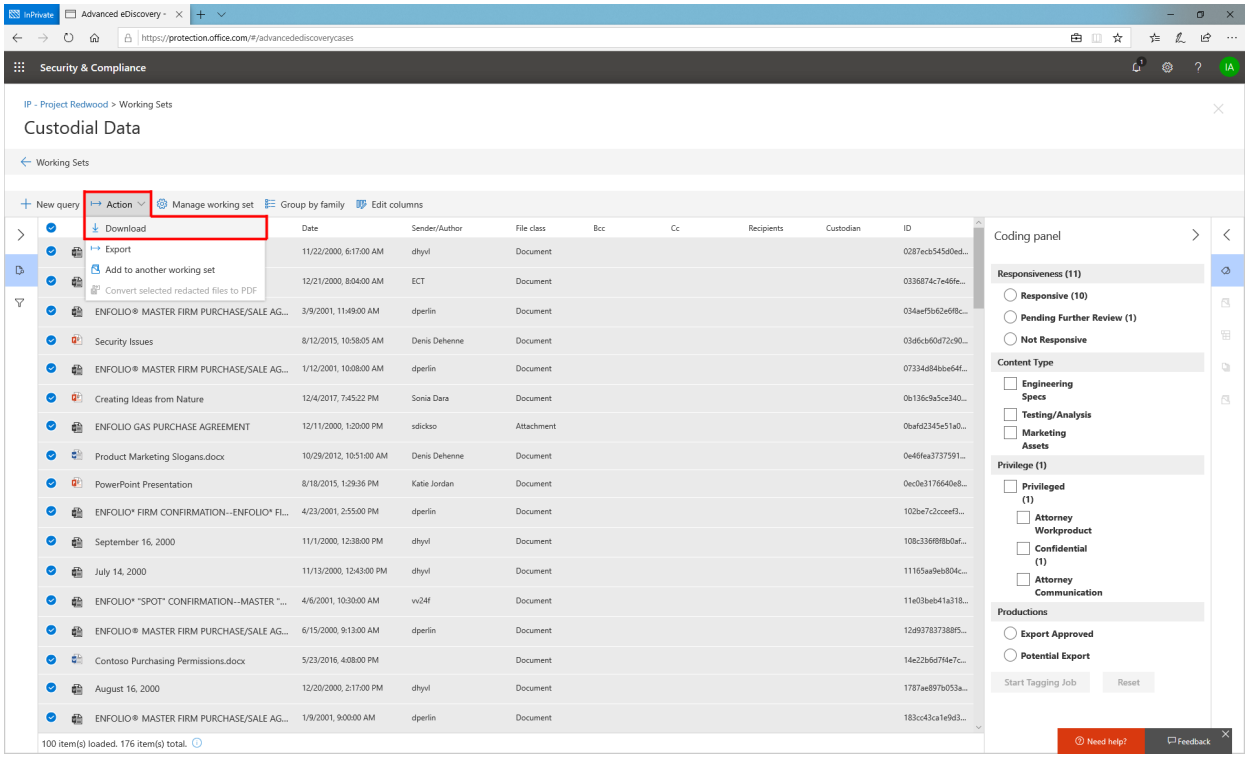

Export documents from a review set in Advanced eDiscovery

11/2/2020 • 3 minutes to read • Edit Online

Export allows users to customize the content that is included in the download package. The Export tool provides a configuration page with the following settings:

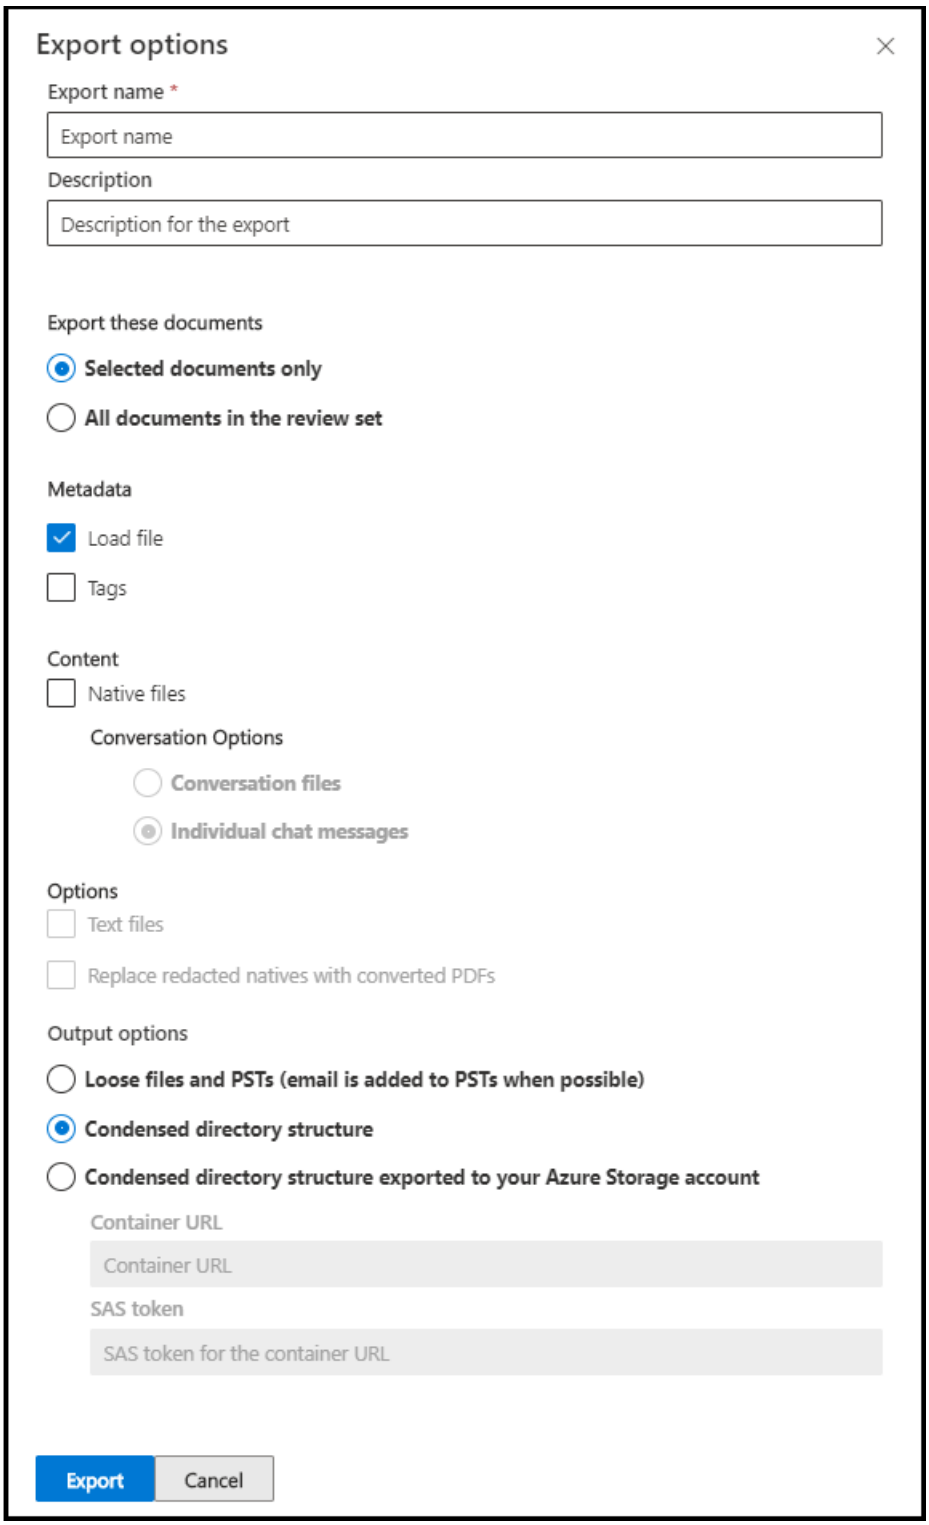

Export options

- Export name: Name of the export job.
- Description: Free-text field for you to add a description.
- Export these documents:
  - Selected documents only - Exports only the documents that are currently selected.
  - All documents in the review set - Exports all documents in the review set
- Metadata
  - Load file - This file contains metadata for each file. see Document metadata fields in Advanced eDiscovery for more information about what fields are included. This file can typically be ingested by third-party eDiscovery tools.
  - Tags - When selected, tagging information will be included in the load file.
- Content
  - Native files - Select this checkbox to include the native files.
  - Conversation options
    - Conversation files - Export reconstructed chat messages. This format presents conversations in a form that resembles what users see in the native application.
    - Individual chat messages - Export the original conversation files as they are stored in Microsoft 365.
- Options
  - Text files - Include extracted text versions of native files.
  - Replace redacted natives with converted PDFs - If redacted PDF files are generated during review, these files are available for export. You can choose to export only the native files that were redacted (by not selecting this option) or you can select this option to export the PDF files that contain the actual redactions.
- Output options (Exported content is either available for download directly through a web browser or can be sent to an Azure Storage account. The first two options enable direct download.)
  - Loose files and PSTs (email is added to PSTs when possible) - Files are exported in a format that resembles the original directory structure seen by users in their native applications. For more information, see the Loose files and PST export structure section.
  - Condensed directory structure - Files are exported and included in the download.
  - Condensed directory structure exported to your Azure Storage account - Files are exported to your organization's Azure Storage accouunt.

Loose files and PST export structure

If you select this export option, the exported content is organized in the following structure:

- Root folder – This folder in named ExportName.zip
  - Export_load_file.csv - Metadata file.
  - Summary.csv - A summary file that also contains export statistics.
  - Exchange - This folder contains all content from Exchange in native file format. Natives files are

replaced with redacted PDFs if you selected the **Replace redacted natives with converted PDFs** option.

- SharePoint = This folder contains all native content from SharePoint in a native file format. Natives files are replaced with redacted PDFs if you selected the **Replace redacted natives with converted PDFs** option.

Condensed directory structure

- Root folder - This folder is named ExportName.zip
  - Export_load_file.csv - Metadata file.
  - Summary.txt - A summary file that also contains export statistics.
  - Input_or_native_files - This folder contains all the native files that were exported. If you export redacted PDF files, they are not put in PST files. Instead, they're added to a separated folder.
  - Error_files - This folder contains the following error files, if they are included in the export:
    - ExtractionError. A CSV file that contains any available metadata of files that weren't properly extracted from parent files.
    - ProcessingError – This file contains a list of documents with processing errors. This content is item-level, meaning if an attachment resulted in a processing error, the email message that contains the attachment is included in this folder.
  - Extracted_text_files - This folder contains all of the extracted text files that were generated at processing.

> **NOTE**
> Export jobs are retained for the life of the case and can be downloaded as long as the case isn't deleted.

Manage jobs in Advanced eDiscovery

2/18/2021 • 5 minutes to read • Edit Online

Here's a list of the jobs (which are typically long-running processes) that are tracked on the **Jobs** tab of a case in Advanced eDiscovery. These jobs are triggered by user actions when using and managing cases.

JOB TYPE	DESCRIPTION
Adding data to a review set	A user adds the results of a search to a review set. This job consists of two sub jobs: • **GatheringItems** - A list of items that match the search query (and the Microsoft 365 data source that they're located in) is generated. • **Ingestion & Indexing** - The items that match the search query are copied to an Azure Storage location (in a process called *ingestion*) and then those items in the Azure Storage location are reindexed. This new index is used when querying and analyzing items in the data set. For more information, see Add search results to a review set.
Adding data to another review set	A user adds documents from one review set to a different review set in the same case. For more information, see Add data to a review set from another review set.
Adding non-Microsoft 365 data to a review set	A user uploads non-Microsoft 365 data to a review set. The data is also indexed during this process. For example, files from an on-premises file server or a client computer are uploaded to a review set. For more information, see Load non-Microsoft 365 data into a review set.
Adding remediated data to a review set	Data with processing errors is remediated and loaded back into a review set. For more information, see: • Error remediation when processing data • Single item error remediation
Comparing load sets	A user looks at the differences between different load sets in a review set. A load set is an instance of adding data to a review set. For example, if you add the results of two different searches to the same review set, each would represent a load set.
Conversation reconstruction	When a user adds the results of a search to a conversation review set, instant message conversations (also called *threaded conversations*) in services like Microsoft Teams are reconstructed in a PDF file. This job is also triggered when a user clicks **Action > Create conversation PDFs** in a review set. For more information, see Review conversations in Advanced eDiscovery.
Converting redacted documents to PDF	After a user annotates a document in a review set and redacts a portion of it, they can choose to convert the redacted document to a PDF file. This ensures that the redacted portion will not be visible if the document is exported for presentation. For more information, see View documents in a review set.

JOB TYPE	DESCRIPTION
Estimating search results	After a user creates and runs a new search (or reruns an existing search) the search tool searches the index for items that match the search query and prepares an estimate that includes the number and total size of all items by the search, and the number of data sources searched. For more information, see Collect data for a case.
Preparing data for export	A user exports documents from a review set. When the export process is complete, they can download the exported data to a local computer. For more information, see Export case data.
Preparing for error resolution	When a user selects a file and creates a new error remediation in the Error view on the **Processing** tab of a case, the first step in the process is to upload the file that has the processing error to an Azure Storage location in the Microsoft cloud. This job tracks the progress of the upload process. For more information about the error remediation workflow, see Error remediation when processing data.
Preparing search preview	After a user creates and runs a new search (or reruns an existing search), the search tool prepares a sample subset of items (that match the search query) that can be previewed. Previewing search results help you determine the effectiveness of the search. For more information, see Collect data for a case.
Re-indexing custodian data	When you add a custodian to a case, all partially indexed items in the custodian's selected data sources are reindexed by a process called *Advanced indexing*. This job is also triggered when you click **Update index** on the **Processing** tab of a case, and when you update the index for a specific custodian on the custodian properties flyout page. For more information, see Advanced indexing of custodian data.
Running analytics	A user analyzes data in a review set by running Advanced eDiscovery analytics tools such as near duplicate detection, email threading analysis, and themes analysis. For more information, see Analyze data in a review set.
Tagging documents	This job is triggered when a user clicks **Start tagging job** in the **Tagging panel** when reviewing documents in a review set. A user can start this job after tagging documents in a review set and then bulk-selecting them in the view document panel. For more information, see Tag documents in a review set.

Job status

The following table describes the different status states for jobs.

STATUS	DESCRIPTION

STATUS	DESCRIPTION
Submitted	A new job was created. The date and time that the job was submitted is displayed in the **Created** column on the **Jobs** tab.
Submission failed	The job submission failed. You should attempt to rerun the action that triggered the job.
In progress	The job is in progress, you can monitor the progress of the job in the **Jobs** tab.
Successful	The job was successfully completed. The date and time that the job completed is displayed in the **Completed** column on the **Jobs** tab.
Partially successful	The job was partially successful. This status is typically returned when the job didn't find any partially indexed data (also called *unindexed data*) in some of the custodian data sources.
Failed	The job failed. You should attempt to rerun the action that triggered the job. If the job fails a second time, we recommend that you contact Microsoft Support and provide the support information from the job.

Close or delete an Advanced eDiscovery case

11/2/2020 • 2 minutes to read • Edit Online

When the legal case or investigation supported by an Advanced eDiscovery case is completed, you can close or delete a case. You can also reopen a closed case.

Close a case

Here's what happens when you close an Advanced eDiscovery case:

- If the case contains any content locations on hold, those holds will be turned off. After the hold is turned off, a 30-day grace period (called a *delay hold*) is applied to content locations that were on hold. This helps prevent content from being immediately deleted and gives admins an opportunity to search for or recover content that will be permanently deleted after the delay hold period expires. For more information, see Removing content locations from an eDiscovery hold.
- Closing a case only turns off the holds that are associated with that case. If other holds are place on a content location (such as a Litigation Hold, Core eDiscovery hold, or a hold from a different Advanced eDiscovery case) those holds will still be maintained.
- The case is still listed on the eDiscovery page in the Microsoft 365 compliance center. The details, holds, searches, and members of a closed case are retained.
- You can edit a case after it's closed. For example, you can add or removing members, create searches, export search results, and prepare search results for analysis in Advanced eDiscovery. The primary difference between active and closed cases is that holds are turned off when a case is closed.

To close a case:

1. On the **Advanced eDiscovery** page, select the case that you want to close.
2. On the **Settings** tab, under **Case Information**, click **Select**.
3. Click **Close case**.

   It might take up to 60 minutes for the closing process to complete.

Reopen a closed case

When you reopen an Advanced eDiscovery case, any holds that were in place when the case was closed won't be automatically reinstated. After the case is reopened, you'll have to go to the **Holds** tab and turn on the previous holds. To turn on a hold, select it to display the flyout page, and then set the **Status** toggle to **On**.

To reopen a closed case:

1. On the **Advanced eDiscovery** page, select the case that you want to reopen.
2. On the **Settings** tab, under **Case Information**, click **Select**.
3. Click **Reopen case**.

   It might take up to 60 minutes for the reopening process to complete.

Delete a case

You can delete both active and closed Advanced eDiscovery cases. When you delete a case, all components associated with the case, such as the list of custodians, communications, searches, review sets, and export job are deleted. The case is removed from the list of cases on the **Advanced eDiscovery** page in the Microsoft 365 compliance center. You can't recover or reopen a deleted case.

> **NOTE**
>
> In data spillage scenarios, the only way to remove items in a review set is to delete the Advanced eDiscovery case. Other "search and purge" methods don't remove items from a review set.

Before you can delete a case (whether it's active or closed), you must first delete *all* holds associated with the case. That includes deleting holds with a status of **Off**.

To delete holds associated with a case:

1. Go the **Holds** tab in the Advanced eDiscovery case that you want to delete.
2. Click the hold that you want to delete.
3. On the flyout page, click **Delete hold**.

To delete a case:

1. On the **Advanced eDiscovery** page, select the case that you want to delete.
2. On the **Settings** tab, under **Case Information**, click **Select**.
3. Click **Delete case**.

Add or remove members from a case

11/2/2020 • 2 minutes to read • Edit Online

You can add or remove members to manage who can access the case. However, before a member can access a Advanced eDiscovery case (and perform tasks in the case), you must add the user to the eDiscovery Manager role group on the **Permissions** page in the security and compliance center. For more information, see Assign eDiscovery permissions in the Security & Compliance Center.

1. On the **Advanced eDiscovery** page, go to the case that you want to add a member to.
2. Click the **Settings** tab and then click **Select** in the **Access & permissions** tile.
3. Click **Update**.
4. Under **Manage members**, click **Add** to add members to the case. You can also choose to add a role group to the case by clicking **Add** under **Manage role groups**.
5. In the list of people or role groups that can be added as members of the case, select the check box next to the names of the people or role groups that you want to add.
6. After you've selected the people or role groups to add as members of the case, click **Add**.
7. In the **Manage this case** flyout page, click **Save** to save the new list of case members.

Configure search and analytics settings in Advanced eDiscovery

2/18/2021 • 3 minutes to read • Edit Online

You can configure settings for each Advanced eDiscovery case to control the following functionality.

- Near duplicates and email threading
- Themes
- Autogenerated review set query
- Ignore text
- Optical character recognition

To configure search and analytics settings for a case:

1. On the **Advanced eDiscovery** page, select the case.
2. On the **Settings** tab, under **Search & analytics**, click **Select**.

   The case settings page is displayed. These settings are applied to all review sets in a case.

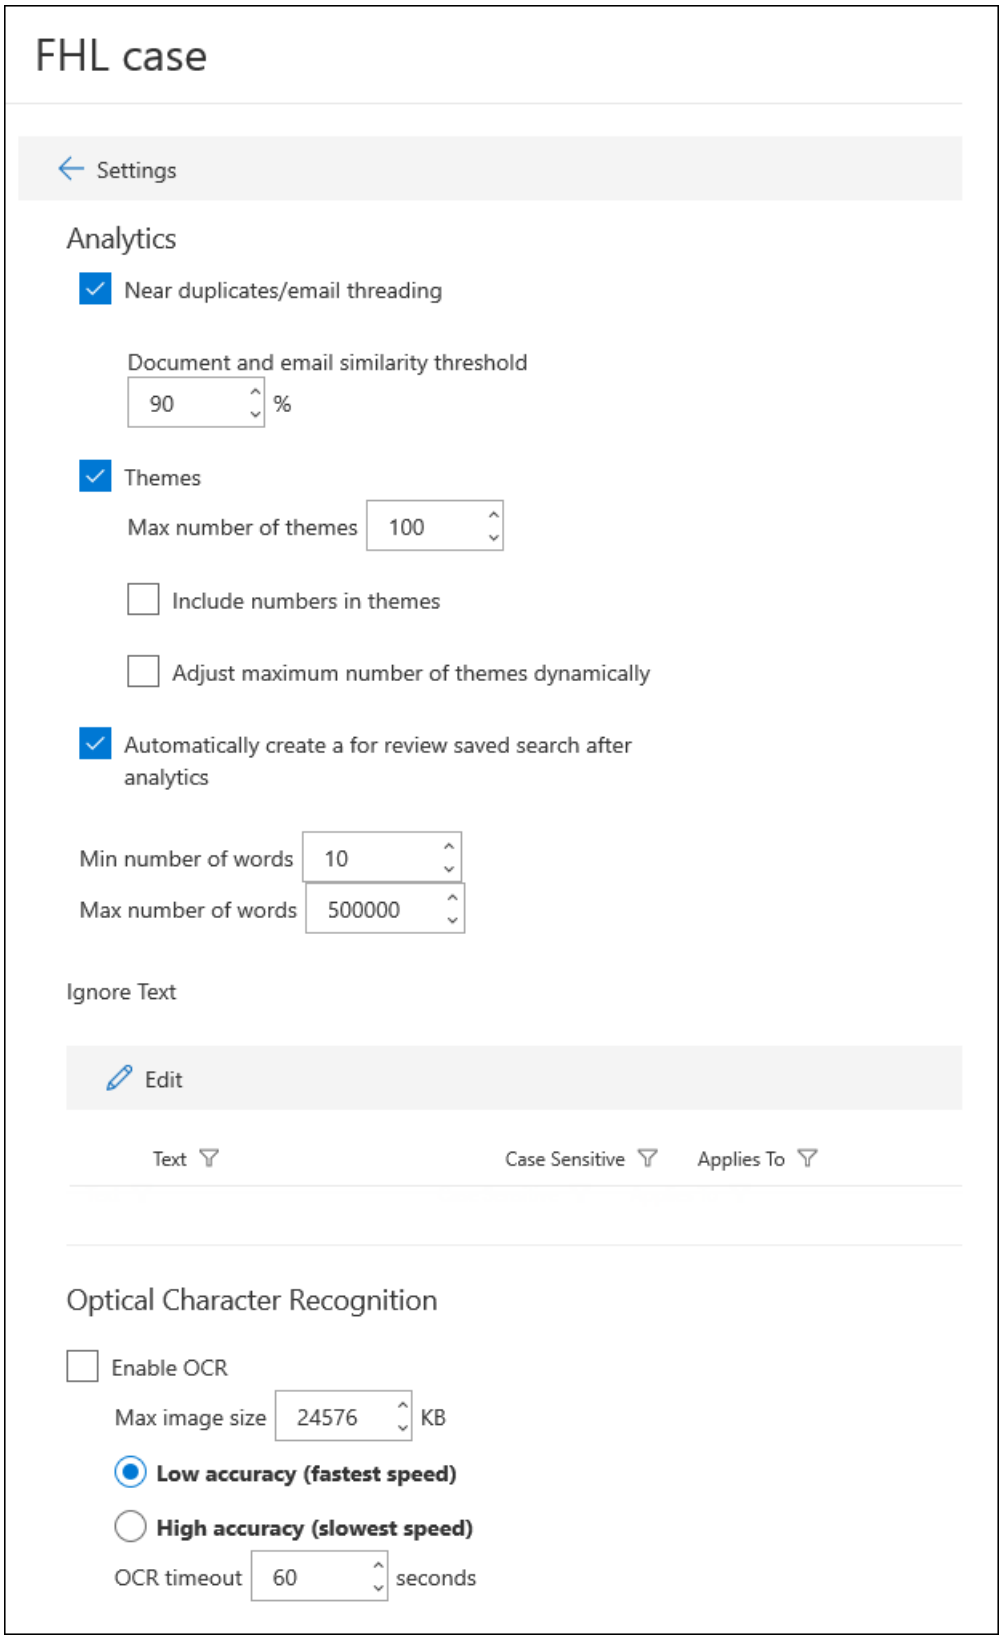

Near duplicates and email threading

In this section, you can set parameters for duplicate detection, near duplicate detection, and email threading. For more information, see Near duplicate detection and Email threading.

- **Near duplicates/email threading:** When turned on, duplicate detection, near duplicate detection, and email threading are included as part of the workflow when you run analytics on the data in a review set.
- **Document and email similarity threshold:** If the similarity level for two documents is above the threshold, both documents are put in the same near duplicate set.

- **Minimum/maximum number of words:** These settings specify that near duplicates and email threading analysis are performed only on documents that have at least the minimum number of words and at most the maximum number of words.

Themes

In this section, you can set parameters for themes. For more information, see Themes.

- **Themes:** When turned on, themes clustering is performed as part of the workflow when you run analytics on the data in a review set.
- **Maximum number of themes:** Specifies the maximum number of themes that can be generated when you run analytics on the data in a review set.
- **Include numbers in themes:** When turned on, numbers (that identify a theme) are included when generating themes.
- **Adjust maximum number of themes dynamically:** In certain situations, there may not be enough documents in a review set to produce the desired number of themes. When this setting is enabled, Advanced eDiscovery adjusts the maximum number of themes dynamically rather than attempting to enforce the maximum number of themes.

Review set query

If you select the **Automatically create a For Review saved search after analytics** checkbox, Advanced eDiscovery autogenerates review set query named **For Review.**

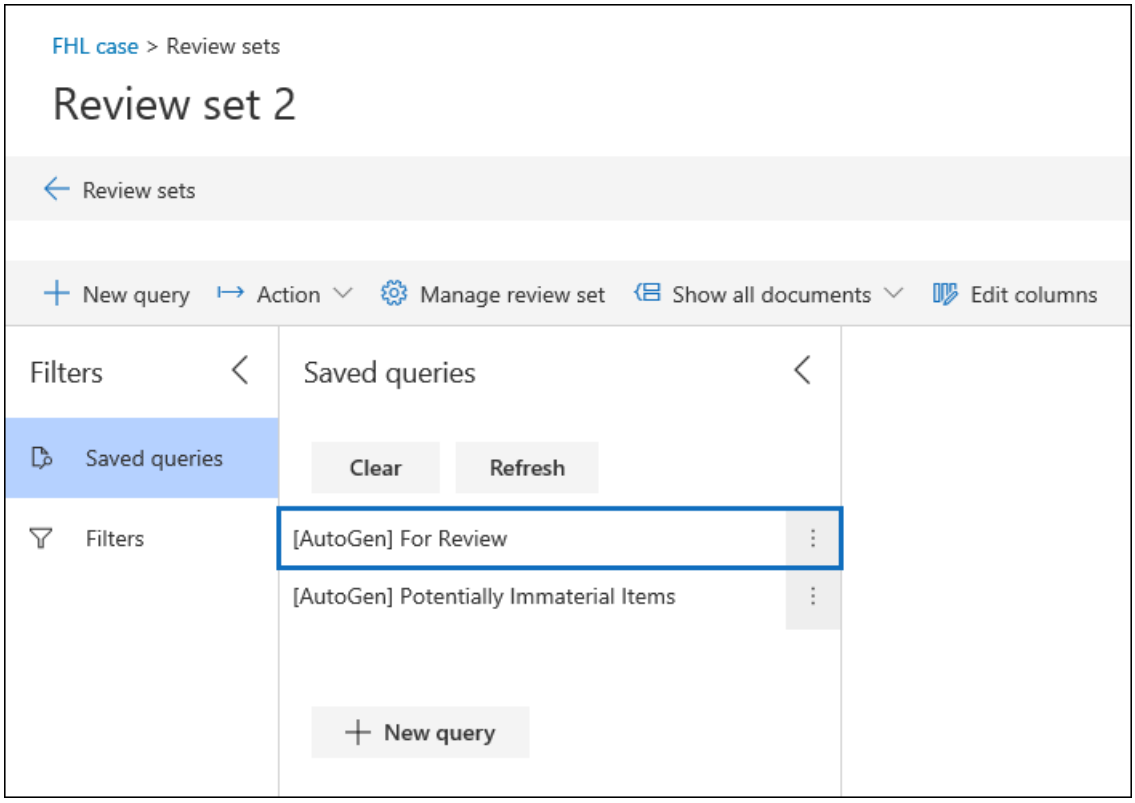

This query basically filters out duplicate items from the review set. This lets you review the unique items in the review set. This query is created only when you run analytics for a review set in the case. For more information, about review set queries, see Query the data in a review set.

Ignore text

There are situations where certain text will diminish the quality of analytics, such as lengthy disclaimers that get

added to email messages regardless of the content of the email. If you know of text that should be ignored, you can exclude it from analytics by specifying the text string and the analytics functionality (Near-duplicates, Email threading, Themes, and Relevance) that the text should be excluded for. Using regular expressions (RegEx) as ignored text is also supported.

Optical character recognition (OCR)

When this setting is turned on, OCR processing will be run on image files. OCR processing is run in the following situations:

- When custodians and non-custodial data sources are added to a case. OCR processing is performed during the Advanced indexing process. This means that text in image files that matches the search criteria will be returned in a collection search.
- When content from other data sources (that aren't associated with a custodian and added to the case in a non-custodial data source) is added to a review set.

After data is added to a review set, image text can be reviewed, searched, tagged, and analyzed. You can view the extracted text in the Text viewer of the selected image file in the review set. For more information, see:

- Advanced indexing of custodian data
- Add search results to a review set
- Supported image file types

Advanced eDiscovery alignment with the Electronic Discovery Reference Model

2/18/2021 • 2 minutes to read • Edit Online

The built-in workflow of Advanced eDiscovery in Microsoft 365 aligns with the eDiscovery process outlined by the Electronic Discovery Reference Model (EDRM).

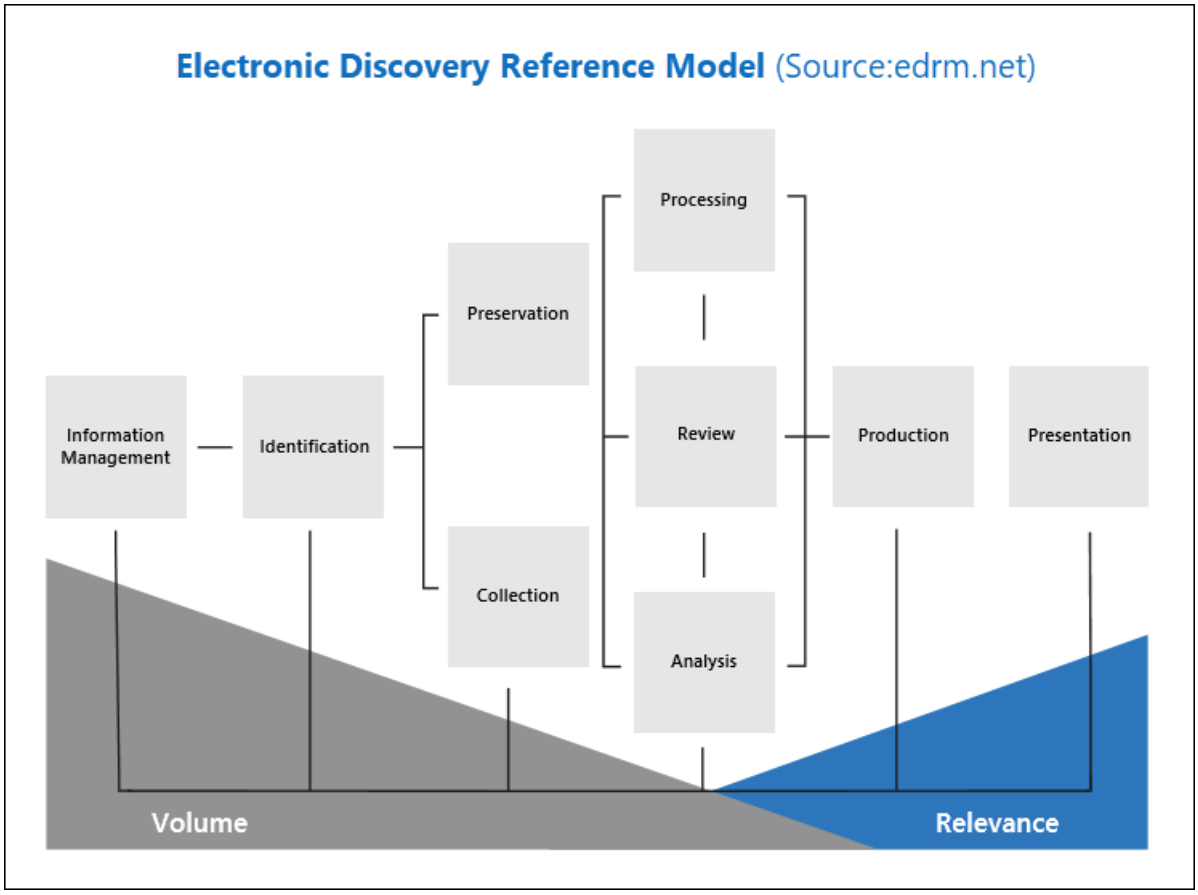

(Image source courtesy of edrm.net. The source image was made available under Creative Commons Attribution 3.0 Unported License.)

At a high level, here's how Advanced eDiscovery supports the EDRM workflow:

- **Identification.** After you identify potential persons of interest in an investigation, you can add them as custodians (also called *data custodians*, because they may possess information that's relevant to the investigation) to an Advanced eDiscovery case. After users are added as custodians, it's easy to preserve, collect, and review custodian documents.
- **Preservation.** To preserve and protect data that's relevant to an investigation, Advanced eDiscovery lets you place a legal hold on the data sources associated with the custodians in a case. You can also place non-custodial data on hold. Advanced eDiscovery also has a built-in communications workflow so you can send legal hold notifications to custodians and track their acknowledgments.
- **Collection.** After you identified (and preserved) the data sources relevant to the investigation, you can use the built-in search tool in Advanced eDiscovery search for and collect live data from the custodial data sources (and non-custodial data sources, if applicable) that may be relevant to the case.
- **Processing.** After you've collected all data relevant to the case, the next step is process it for further

review and analysis. In Advanced eDiscovery, the in-place data that you identified in the collection phase is copied to an Azure Storage location (called a *review set*), which provides you with a static view of the case data.

- **Review.** After data has been added to a review set, you can view specific documents and run additional queries to reduce the data to what is most relevant to the case. Also, can annotate and tag specific documents.
- **Analysis.** Advanced eDiscovery provides integrated analytics tool that helps you further cull data from the review set that you determine isn't relevant to the investigation. In addition to reducing the volume of relevant data, Advance eDiscovery also helps you save legal review costs by letting you organize content to make the review process easier and more efficient.
- **Production** and **Presentation.** When you're ready, you can export documents from a review set for legal review. You can export documents in their native format or in an EDRM-specified format so they can be imported into third-party review applications.

More information

To get started using Advanced eDiscovery, see:

- Set up Advanced eDiscovery
- Create and manage an Advanced eDiscovery case

Limits in Advanced eDiscovery

2/18/2021 • 4 minutes to read • Edit Online

This article describes the limits in the Advanced eDiscovery solution in Microsoft 365.

Case and review set limits

The following table lists the limits for cases and review sets in Advanced eDiscovery.

DESCRIPTION OF LIMIT	LIMIT
Total number of documents that can be added to a case (for all review sets in a case).	3 million
Total file size per load set. This includes loading non-Office 365 into a review set.	300 GB
Total amount of data loaded into all review sets in the organization per day.	2 TB
Maximum number of loads sets per case.	200
Maximum number of review sets per case.	20
Maximum number of tag groups per case.	1000
Maximum number of tags per case.	1000

Indexing limits

The following table lists the indexing limits in Advanced eDiscovery.

DESCRIPTION OF LIMIT	LIMIT
Maximum number of characters extracted from a single file.	10 million[1]
Maximum size of a single file.	100 MB[1]
Maximum depth of embedded items in a document.	25[1]
Maximum size of files processed by Optical Character Recognition (OCR).	24 MB[1]
Maximum number of indexing jobs per organization per day.	10^6

Search limits

The limits described in this section are related to using the search tool on the **Searches** tab to collect data for a case. For more information, see Collect data for a case in Advanced eDiscovery.

DESCRIPTION OF LIMIT	LIMIT
Maximum number of mailboxes or sites that can be searched in a single search.	No limit
Maximum number of searches that can run at the same time.	No limit
Maximum number of searches that a single user can start at the same time.	10
Maximum number of characters for a search query (including operators and conditions).	10,000 [2]
Minimum number of alpha characters for prefix wildcards; for example, **one*** or **set***.	3
Maximum variants returned when using prefix wildcard to search for an exact phrase or when using a prefix wildcard and the **NEAR** Boolean operator.	10,000 [3]
Maximum number of items per user mailbox that are displayed on preview page for searches. The newest items are displayed.	100
Maximum number of items from all mailboxes displayed on preview page for searches.	1,000
Maximum number of mailboxes that can be previewed for search results. If there are more than 1000 mailboxes that contain items that match the search query, only the top 1,000 mailboxes with the most results are available for preview.	1,000
Maximum number of items from SharePoint and OneDrive for Business sites displayed on preview page for searches. The newest items are displayed.	200
Maximum number of SharePoint and OneDrive for Business sites that can be previewed for search results. If there are more than 200 sites that contain items that match the search query, only the top 200 sites with the most results are available for preview.	200
Maximum number of items per public folder mailbox displayed on preview page for searches.	100
Maximum number of items found in all public folder mailbox items displayed on preview page for searches.	200

DESCRIPTION OF LIMIT	LIMIT
Maximum number of public folder mailboxes that can be previewed for search results. If there are more than 500 public folder mailboxes that contain items that match the search query, only the top 500 mailboxes with the most results are available for preview.	500

Viewer limits

DESCRIPTION OF LIMIT	LIMIT
Maximum size of Excel file that can be viewed in the native viewer.	4 MB

Export limits

DESCRIPTION OF LIMIT	LIMIT
Maximum size of a single export.	3 million documents or 100 GB, whichever is smaller
Maximum amount of data in a single day.	2 TB
Maximum concurrent exports in your organization.	10 [4]
Maximum concurrent exports per user.	3
Maximum size of a single PST file.	10 GB
Maximum concurrent exports per review set.	1

Review set download limits

DESCRIPTION OF LIMIT	LIMIT
Total file size or maximum number of documents downloaded from a review set.	3 MB or 50 documents [5]

NOTE

[1] Any item that exceeds a single file limit will show up as a processing error.

[2] When searching SharePoint and OneDrive for Business locations, the characters in the URLs of the sites being searched count against this limit.

[3] For non-phrase queries (a keyword value that doesn't use double quotation marks) we use a special prefix index. This tells us that a word occurs in a document, but not where it occurs in the document. To do a phrase query (a keyword value with double quotation marks), we need to compare the position within the document for the words in the phrase. This means that we can't use the prefix index for phrase queries. In this case, we internally expand the query with all possible words that the prefix expands to; for example, **time*** can expand to **"time OR timer OR times OR timex OR timeboxed OR ..."**. The limit of 10,000 is the maximum number of variants the word can expand to, not the number of documents matching the query. There is no upper limit for non-phrase terms.

[4] This limit is shared across all eDiscovery tools. This means that concurrent exports in Content search, Core eDiscovery, and Advanced eDiscovery are applied against this limit.

[5] This limit applies to downloading selected documents from a review set. It doesn't apply to exporting documents from a review set. For more information about downloading and exporting documents, see Export case data in Advanced eDiscovery.

[6] Indexing limits per organization per day. As a workaround, you can select multiple custodians on the **Data sources** tab in a case and then click **Update index** to avoid creating a separate index job for each custodian.

Supported file types in Advanced eDiscovery

11/2/2020 • 4 minutes to read • Edit Online

Advanced eDiscovery supports many file types at many different levels. The support files types are described in the following tables in this article. This list isn't finalized, and we will add new file types as we continue our validation testing. These tables indicate if a file type is supported for text extraction (and Optical Character Recognition or OCR text extraction for image files), viewable in the native viewer and also support in the Annotate viewer in Advanced eDiscovery.

Archive / Container

MIME TYPE	FILE IDENTIFICATION	METADATA EXTRACTION	CONTAINER EXTRACTION	POSSIBLE EXTENSIONS		
application/x-7z-compressed	Yes	Yes	Yes	.7z		
application/x-rar-compressed	Yes	Yes	Yes	.rar		
application/x-tar	Yes	Yes	Yes	.tar		
application/zip	Yes	Yes	Yes	.zip		

Audio / Video

MIME TYPE	FILE IDENTIFICATION	METADATA EXTRACTION	TEXT EXTRACTION	NATIVE VIEWER	ANNOTATE VIEWER	POSSIBLE EXTENSIONS
application/mp4	Yes	Yes	No	Yes	No	.f4v; .m4a; .m4v; .mp4; .mp4v; .mpeg; .mpeg4
audio/mpeg	Yes	Yes	No	Yes	No	.mpeg
video/3gpp	Yes	Yes	No	Yes	No	.3gp
video/3gpp2	Yes	Yes	No	Yes	No	.3g2; .3gp2
video/quicktime	Yes	Yes	No	Yes	No	.moov; .mov; .qt
video/x-m4v	Yes	Yes	No	Yes	No	.m4v

MIME TYPE	FILE IDENTIFICATION	METADATA EXTRACTION	TEXT EXTRACTION	NATIVE VIEWER	ANNOTATE VIEWER	POSSIBLE EXTENSIONS

Database

MIME TYPE	FILE IDENTIFICATION	METADATA EXTRACTION	TEXT EXTRACTION	NATIVE VIEWER	ANNOTATE VIEWER	POSSIBLE EXTENSIONS
application/x-msaccess	Yes	Yes	Yes	No	No	.mdb

Email

MIME TYPE	FILE IDENTIFICATION	METADATA EXTRACTION	TEXT EXTRACTION	NATIVE VIEWER	ANNOTATE VIEWER	POSSIBLE EXTENSIONS
application/vnd.ms-outlook	Yes	Yes	Yes	Yes	Yes	.msg
message/rfc822	Yes	Yes	Yes	Yes	Yes	.eml
text/vcard-contact	Yes	Yes	Yes	Yes	Yes	.vcf

Email Container

MIME TYPE	FILE IDENTIFICATION	METADATA EXTRACTION	CONTAINER EXTRACTION	POSSIBLE EXTENSIONS		
application/mbox	Yes	Yes	Yes	.mbox		
application/vnd.ms-outlook-pst	Yes	Yes	Yes	.pst		

HTML

MIME TYPE	FILE IDENTIFICATION	METADATA EXTRACTION	TEXT EXTRACTION	NATIVE VIEWER	ANNOTATE VIEWER	POSSIBLE EXTENSIONS
application/xhtml+xml	Yes	Yes	Yes	Yes	Yes	.xhtml

MIME TYPE	FILE IDENTIFICATION	METADATA EXTRACTION	TEXT EXTRACTION	NATIVE VIEWER	ANNOTATE VIEWER	POSSIBLE EXTENSIONS
application/xml	Yes	Yes	Yes	Yes	Yes	.xml
text/html	Yes	Yes	Yes	Yes	Yes	.htm; .html; .shtml

Image

MIME TYPE	FILE IDENTIFICATION	METADATA EXTRACTION	OCR TEXT EXTRACTION	NATIVE VIEWER	ANNOTATE VIEWER	POSSIBLE EXTENSIONS
image/bmp	Yes	Yes	Yes	Yes	Yes	.bmp
image/emf	Yes	Yes	Yes	Yes	Yes	.emf
image/gif	Yes	Yes	Yes	Yes	Yes	.gif
image/jpeg	Yes	Yes	Yes	Yes	Yes	.jpeg; .jpg
image/png	Yes	Yes	Yes	Yes	Yes	.png
image/svg+xml	Yes	Yes	Yes	Yes	No	.svg
image/tiff	Yes	Yes	Yes	Yes	Yes	.tif
image/vnd.dwg	Yes	Yes	Yes	Yes	Yes	.dwg; .dxf
image/wmf	Yes	Yes	Yes	Yes	Yes	.wmf

Microsoft Excel

MIME TYPE	FILE IDENTIFICATION	METADATA EXTRACTION	TEXT EXTRACTION	NATIVE VIEWER	ANNOTATE VIEWER	POSSIBLE EXTENSIONS
application/vnd.ms-excel	Yes	Yes	Yes	Yes	Yes	.dat; .xls
application/vnd.ms-excel.sheet.binary.macroenabled.12	Yes	Yes	Yes	Yes	No	.xlsb

MIME TYPE	FILE IDENTIFICATION	METADATA EXTRACTION	TEXT EXTRACTION	NATIVE VIEWER	ANNOTATE VIEWER	POSSIBLE EXTENSIONS
application/vnd.ms-excel.sheet.macroenabled.12	Yes	Yes	Yes	Yes	Yes	.xlsm
application/vnd.ms-excel.template.macroenabled.12	Yes	Yes	Yes	No	No	.xltm
application/vnd.openxmlformats-officedocument.spreadsheetml.sheet	Yes	Yes	Yes	Yes	Yes	.xlsx
application/vnd.openxmlformats-officedocument.spreadsheetml.template	Yes	Yes	Yes	Yes	Yes	.xltx

Microsoft OneNote

MIME TYPE	FILE IDENTIFICATION	METADATA EXTRACTION	TEXT EXTRACTION	NATIVE VIEWER	ANNOTATE VIEWER	POSSIBLE EXTENSIONS
application/onenote	Yes	Yes	Yes	Yes	No	.one

Microsoft PowerPoint

MIME TYPE	FILE IDENTIFICATION	METADATA EXTRACTION	TEXT EXTRACTION	NATIVE VIEWER	ANNOTATE VIEWER	POSSIBLE EXTENSIONS
application/vnd.ms-powerpoint	Yes	Yes	Yes	Yes	Yes	.pot; .pps; .ppt

MIME TYPE	FILE IDENTIFICATION	METADATA EXTRACTION	TEXT EXTRACTION	NATIVE VIEWER	ANNOTATE VIEWER	POSSIBLE EXTENSIONS
application/vnd.openxmlformats-officedocument.presentationml.presentation	Yes	Yes	Yes	Yes	Yes	.pptx
application/vnd.openxmlformats-officedocument.presentationml.slideshow	Yes	Yes	Yes	Yes	Yes	.ppsx
application/vnd.openxmlformats-officedocument.presentationml.template	Yes	Yes	Yes	Yes	Yes	.potx

Microsoft Project

MIME TYPE	FILE IDENTIFICATION	METADATA EXTRACTION	TEXT EXTRACTION	NATIVE VIEWER	ANNOTATE VIEWER	POSSIBLE EXTENSIONS
application/vnd.ms-project	Yes	Yes	Yes	No	Yes	.mpp

Microsoft Publisher

MIME TYPE	FILE IDENTIFICATION	METADATA EXTRACTION	TEXT EXTRACTION	NATIVE VIEWER	ANNOTATE VIEWER	POSSIBLE EXTENSIONS
application/x-mspublisher	Yes	Yes	Yes	Yes	Yes	.pub

Microsoft Visio

MIME TYPE	FILE IDENTIFICATION	METADATA EXTRACTION	TEXT EXTRACTION	NATIVE VIEWER	ANNOTATE VIEWER	POSSIBLE EXTENSIONS
application/vnd.ms-visio.drawing	Yes	Yes	Yes	Yes	No	

MIME TYPE	FILE IDENTIFICATION	METADATA EXTRACTION	TEXT EXTRACTION	NATIVE VIEWER	ANNOTATE VIEWER	POSSIBLE EXTENSIONS
application/vnd.visio	Yes	Yes	Yes	Yes	Yes	.vsd

Microsoft Word

MIME TYPE	FILE IDENTIFICATION	METADATA EXTRACTION	TEXT EXTRACTION	NATIVE VIEWER	ANNOTATE VIEWER	POSSIBLE EXTENSIONS
application/msword	Yes	Yes	Yes	Yes	Yes	.dat; .doc
application/rtf	Yes	Yes	Yes	Yes	Yes	.doc; .rtf
application/vnd.ms-word.document.macroenabled.12	Yes	Yes	Yes	Yes	Yes	.docm
application/vnd.ms-word.template.macroenabled.12	Yes	Yes	Yes	Yes	Yes	.dotm
application/vnd.openxmlformats-officedocument.wordprocessingml.document	Yes	Yes	Yes	Yes	Yes	.docx
application/vnd.openxmlformats-officedocument.wordprocessingml.template	Yes	Yes	Yes	Yes	Yes	.dotx

Microsoft Works

MIME TYPE	FILE IDENTIFICATION	METADATA EXTRACTION	TEXT EXTRACTION	NATIVE VIEWER	ANNOTATE VIEWER	POSSIBLE EXTENSIONS
application/vnd.ms-works-ss	Yes	Yes	No	No	No	.wps

MIME TYPE	FILE IDENTIFICATION	METADATA EXTRACTION	TEXT EXTRACTION	NATIVE VIEWER	ANNOTATE VIEWER	POSSIBLE EXTENSIONS
application/vnd.ms-works-wp	Yes	Yes	No	No	No	.wps

Open Document Format

MIME TYPE	FILE IDENTIFICATION	METADATA EXTRACTION	TEXT EXTRACTION	NATIVE VIEWER	ANNOTATE VIEWER	POSSIBLE EXTENSIONS
application/vnd.oasis.opendocument.text	Yes	Yes	Yes	Yes	Yes	.odt

Other

MIME TYPE	FILE IDENTIFICATION	METADATA EXTRACTION	TEXT EXTRACTION	NATIVE VIEWER	ANNOTATE VIEWER	POSSIBLE EXTENSIONS
application/json	Yes	Yes	Yes	Yes	Yes	n/a
application/vnd.ms-graph	Yes	Yes	No	No	No	
application/winhlp	Yes	Yes	No	No	No	.hlp
application/x-tnef	Yes	Yes	No	No	No	

Plain Text

MIME TYPE	FILE IDENTIFICATION	METADATA EXTRACTION	TEXT EXTRACTION	NATIVE VIEWER	ANNOTATE VIEWER	POSSIBLE EXTENSIONS
text/csv	Yes	Yes	Yes	Yes	Yes	.csv
text/plain	Yes	Yes	Yes	Yes	Yes	.con; .css; .csv; .dat; .pl; .txt

Portable Document Format

MIME TYPE	FILE IDENTIFICATION	METADATA EXTRACTION	TEXT EXTRACTION	NATIVE VIEWER	ANNOTATE VIEWER	POSSIBLE EXTENSIONS
application/pdf	Yes	Yes	Yes	Yes	Yes	.pdf

Word Perfect

MIME TYPE	FILE IDENTIFICATION	METADATA EXTRACTION	TEXT EXTRACTION	NATIVE VIEWER	ANNOTATE VIEWER	POSSIBLE EXTENSIONS
application/vnd.wordperfect; version=5.0	Yes	Yes	Yes	No	No	.wpd
application/vnd.wordperfect; version=5.1	Yes	Yes	Yes	No	No	.wpd
application/vnd.wordperfect; version=6.x	Yes	Yes	Yes	No	No	.wpd

Word Pro

MIME TYPE	FILE IDENTIFICATION	METADATA EXTRACTION	TEXT EXTRACTION	NATIVE VIEWER	ANNOTATE VIEWER	POSSIBLE EXTENSIONS
application/vnd.lotus-wordpro	Yes	Yes	No	No	No	.lwp

Document metadata fields in Advanced eDiscovery

11/2/2020 • 10 minutes to read • Edit Online

The following table lists the metadata fields for documents in a review set in a case in Advanced eDiscovery. The table provides the following information:

- **Field name** and **Display field name:** The name of the metadata field and the name of the field that's displayed when viewing the file metadata of a selected document in a review set. Some metadata fields aren't included when viewing the file metadata of a document. These fields are highlighted with an asterisk (*).
- **Searchable field name:** The name of the property that you can search for when running a review set query. A blank cell means that you can't search for the field in a review set query.
- **Exported field name:** The name of the metadata field that included when documents are exported. A blank cell means the field isn't included with the exported metadata.
- **Description:** A description of the metadata field.

> **NOTE**
>
> The **Keywords** field in review set search uses Keyword Query Language (KQL). The fields listed in the **Searchable field name** column can be used in the **Keywords** field in a review set search to form complex queries without you having to use the query builder. For more information about KQL, see Keyword Query Language syntax reference.

FIELD NAME AND DISPLAY FIELD NAME	SEARCHABLE FIELD NAME	EXPORTED FIELD NAME	DESCRIPTION
Attachment Content Id	AttachmentContentId		Attachment content Id of the item.
Attachment Names	AttachmentNames	Attachment_Names	List of names of attachments.
Attorney client privilege score	AttorneyClientPrivilegeScore		Attorney-client privilege model content score.
Author	Author	Doc_authors	Author from the document metadata.
BCC	Bcc	Email_bcc	Bcc field for message types. Format is **DisplayName <SMTPAddress>**.
CC	Cc	Email_cc	Cc field for message types. Format is **DisplayName <SMTPAddress>**.
Compliance labels	ComplianceLabels	Compliance_labels	Retention labels applied to content in Office 365.

FIELD NAME AND DISPLAY FIELD NAME	SEARCHABLE FIELD NAME	EXPORTED FIELD NAME	DESCRIPTION
Compound Path	CompoundPath	Compound_path	Human readable path that describes the source of the item.
Content*	Content		Extracted text of the item.
Conversation Body	Conversation Body		Conversation body of the item.
Conversation Topic	Conversation Topic		Conversation topic of the item.
Conversation ID	ConversationId	Conversation_ID	Conversation Id from the message.
Conversation Index		Conversation_index	Conversation index from the message.
Conversation Pdf Time	ConversationPdfTime		Date when the PDF version of the conversation was created.
Conversation Redaction Burn Time	ConversationRedactionBurn Time		Date when the PDF version of the conversation was created for Chat.
Document date created	CreatedTime	Doc_date_created	Create date from document metadata.
Custodian	Custodian	Custodian	Name of the custodian the item was associated with.

FIELD NAME AND DISPLAY FIELD NAME	SEARCHABLE FIELD NAME	EXPORTED FIELD NAME	DESCRIPTION
Date	Date	Date	Date is a computed field that depends on the file type. Email: Sent date Email attachments: Last modified date of the document;if not available, the parent's Sent date Embedded documents: Last modified date of the document; if not available, the parent's last modified date SPO documents (includes modern attachments): SharePoint Last modified date; if not available, the documents last modified date Non-Office 365 documents: Last modified date Meetings: Meeting start date VoiceMail: Sent date IM: Sent date
Other paths	Dedupedcompoundpath	Deduped_compound_path	List of compound paths of documents that are exact duplicates (email: based on content, documents: based on hash).
Other custodians	DedupedCustodians	Deduped_custodians	List of custodians of documents that are exact duplicates (for email, based on content; for documents, based on hash).
Other file IDs	DedupedFileIds	Deduped_file_IDs	List of file IDs of documents that are exact duplicates (for email, based on content; for documents, based on hash).
Document comments	DocComments	Doc_comments	Comments from the document metadata.
Document company		Doc_company	Company from the document metadata.
DocIndex*			The index in the family. **-1** or **0** means it is the root.
Document keywords		Doc_keywords	Keywords from the document metadata.

FIELD NAME AND DISPLAY FIELD NAME	SEARCHABLE FIELD NAME	EXPORTED FIELD NAME	DESCRIPTION
Document modified by		Doc_modified_by	Last modified date by from document metadata.
Document Revision		Doc_revision	Revision from the document metadata.
Document subject		Doc_subject	Subject from the document metadata.
Document template		Doc_template	Template from the document metadata.
Dominant theme	DominantTheme	Dominant_theme	Dominant theme as calculated for analytics.
Duplicate subset		Duplicate_subset	Group ID for exact duplicates.
EmailAction*		Email_action	Values are **None**, **Reply**, or **Forward**; based on the subject line of a message.
Email Delivery Receipt		Email_delivery_receipt	Email address supplied in Internet Headers for delivery receipt.
Importance	EmailImportance	Email_importance	Importance of the message: **0** - Low; **1** - Normal; **2** - High
EmailLevel*		Email_level	Indicates a message's level within the email thread it belongs to; attachments inherit its parent message's value.
Email Message Id		Email_message_ID	Internet message Id from the message.
EmailReadReceipt*		Email_read_receipt	Email address supplied in Internet Headers for read receipt.
Email Security	EmailSecurity	Email_security	Security setting of the message: **0** - None; **1** - Signed; **2** - Encrypted; **3** - Encrypted and signed.
Email Sensitivity	EmailSensitivity	email_sensitivity	Sensitivity setting of the message: **0** - None; **1** Personal; **2** - Private; **3** - CompanyConfidential.

FIELD NAME AND DISPLAY FIELD NAME	SEARCHABLE FIELD NAME	EXPORTED FIELD NAME	DESCRIPTION
Email set	EmailSet	Email_set	Group ID for all messages in the same email set.
EmailThread*		Email_thread	Position of the message within the email set; consists of node IDs from the root to the current message and are separated by periods (.).
Extracted content type		Extracted_content_type	Extracted content type, in the form of mime type; for example, **image/jpeg**
ExtractedTextLength*		Extracted_text_length	Number of characters in the extracted text.
Family relevance score Case issue 1*		Family_relevance_score_case_issue_1	Family relevance score Case issue 1 from Relevance.
FamilyDuplicateSet*		Family_duplicate_set	Numeric identifier for families that are exact duplicates of each other (same content and all the same attachments).
Family ID	FamilyId	Family_ID	Family Id groups together all items; for email, this includes the message and all attachments; for documents, this includes the document and any embedded items.
Family Size		Family_size	Number of documents in the family.
File relevance score Case issue 1*		File_relevance_score_case_issue_1	File relevance score Case issue 1 from Relevance.
File class	FileClass	File_class	For content from SharePoint and OneDrive: **Document**; for content from Exchange: **Email** or **Attachment**.
File ID	FileId	File_ID	Document identifier unique within the case.
File system date created		File_system_date_created	Created date from file system (only applies to non-Office 365 data).

FIELD NAME AND DISPLAY FIELD NAME	SEARCHABLE FIELD NAME	EXPORTED FIELD NAME	DESCRIPTION
File system date modified		File_system_date_modified	Modified date from file system (only applies to non-Office 365 data).
File Type	FileType		File type of the item based on file extension.
Group Id	GroupID		Group ID for grouped content.
Has attachment	HasAttachment	Email_has_attachment	Indicates whether or not the message has attachments.
Has attorney	HasAttorney		**True** when at least one of the participants is found in the attorney list; otherwise, the value is **False**.
HasText*		Has_text	Indicates whether or not the item has text; possible values are **True** and **False**.
Immutable ID		Immutable_ID	This Id is used to uniquely identify a document within a review set. This field can't be used in a review set search and the Id can't be used to access a document in its native location.
Inclusive type	InclusiveType	Inclusive_type	Inclusive type calculated for analytics: **0** - not inclusive; **1** - inclusive; **2** - inclusive minus; **3** - inclusive copy.
In Reply To Id		In_reply_to_ID	In reply to Id from the message.
Is modern attachment	IsModernAttachment		This file is a modern attachment or linked file.
Is from document version	IsFromDocumentVersion		Current document is from a different version of another document.
Is email attachment	IsEmailAttachment		This item is from an email attachment that shows up as an attached item to the message.
Is inline attachment	IsInlineAttachment		This was attached inline and shows up in the body of the message.

FIELD NAME AND DISPLAY FIELD NAME	SEARCHABLE FIELD NAME	EXPORTED FIELD NAME	DESCRIPTION
Is Representative	IsRepresentative	Is_representative	One document in every set of exact duplicates is marked as representative.
Item class	ItemClass	Item_class	Item class supplied by exchange server; for example, **IPM.Note**
Last modified date	LastModifiedDate	Doc_date_modified	Last modified date from document metadata.
Load ID	LoadId	Load_ID	The Id of the load set in which the item was added to a review set.
Location	Location	Location	String that indicates the type of location that documents were sourced from. **Imported Data** - Non-Office 365 data **Teams** - Microsoft Teams **Exchange** - Exchange mailboxes **SharePoint** - SharePoint sites **OneDrive** - OneDrive accounts
Location name	LocationName	Location_name	String that identifies the source of the item. For exchange, this will be the SMTP address of the mailbox; for SharePoint and OneDrive, the URL for the site collection.
Marked as representative	MarkAsRepresentative		One document from each set of exact duplicates is marked as representatives.
Marked as pre tagged Case issue 1*		Marked_as_pre_tagged_Case_issue_1	Marked as pre-tagged Case issue 1 from Relevance.
Marked as seed Case issue 1*		Marked_as_seed_Case_issue_1	Marked as seed Case issue 1 from Relevance.
Meeting End Date	MeetingEndDate	Meeting_end_date	Meeting end date for meetings.
Meeting Start Date	MeetingStartDate	Meeting_start_date	Meeting start date for meetings.

FIELD NAME AND DISPLAY FIELD NAME	SEARCHABLE FIELD NAME	EXPORTED FIELD NAME	DESCRIPTION
Message kind	MessageKind	Message_kind	The type of message to search for. Possible values: **contacts** **docs** **email** **externaldata** **faxes** **im** **journals** **meetings** **microsoftteams** (returns items from chats, meetings, and calls in Microsoft Teams) **notes** **posts** **rssfeeds** **tasks** **voicemail**
Native Extension	NativeExtension	Native_extension	Native extension of the item.
Native file name	NativeFileName	Native_file_name	Native file name of the item.
NativeMD5		Native_MD5	MD5 hash (128-bit hash value) of the file stream.
NativeSHA256		Native_SHA_256	SHA256 hash (256-bit hash value) of the file stream.
ND/ET Sort: Excluding attachments	NdEtSortExclAttach	ND_ET_sort_excl_attach	Concatenation of the email thread (ET) set and Near-duplicate (ND) set. This field is used for efficient sorting at review time. A **D** is prefixed to ND sets and an **E** is prefixed to ET sets.
ND/ET Sort: Including attachments	NdEtSortInclAttach	ND_ET_sort_incl_attach	Concatenation of an email thread (ET) set and near-duplicate (ND) set. This field is used for efficient sorting at review time. A **D** is prefixed to ND sets and an **E** is prefixed to ET sets. Each email item in an ET set is followed by its appropriate attachments.
Normalized relevance score Case issue 1		Normalized_relevance_score_case_issue_1	Normalized relevance score Case issue 1 from Relevance.
O365 authors		O365_authors	Author from SharePoint.

FIELD NAME AND DISPLAY FIELD NAME	SEARCHABLE FIELD NAME	EXPORTED FIELD NAME	DESCRIPTION
O365 created by		O365_created_by	Created by from SharePoint.
O365 date created		O365_date_created	Created date from SharePoint.
O365 date modified		O365_date_modified	Last modified date from SharePoint.
O365 modified by		O365_modified_by	Modified by from SharePoint.
Parent ID	ParentId	Parent_ID	Id of the item's parent.
ParentNode		Parent_node	The closest preceding email message in the email thread.
Parent path	ParentPath	Parent_path	Compound path of the direct parent of the item.
Participant domains	ParticipantDomains	Email_participant_domains	List of all domains of participants of a message.
Participants	Participants	Email_participants	List of all participants of a message; for example, Sender, To, Cc, Bcc.
Pivot ID	PivotId	Pivot_ID	The ID of a pivot.
Potentially privileged	PotentiallyPrivileged	Potentially_privileged	True if attorney-client privilege detection model considers the document potentially privileged
Processing status	ProcessingStatus	Error_code	Processing status after the item was added to a review set.
Read percent Case issue 1		Read_percent_Case_issue_1	Read percent Case issue 1 from Relevance.
Read percentile	ReadPercentile		Read percentile for the document based on Relevance.
Recipient Count		Recipient_count	Number of recipients in the message.
Recipient domains	RecipientDomains	Email_recipient_domains	List of all domains of recipients of a message.
Recipients	Recipients	Email_recipients	List of all recipients of a message (To, Cc, Bcc).

FIELD NAME AND DISPLAY FIELD NAME	SEARCHABLE FIELD NAME	EXPORTED FIELD NAME	DESCRIPTION
Relevance load group Case issue 1		Relevance_load_group_case_issue_1	Relevance load group Case issue 1 from Relevance.
Relevance status description Case issue 1		Relevance_status_description_Case_issue_1	Relevance status description Case issue 1 from Relevance.
Relevance tag Case issue 1		Relevance_tag_case_issue_1	Relevance tag Case issue 1 from Relevance.
Relevance Comment		Relevance_comment	Comment field from Relevance.
Relevance score	RelevanceScore		Relevance score of a document based on Relevance.
Relevance tag	RelevanceTag		Relevance score of a document based on Relevance.
Representative ID	RepresentativeId		Numeric identifier of each set of exact duplicates.
Sender	Sender	Email_sender	Sender (From) field for message types. Format is **DisplayName <SmtpAddress>**.
Sender/Author	SenderAuthor		Calculated field comprised of the sender or author of the item.
Sender domain	SenderDomain	Email_sender_domain	Domain of the sender.
Sent	Sent	Email_date_sent	Sent date of the message.
Set Order: Inclusive First	SetOrderInclusivesFirst	Set_order_inclusives_first	Sorting field - email and attachments: counter-chronological; documents: pivot first then by descending similarity score.
SimilarityPercent		Similarity_percent	Indicates how similar a document is to the pivot of the near duplicate set.
Native file size	Size	Native_size	Number of bytes of the native item.
Subject	Subject	Email_subject	Subject of the message.

FIELD NAME AND DISPLAY FIELD NAME	SEARCHABLE FIELD NAME	EXPORTED FIELD NAME	DESCRIPTION
Subject/Title	SubjectTitle		Calculated field comprised of the subject or title of the item.
Tagged by Case issue 1		Tagged_by_Case_issue_1	User who tagged this document for Case issue 1 in Relevance.
Tags	Tags	Tags	Tags applied in a review set.
Themes list	ThemesList	Themes_list	Themes list as calculated for analytics.
Title	Title	Doc_title	Title from the document metadata.
To	To	Email_to	To field for message types. Format is **DisplayName<SmtpAddress>**
Unique in email set	UniqueInEmailSet		**False** if there's a duplicate of the attachment in its email set.
Was Remediated	WasRemediated	Was_Remediated	**True** if the item was remediated, otherwise **False**.
Word count	WordCount	Word_count	Number of words in the item.

NOTE

For more information about searchable properties when searching Office 365 content locations when you're collecting data for an Advanced eDiscovery case, see Keyword queries and search conditions for Content Search.

Set up attorney-client privilege detection in Advanced eDiscovery

11/2/2020 • 5 minutes to read • Edit Online

A major and costly aspect of the review phase of any eDiscovery process is reviewing documents for privileged content. Advanced eDiscovery provides machine learning-based detection of privileged content to make this process more efficient. This feature is called *attorney-client privilege detection*.

How does it work?

When attorney-client privilege detection is enabled, all documents in a review set will be processed by the attorney-client privilege detection model when you analyze the data in the review set. The model looks for two things:

- Privileged content – The model uses machine learning to determine the likelihood that the document contains content that is legal in nature.
- Participants – As part of setting up attorney-client privilege detection, you have to submit a list of attorneys for your organization. The model then compares the participants of the document with the attorney list to determine if a document has at least one attorney participant.

The model produces the following three properties for every document:

- **AttorneyClientPrivilegeScore:** The likelihood the document is legal in nature; the values for the score are between **0** and **1**.
- **HasAttorney:** This property is set to **true** if one of the document participants is listed in the attorney list; otherwise the value is **false**. The value is also set to **false** if your organization didn't upload an attorney list.
- **IsPrivilege:** This property is set to **true** if the value for **AttorneyClientPrivilegeScore** is above the threshold *or* if the document has an attorney participant; otherwise the value is set to **false**.

These properties (and their corresponding values) are added to the file metadata of the documents in a review set, as shown in the following screenshot:

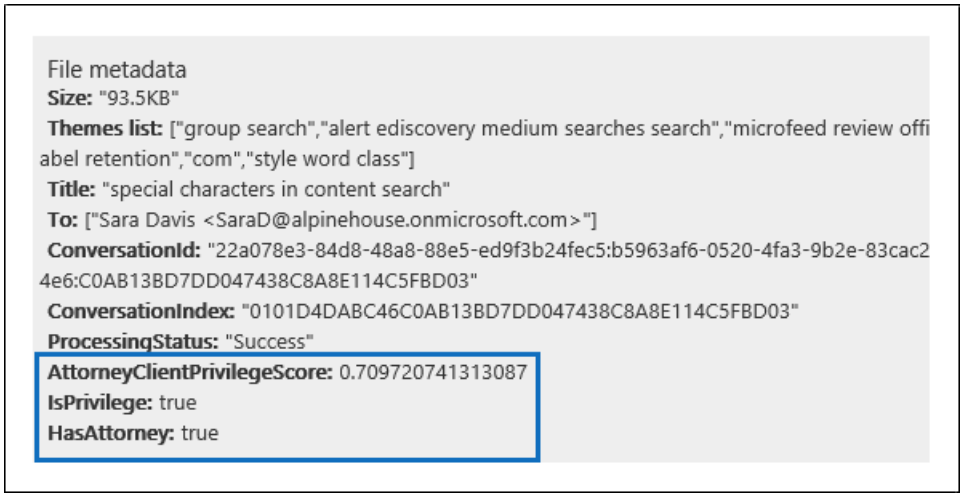

These three properties are also searchable within a review set. For more information, see Query the data in a review set.

Set up the attorney-client privilege detection model

To enable the attorney-client privilege detection model, your organization has to turn it on and then upload an attorney list.

Step 1: Turn on attorney-client privilege detection

A person who is an eDiscovery Administrator in your organization (a member of the eDiscovery Administrator subgroup in the eDiscovery Manager role group) must make the model available in your Advanced eDiscovery cases.

1. In the Security & Compliance Center, go to **eDiscovery > Advanced eDiscovery**.
2. On the **Advanced eDiscovery** home page, in the **Settings** tile, click **Configure global analytics settings**.

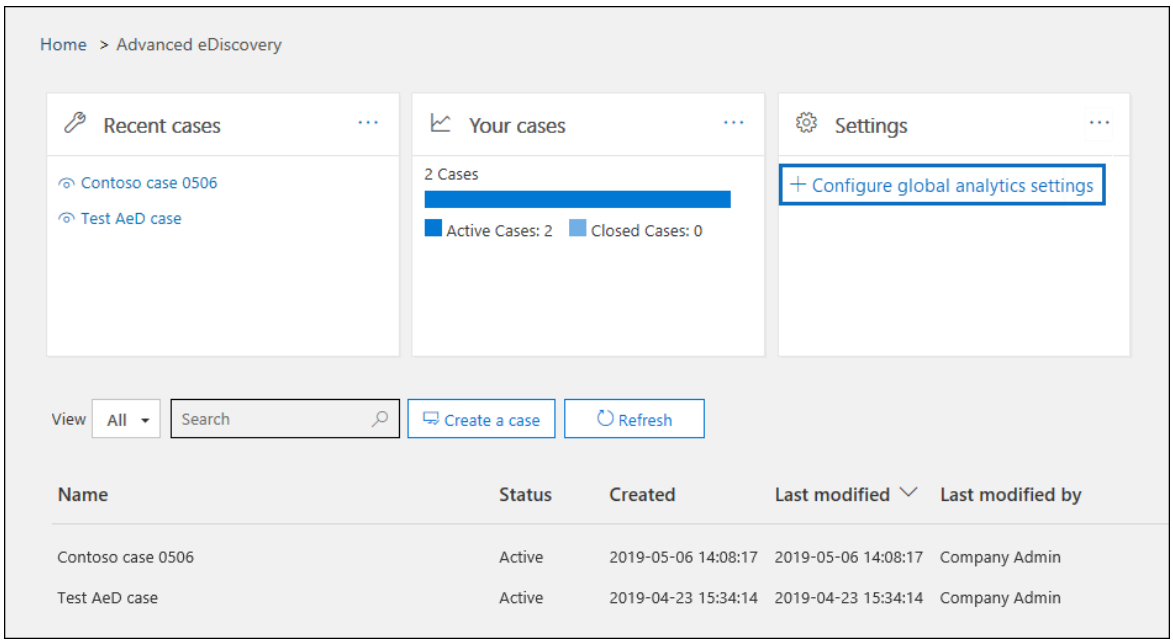

3. On the **Analytics settings** tab, select **Manage attorney-client privilege setting**.
4. On the **Attorney-client privilege** flyout page, use the toggle to turn on the feature and then select **Save**.

Step 2: Upload a list of attorneys (optional)

To take full advantage of the attorney-client privilege detection model and use the results of the **Has Attorney** or **Potentially Privileged** detection that was previously described, we recommend that you upload a list of email addresses for the lawyers and legal personnel who work for your organization.

To upload an attorney list for use by the attorney-client privilege detection model:

1. Create a .csv file (without a header row) and add the email address for each appropriate person on a separate line. Save this file to your local computer.
2. On the **Advanced eDiscovery** home page, in the **Settings** tile, select **Configure experimental features**, and then select **Manage attorney-client privilege setting**.

   The **Attorney-client privilege** page is displayed, and the **Attorney-client privilege detection** toggle is turned on.

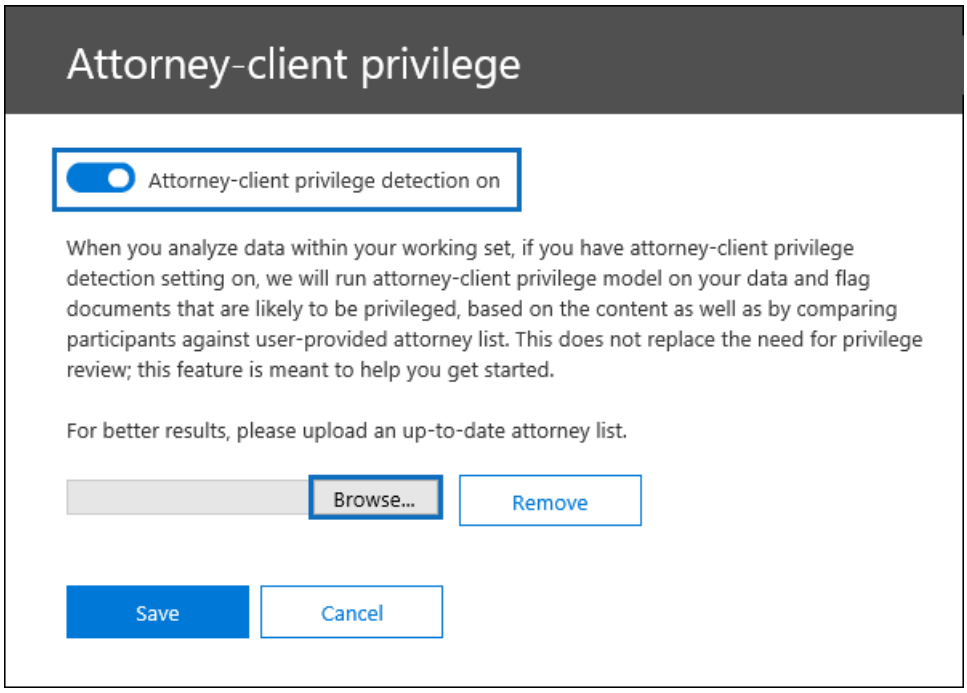

3. Select **Browse** and then find and select the .csv file that you created in step 1.
4. Select **Save** to upload the attorney list.

Use the attorney-client privilege detection model

Follow the steps in this section to use attorney-client privilege detection for documents in a review set.

Step 1: Create a smart tag group with attorney-client privilege detection model

One of the primary ways to see the results of attorney-client privilege detection in your review process is by using a smart tag group. A smart tag group indicates the results of the attorney-client privilege detection and shows the results in-line next to the tags in a smart tag group. This lets you quickly identify potentially privileged documents during document review. Additionally, you can also use the tags in the smart tag group to tag documents as privileged or non-privileged. For more information about smart tags, see Set up smart tags in Advanced eDiscovery.

1. In the review set that contains the documents that you analyzed in Step 1, select **Manage review set** and then select **Manage tags**.
2. Under **Tags**, select the pull-down next to **Add group** and then select **Add smart tag group**.

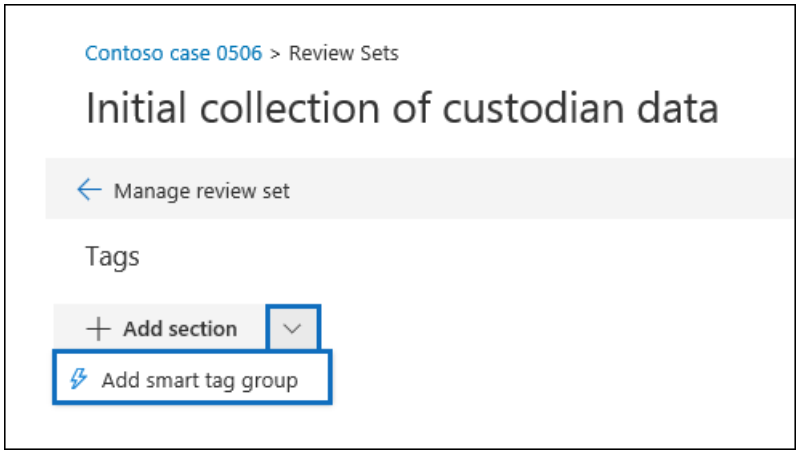

3. On the **Choose a model for your smart tag** page, choose **Select** next to **Attorney-client privilege**.

   A tag group named **Attorney-client privilege** is displayed. It contains two child tags named **Positive** and **Negative**, which correspond to the possible results produced by the model.

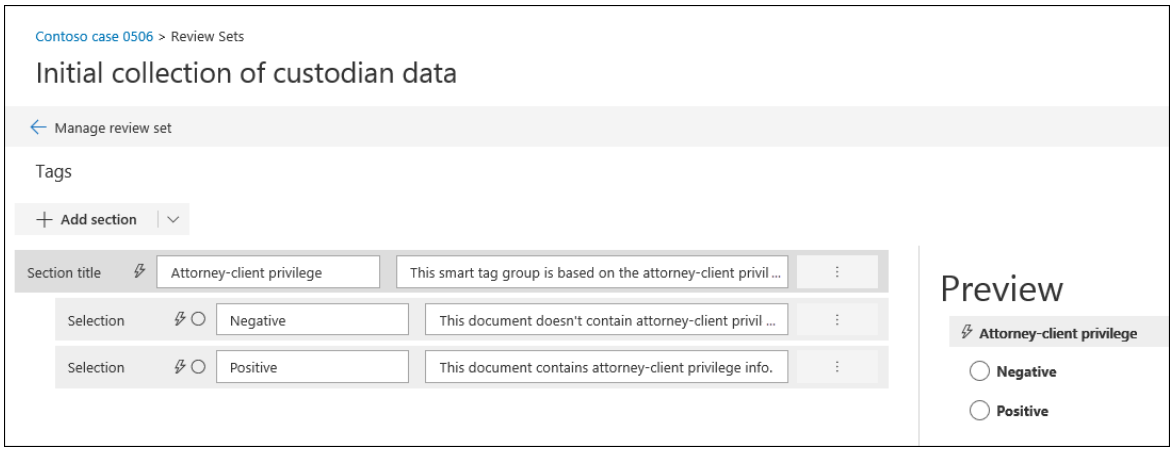

4. Rename the tag group and tags as appropriate for your review. For example, you can rename **Positive** to **Privileged** and **Negative** to **Not privileged**.

Step 2: Analyze a review set

When you analyze the documents in a review set, the attorney-client privilege detection model will also run and the corresponding properties (described in How does it work? will be added to every document in the review set. For more information about analyzing data in review set, see Analyze data in a review set in Advanced eDiscovery.

Step 3: Use the smart tag group for review of privileged content

After analyzing the review set and setting up smart tags, the next step is to review the documents. If the model has determined the document is potentially privileged, the corresponding smart tag in the **Tagging panel** will indicate the following results produced by the attorney-client privilege detection:

- If the document has content that may be legal in nature, the label **Legal content** is displayed next to the corresponding smart tag (which in this case is the default **Positive** tag).
- If the document has a participant who is found in your organization's attorney list, the label **Attorney** is displayed next to the corresponding smart tag (which in this case is also the default **Positive** tag).
- If the document has content that may be legal in nature *and* has a participant found in the attorney list, both the **Legal content** and **Attorney** labels are displayed.

If the model determines that a document doesn't contain content that is legal in nature or doesn't contain a participant from the attorney list, then neither label is displayed in the tagging panel.

For example, the following screenshots show two documents. The first one contains content that is legal in nature and has a participant found in the list of attorneys. The second contains neither and therefore doesn't display any labels.

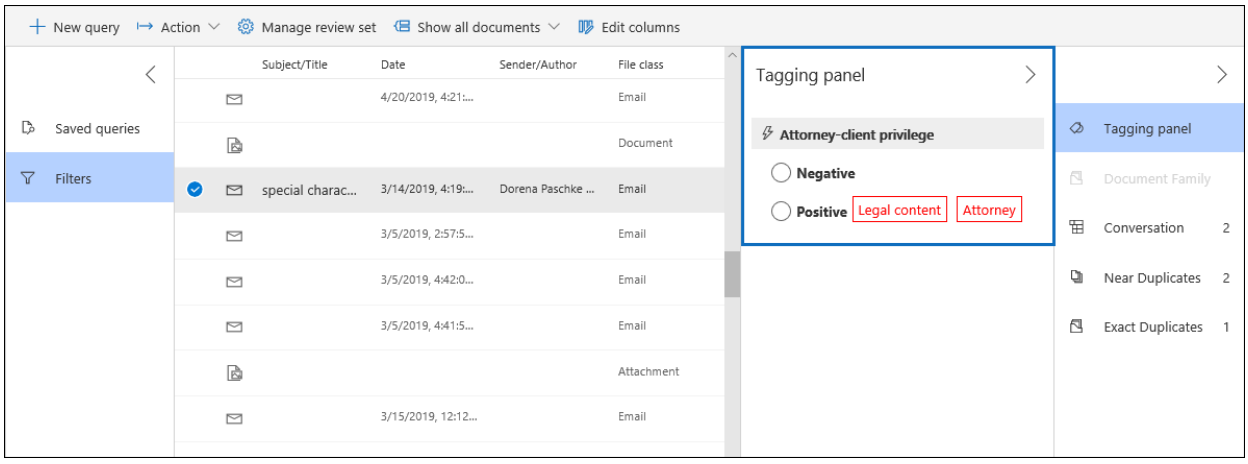

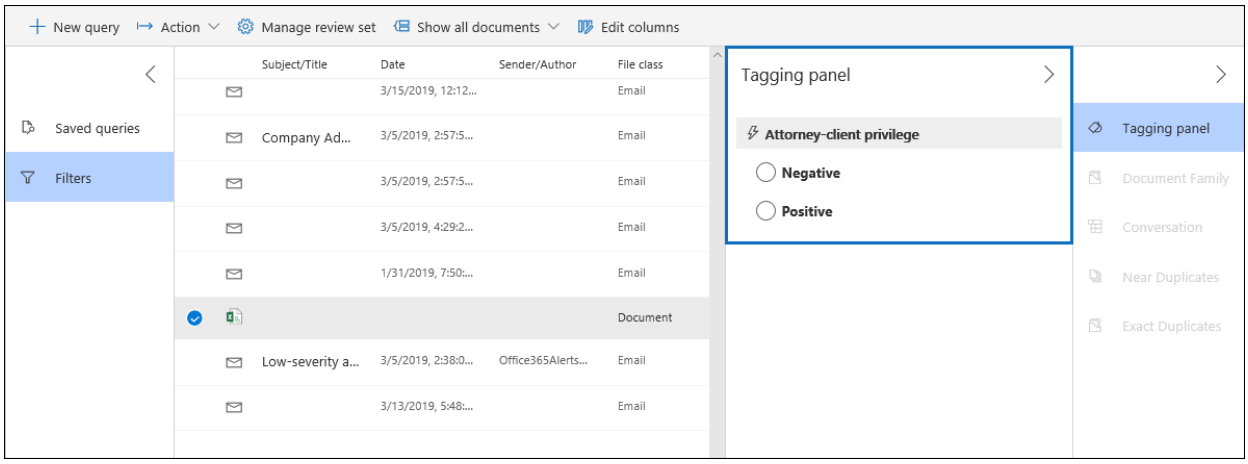

After you review a document to see if it contains privileged content, you can tag the document with the appropriate tag.

Review conversations in Advanced eDiscovery

11/2/2020 • 6 minutes to read • Edit Online

Instant messaging is a convenient way to ask questions, share ideas, or quickly communicate across large audiences. As instant messaging platforms, like Microsoft Teams, become core to enterprise collaboration, organizations must evaluate how their eDiscovery workflow addresses these new forms of communication and collaboration.

The Conversation Reconstruction feature in Advanced eDiscovery is designed to help you identify contextual content and produce distinct conversation views. This capability allows you to efficiently and rapidly review complete instant message conversations (also called *threaded conversations*) that are generated in platforms like Microsoft Teams.

With Conversation Reconstruction, you can use built-in capabilities to reconstruct, review, and export threaded conversations. Use Advanced eDiscovery Conversation Reconstruction to:

- Preserve unique message-level metadata across all messages within a conversation.
- Collect contextual messages around your search results.
- Review, annotate, and redact threaded conversations.
- Export individual messages or threaded conversations

Terminology

Here are few definitions to help you get start using Conversation Reconstruction.

- **Messages:** Represent the smallest unit of a conversation. Messages may vary in size, structure, and metadata.
- **Conversation:** Represents a grouping of one or more messages. Across different applications, conversations may be represented in different ways. In some applications, there is an explicit action that results from replying to an existing message. Conversations are formed explicitly as a result of this user action. For example, here is a screenshot of a channel conversation in Microsoft Teams.

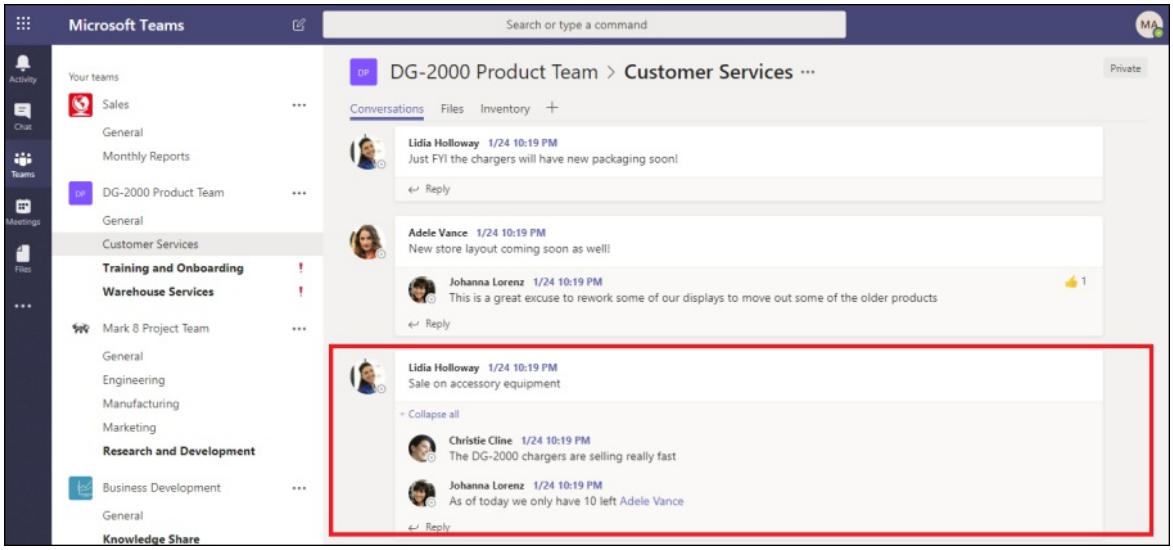

In other apps (such as 1xN chat messages in Teams), there is not a formal reply chain and instead messages appear as a "flat river of messages" within a single thread. In these types apps, conversations

are inferred from a group of messages that occur within a certain time. This "soft-grouping" of messages (as opposed to a reply chain) represent the "back and forth" conversation about a specific topic of interest.

Step 1: Run a search

After you have identified relevant custodians and content locations, you can create a search to find potentially relevant content. On the **Searches** tab in the Advanced eDiscovery case, you can create a search by clicking **New search** and following the wizard. For information about how you can create a search, build a search query, and view the search results, see Collect data for a case.

Step 2: Create a conversation review set

In a review set, you can search, tag, annotate, and redact documents, email messages, and chat conversations. In Advanced eDiscovery, you can customize your review of conversations, based in individual messages or threaded conversations. This is determined by the type of review set that you add the results of the search created in Step 1 to. There are two different types of review sets:

- **Standard review sets:** Messages in conversations are processed and displayed as individual items.
- **Conversation review sets:** Messages in conversations are processed individually but displayed in a conversation view. In a conversation review set, you can annotate, tag, and redact messages in a threaded conversation view.

For more information about how to review and manage content in a review set, see Manage review sets.

Step 3: Enable conversation retrieval options

After you have reviewed and finalized your search query, you can add the search results to a review set. When you add your search results into a review set, the original data is copied to an Azure Storage area to facilitate the review and analysis process. For more information about adding search results to a review set, see Add search results to a review set.

When you add data from conversations to a review set, you can use the conversation retrieval options to expand your search and include contextual messages. After you set the conversation retrieval options, the following things can happen:

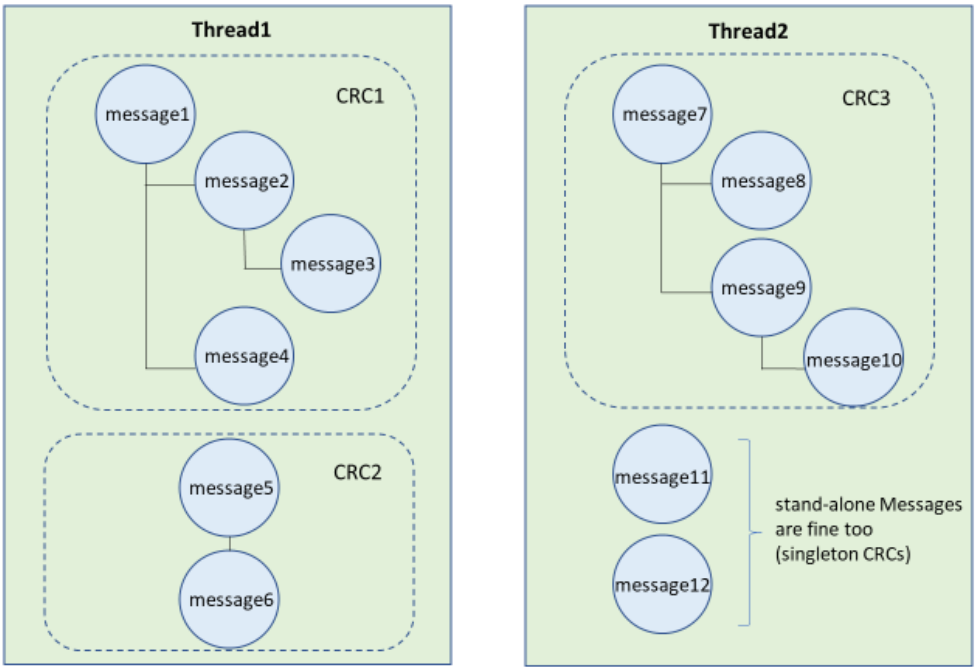

1. Using a keyword and date range query, the search returned a hit on *Message 3*. This message was part of a larger conversation, illustrated by *CRC1*.
2. When you add the data into a review set and enable the conversation retrieval options, Advanced eDiscovery will go back and collect other items in *CRC1*.
3. After the items have been added to the review set, you can review all the individual messages from *CRC1*.

To enable conversation retrieval:

1. On the **Searches** tab in the Advanced eDiscovery case, select a search, and then click **Add to review set** on the flyout page.
2. Select an existing review set or create a review set. You can configure retrieval options when adding search results to a standard or a conversation review set.
3. Under **Collection options**, configure the conversation retrieval options for the content sources that you want to expand in your search, and then click **Add** to start the process.
4. After the **Add to review set** job on the **Jobs** tab has finished, you can start reviewing the conversations.

Step 4: Review and export conversations in a review set

After the content has been processed and added to the review set, you can start reviewing the data in the review set. The review capabilities are different depending on whether the content was added to a standard review set or a conversation review set.

Reviewing conversations in a standard review set

In a standard review set, messages are processed and displayed as individual items, similar to how they're stored in a mailbox folder. In this workflow, each message is processed as a separate item. As a result, the threaded summary and export options aren't available in a standard review set.

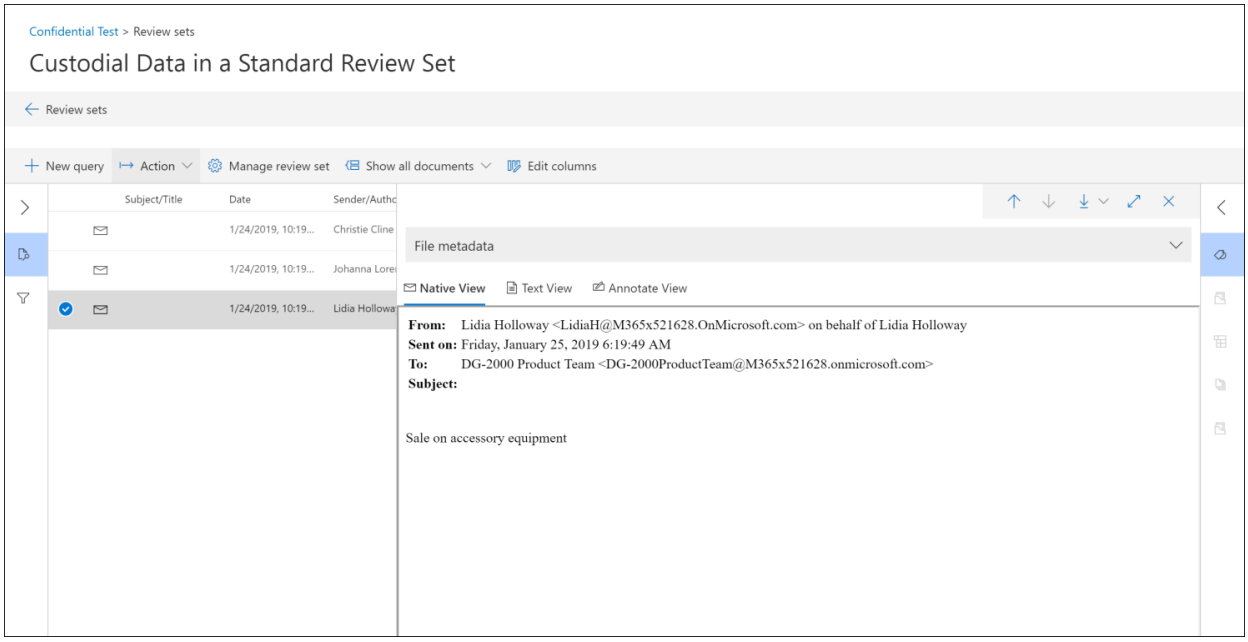

Reviewing conversations in a conversation review set

In a conversation review set, individual messages are threaded together and presented as conversations. This lets you review and export contextual conversations.

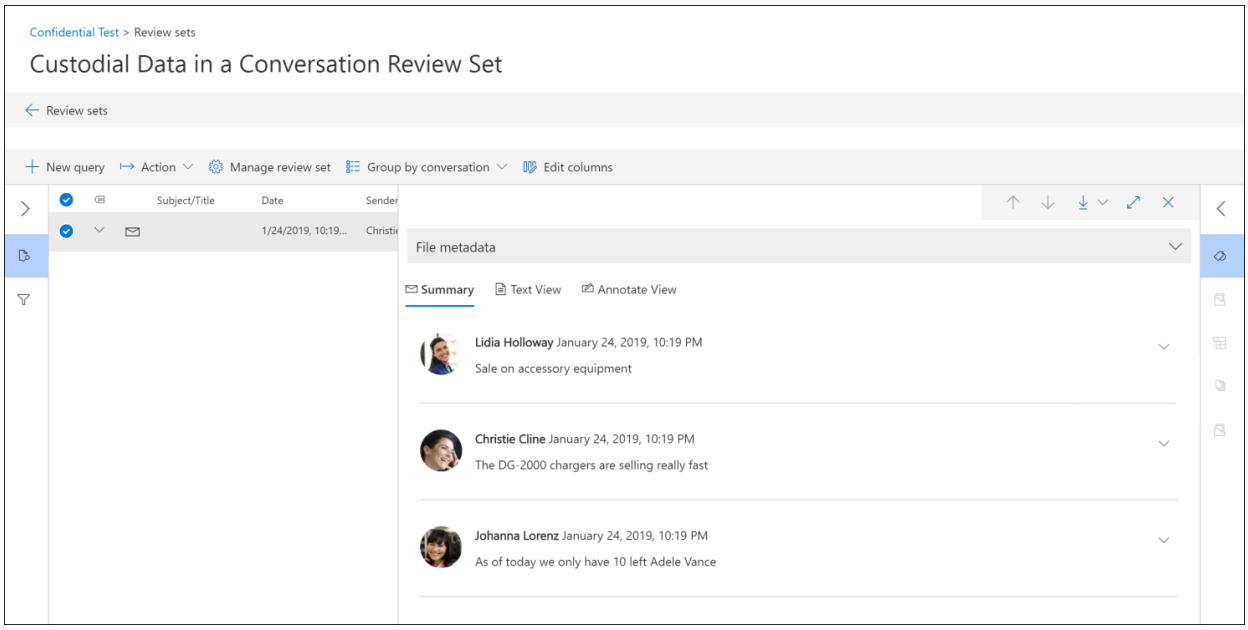

The following sections describe reviewing and exporting conversations in a conversation review set.

Reviewing conversations

In a conversation review set, you can use the following options to facilitate the review process.

- **Group by conversation:** Groups messages within the same conversation together to help users simplify and expedite their review process.
- **Summary view:** Displays the threaded conversation. In this view, you can see the entire conversation and also access the metadata for each individual message.
  - View metadata for individual messages
  - Download individual messages
- **Text view:** Provides the extracted text for the entire conversation.
- **Annotate view:** Lets you markup a threaded view of the conversation. All messages in the conversation share the same annotated document.
- **Tagging:** When viewing conversations in a review set, you can view and apply tags by clicking **Tagging panel** in the Coding panel.
- **Rerun conversation conversion:** When messages are added to a conversation review set, a conversion job is automatically run to create the threaded summary and annotate views. If the Conversation Reconstruction job fails, you can rerun this job by clicking **Action > Create conversation PDFs** in the review set.

Exporting conversations

In a conversation review set, you can set the following options to export conversations:

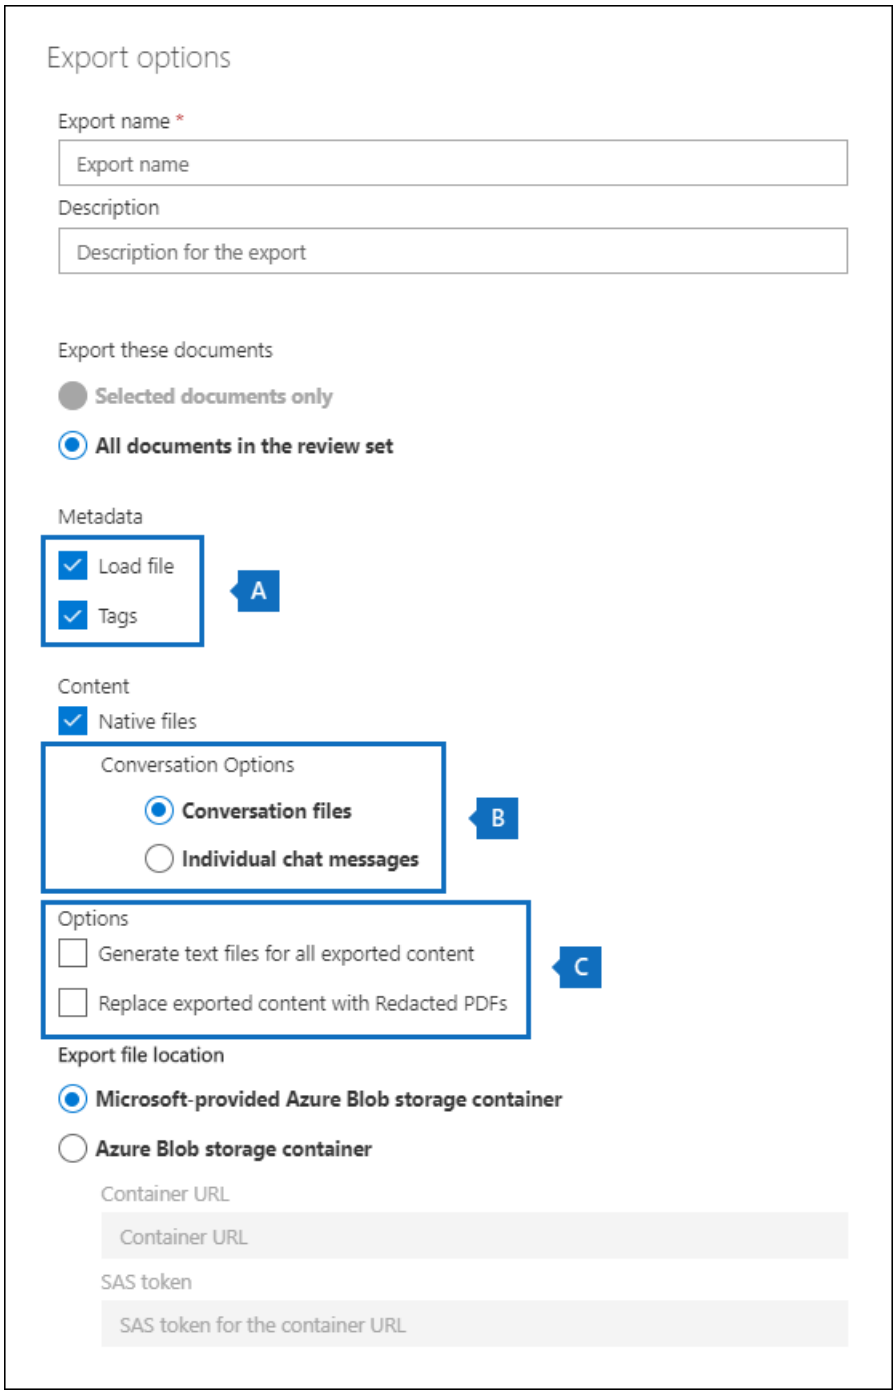

a. Metadata options

- **Load file:** Metadata is included for each individual message, email, and document. There is one row for each message in a conversation.
- **Tags:** Tags from your review process are included in the metadata file. Messages in a conversation share the same tags.

b. Conversation options

- **Conversation files:** When you export conversation files, the annotated view is converted to a PDF file and downloaded to the export folder. Messages in one conversation file point to the PDF version of the same conversation file.
- **Individual chat messages:** When you export individual messages, each unique message in the conversation is exported as a standalone item. The file is exported in the same format that it was saved as in the mailbox. For a specific conversation, you receive multiple .msg files.

> **NOTE**
>
> If you applied annotations to the conversation file, these annotations won't be transferred to the individual messages.

c. Other options

- **Generate text files for all exported content:** Generates a text file for each conversation exported from the review set.
- **Replace exported content with redacted PDFs:** If redacted conversation files are generated during the review process, then these files are available during export. You can decided whether to export only the native files (by not selecting this option) or to replace the native files with the redacted versions of the native files (by selecting this option), which are exported as PDF files.

More information

To learn more about how to review case data in Advanced eDiscovery, see the following articles:

- View case data
- Analyze case data
- Export case data

CJK language support for Advanced eDiscovery

11/2/2020 • 2 minutes to read • Edit Online

Advanced eDiscovery supports double-byte character set languages (these include Simplified Chinese, Traditional Chinese, Japanese, and Korean, which are collectively known as *CJK* languages) for the following advanced scenarios in a review set:

- When you query the data in a review set.
- When you tag documents in a review set.
- When you analyze case data in a review set by using near duplicate detection, email threading, and themes analytics.

Frequently asked questions

How do I create a search to collect items that contains CJK characters?

You can use CJK characters for keyword searches, keyword queries and search conditions when searching for content in Advanced eDiscovery. Searching for CJK characters is also supported when searching for content in Core eDiscovery and Content Search.

We provide CJK support for all search operators and search conditions, including the boolean operators **AND**, **OR**, **NOT**, and **NEAR**.

If you're certain that content locations or items contain CJK characters, but searches aren't returning any results, click the query language-country/region icon and select the corresponding language-country culture code value for the search. The default language/region is neutral.

Can I search for multiple languages at once?

It depends on your search scenario.

- When you query data in a review set in Advanced eDiscovery, you can search for multiple languages.
- When you create a search to collect data, create a separate search for each language you're targeting. For example, if you are searching for a document that contains both Chinese and Korean, select Chinese for your first query and select Korean for your second query.

I don't see the query language-country/region icon to select a language for queries in a review set. How can I specify a query language in a review set search?

For review set queries, you don't need to specify a document language. Advanced eDiscovery automatically detects document languages when you add content to a review set. This helps you optimize your query results in a review set.

Can I see detected languages in file metadata?

No, you can't see detected languages in file metadata.

Can I filter by document languages in a review set?

No, you can't filter, sort, or search by document languages in a review set.

Will this CJK release for review set scenarios affect any of my existing searches and review sets?

No, none of your existing searches and review sets will change. You don't need to reindex existing data, and search results for English text will be the same.

How do I change my display language to Chinese, Japanese, or Korean?

For information about how to change display language and time zone, see How to set language and region settings for Office 365.

Known issues

- OCR doesn't support CJK characters from image files
- Email files (such as *.eml and *.msg) in Annotate view aren't supported for CJK languages.
- Search hit highlighting in Text view isn't supported for CJK languages.
- The Relevance module used to analyze data doesn't support CJK languages.
- Query-based holds aren't supported for CJK languages.

Assign eDiscovery permissions in the Security & Compliance Center

2/18/2021 • 10 minutes to read • Edit Online

If you want people to use any of the eDiscovery-related tools in the Security & Compliance Center in Office 365 or the Microsoft 365 compliance center, you have to assign them the appropriate permissions. The easiest way to do this is to add the person the appropriate role group on the **Permissions** page in the Security & Compliance Center. This topic describes the permissions required to perform eDiscovery- and Content Search-related tasks using the Security & Compliance Center.

The primary eDiscovery-related role group in Security & Compliance Center is called **eDiscovery Manager**. There are two subgroups within this role group.

- **eDiscovery Managers** - An eDiscovery Manager can use the Content Search tool in the Security & Compliance Center to search content locations in the organization, and perform various search-related actions such as preview and export search results. Members can also create and manage cases in Core eDiscovery and Advanced eDiscovery, add and remove members to a case, create case holds, run searches associated with a case, and access case data. eDiscovery Managers can only access and manage the cases they create. They can't access or manage cases created by other eDiscovery Managers.
- **eDiscovery Administrators** - An eDiscovery Administrator is a member of the eDiscovery Manager role group, and can perform the same content search and case management-related tasks that an eDiscovery Manager can perform. Additionally, an eDiscovery Administrator can:
  - Access all cases that are listed on the **eDiscovery** and **Advanced eDiscovery** pages in the Security & Compliance Center.
  - Access case data in Advanced eDiscovery for any case in the organization.
  - Manage any eDiscovery case after they add themselves as a member of the case.

  For reasons why you might want eDiscovery Administrators in your organization, see More information.

> **NOTE**
> To analyze a user's data using Advanced eDiscovery, the user (the custodian of the data) must be assigned an Office 365 E5 or Microsoft 365 E5 license. Alternatively, users with an Office 365 E1 or a Office 365 or Microsoft 365 E3 license can be assigned an Microsoft 365 E5 Compliance or Microsoft 365 eDiscovery and Audit add-on license. Administrators, compliance officers, or legal personnel who are assigned to cases as members and use Advanced eDiscovery to collect, view, and analyze data don't need an E5 license. For more information about Advanced eDiscovery licensing, see Get started with Advanced eDiscovery.

Confirm your roles

- You have to be a member of the Organization Management role group or be assigned the Role Management role to assign eDiscovery permissions in the Security & Compliance Center.
- You can use the Add-RoleGroupMember cmdlet in Security & Compliance Center PowerShell to add a mail-enabled security group as a member of the eDiscovery Managers subgroup in the eDiscovery Manager role group. However, you can't add a mail-enabled security group to the eDiscovery Administrators subgroup. For details, see More information.

Assign eDiscovery permissions in the Security & Compliance Center

1. Go to https://protection.office.com.
2. Sign in using your work or school account.
3. In the left pane of the security and compliance center, select **Permissions**, and then select the checkbox next to **eDiscovery Manager**.
4. On the **eDiscovery Manager** flyout page, do one of the following based on the eDiscovery permissions that you want to assign.

   To make a user an eDiscovery Manager: Next to **eDiscovery Manager**, select **Edit**. In the **Choose eDiscovery Manager** section, select the **Choose eDiscovery Manager** hyperlink, and then select **+ Add**. Select the user (or users) you want to add as an eDiscovery manager, and then select **Add**. When you're finished adding users, select **Done**. Then, on the **Editing Choose eDiscovery Manager** flyout page, select **Save** to save the changes to the eDiscovery Manager membership.

   To make a user an eDiscovery Administrator: Next to **eDiscovery Manager**, select **Edit**. In the **Choose eDiscovery Administrator** section, Under **eDiscovery Administrators**, select **Choose eDiscovery Administrator**, select **Edit**, and then select **+ Add**. Select the user (or users) you want to add as an **eDiscovery Administrator**, and then **Add**. When you're finished adding users, select **Done**. Then, on the **Editing Choose eDiscovery Administrator** flyout page, select **Save** to save the changes to the eDiscovery Administrator membership.

> **NOTE**
>
> You can also use the **Add-eDiscoveryCaseAdmin** cmdlet to make a user an eDiscovery Administrator. However, the user must be assigned the Case Management role before you can use this cmdlet to make them an eDiscovery Administrator. For more information, see Add-eDiscoveryCaseAdmin.

On the **Permissions** page in the Security & Compliance Center, you can also assign users eDiscovery-related permissions by adding them to the Compliance Administrator, Organization Management, and Reviewer role groups. For a description of the eDiscovery-related RBAC roles assigned to each of these role groups, see RBAC roles related to eDiscovery.

RBAC roles related to eDiscovery

The following table lists the eDiscovery-related RBAC roles in the Security & Compliance Center, and indicates the built-in role groups that each role is assigned to by default.

ROLE	COMPLIANCE ADMINISTRATOR	EDISCOVERY MANAGER & ADMINISTRATOR	ORGANIZATION MANAGEMENT	REVIEWER
Case Management	✓	✓	✓	
Communication		✓		
Compliance Search	✓	✓	✓	
Custodian		✓		
Export		✓		

ROLE	COMPLIANCE ADMINISTRATOR	EDISCOVERY MANAGER & ADMINISTRATOR	ORGANIZATION MANAGEMENT	REVIEWER
Hold	✓	✓	✓	
Preview		✓		
Review		✓		✓
RMS Decrypt		✓		
Search And Purge			✓	

The following sections describe each of the eDiscovery-related RBAC roles listed in the previous table.

Case Management

This role lets users create, edit, delete, and control access to Core eDiscovery and Advanced eDiscovery cases in the Security & Compliance Center. As previously explained, a user must be assigned the Case Management role before you can use the **Add-eDiscoveryCaseAdmin** cmdlet to make them an eDiscovery Administrator.

For more information, see:

- Get started with Core eDiscovery
- Get started with Advanced eDiscovery

Communication

This role lets users manage all communications with the custodians identified in an Advanced eDiscovery case. This includes creating hold notifications, hold reminders, and escalations to management. The user can also track custodian acknowledgment of hold notifications and manage access to the custodian portal that is used by each custodian to track communications for the cases where they were identified as a custodian.

For more information, see Work with communications in Advanced eDiscovery.

Compliance Search

This role lets users run the Content Search tool in the Security & Compliance Center to search mailboxes and public folders, SharePoint Online sites, OneDrive for Business sites, Skype for Business conversations, Microsoft 365 groups, and Microsoft Teams, and Yammer groups. This role allows a user to get an estimate of the search results and create export reports, but additional roles are needed to initiate content search actions such as previewing, exporting, or deleting search results.

Users who are assigned the Compliance Search role but don't have the Preview role can preview the results of a search in which the preview action has been initiated by a user who is assigned the Preview role. The user without the Preview role can preview results for up to two weeks after the initial preview action was created.

Similarly, users who are assigned the Compliance Search role but don't have the Export role can download the results of a search in which the export action was initiated by a user who is assigned the Export role. The user without the Export role can download the results of a search for up to two weeks after the initial export action was created. After that, they can't download the results unless someone with the Export role restarts the export.

For more information, see Content search in Office 365.

Custodian

This role lets users identify and manage custodians for Advanced eDiscovery cases and use the information

from Azure Active Directory and other sources to find data sources associated with custodians. The user can associate other data sources such as mailboxes, SharePoint sites, and Teams with custodians in a case. The user can also place a legal hold on the data sources associated with custodians to preserve content in the context of a case.

For more information, see Work with custodians in Advanced eDiscovery.

Export

The role lets users export the results of a Content Search to a local computer. It also lets them prepare search results for analysis in Advanced eDiscovery.

For more information about exporting search results, see Export search results from Security & Compliance Center.

Hold

This role lets users place content on hold in mailboxes, public folders, sites, Skype for Business conversations, and Microsoft 365 groups. When content is on hold, content owners can still modify or delete the original content, but the content will be preserved until the hold is removed or until the hold duration expires.

For more information about holds, see:

- Create a hold in Core eDiscovery
- Create a hold in Advanced eDiscovery

Preview

This role lets users view a list of items that were returned from a Content Search. They can also open and view each item from the list to view its contents.

Review

This role lets users access review sets in Advanced eDiscovery. Users who are assigned this role can see and open the list of cases on the **eDiscovery > Advanced** page in the Microsoft 365 compliance center that they're members of. After the user accesses an Advanced eDiscovery case, they can select **Review sets** to access case data. This role doesn't allow the user to preview the results of a collection search that's associated with the case or do other search or case management tasks. Users with this role can only access the data in a review set.

RMS Decrypt

This role lets users view rights-protected email messages when previewing search results and export decrypted rights-protected email messages. This role also lets users view (and export) a file that's encrypted with a Microsoft encryption technology when the encrypted file is attached to an email message that's included in the results of an eDiscovery search. Additionally, this role lets users review and query encrypted email attachments that are added to a review set in Advanced eDiscovery. For more information about decryption in eDiscovery, see Decryption in Microsoft 365 eDiscovery tools.

Search And Purge

This role lets users perform bulk removal of data matching the criteria of a content search. For more information, see Search for and delete email messages in your organization.

More information

- **Why create an eDiscovery Administrator?** As previously explained, an eDiscovery Administrator is member of the eDiscovery Manager role group who can view and access all eDiscovery cases in your organization. This ability to access all the eDiscovery cases has two important purposes:
    - If a person who is the only member of an eDiscovery case leaves your organization, no one (including members of the Organization Management role group or another member of the

eDiscovery Manager role group) can access that eDiscovery case because they aren't a member of a case. In this situation, there would be no way to access the data in the case. But because an eDiscovery Administrator can access all eDiscovery cases in the organization, they can view the case and add themselves or another eDiscovery manager as a member of the case.

  - Because an eDiscovery Administrator can view and access all Core eDiscovery and Advanced eDiscovery cases, they can audit and oversee all cases and associated compliance searches. This can help to prevent any misuse of compliance searches or eDiscovery cases. And because eDiscovery Administrators can access potentially sensitive information in the results of a compliance search, you should limit the number of people who are eDiscovery Administrators.

- **Can I add a group as a member of the eDiscovery Manager role group?** As previously explained, you can add a mail-enabled security group as a member of the eDiscovery Managers subgroup in the eDiscovery Manager role group by using the **Add-RoleGroupMember** cmdlet in Security & Compliance Center PowerShell. For example, you can run the following command to add a mail-enabled security group to the eDiscovery Manager role group.

  ```
  Add-RoleGroupMember "eDiscovery Manager" -Member <name of security group>
  ```

  Exchange distribution groups and Microsoft 365 Groups aren't supported. You must use a mail-enabled security group, which you can create in Exchange Online PowerShell by running `New-DistributionGroup -Type Security`. You can also create a mail-enabled security group (and add members) in the Exchange admin center or in the Microsoft 365 admin center. It might take up to 60 minutes after you create it for a new mail-enabled security to be available to add to the eDiscovery Managers role group.

  Also as previously stated, you can't make a mail-enabled security group an eDiscovery Administrator by using the **Add-eDiscoveryCaseAdmin** cmdlet in Security & Compliance Center PowerShell. You can only add individual users as eDiscovery Administrators.

  You also can't add a mail-enabled security group as a member of a case.

Keyword queries and search conditions for Content Search and eDiscovery

2/18/2021 • 34 minutes to read • Edit Online

This topic describes the email and document properties that you can search for in email items in Exchange Online and documents stored on SharePoint and OneDrive for Business sites by using the Content Search feature in the Microsoft 365 compliance center. You can also use the ***-ComplianceSearch** cmdlets in Security & Compliance Center PowerShell to search for these properties. The topic also describes:

- Using Boolean search operators, search conditions, and other search query techniques to refine your search results.
- Searching for sensitive data types and custom sensitive data types in SharePoint and OneDrive for Business.
- Searching for site content that's shared with users outside of your organization

For step-by-step instructions on how to create a Content Search, see Content Search.

> **NOTE**
> Content Search in the Microsoft 365 compliance center and the corresponding ***-ComplianceSearch** cmdlets in Security & Compliance Center PowerShell use the Keyword Query Language (KQL). For more detailed information, see Keyword Query Language syntax reference.

Searchable email properties

The following table lists email message properties that can be searched by using the Content Search feature in the Microsoft 365 compliance center or by using the **New-ComplianceSearch** or the **Set-ComplianceSearch** cmdlet. The table includes an example of the *property:value* syntax for each property and a description of the search results returned by the examples. You can type these `property:value` pairs in the keywords box for a Content Search.

> **NOTE**
> When searching email properties, it's not possible to search for items in which the specified property is empty or blank. For example, using the *property:value* pair of **subject:""** to search for email messages with an empty subject line will return zero results. This also applies when searching site and contact properties.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLES	SEARCH RESULTS RETURNED BY THE EXAMPLES

AttachmentNames	The names of files attached to an email message.	`attachmentnames:annualreport.ppt`
		`attachmentnames:annual*`
		`attachmentnames:.pptx`
			Messages that have an attached file named annualreport.ppt. In the second example, using the wildcard returns messages with the word "annual" in the file name of an attachment. The third example returns all attachments with the pptx file extension.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLES	SEARCH RESULTS RETURNED BY THE EXAMPLES

Bcc	The Bcc field of an email message.[1]	`bcc:pilarp@contoso.com`
		`bcc:pilarp`
		`bcc:"Pilar Pinilla"`
			All examples return messages with Pilar Pinilla included in the Bcc field.

Category	The categories to search. Categories can be defined by users by using Outlook or Outlook on the web (formerly known as Outlook Web App). The possible values are: blue green orange purple red yellow	`category:"Red Category"`	Messages that have been assigned the red category in the source mailboxes.
Cc	The Cc field of an email message.[1]	`cc:pilarp@contoso.com` `cc:"Pilar Pinilla"`	In both examples, messages with Pilar Pinilla specified in the Cc field.
Folderid	The folder ID (GUID) of a specific mailbox folder. If you use this property, be sure to search the mailbox that the specified folder is located in. Only the specified folder will be searched. Any subfolders in the folder won't be searched. To search subfolders, you need to use the Folderid property for the subfolder you want to search. For more information about searching for the Folderid property and using a script to obtain the folder IDs for a specific mailbox, see Use Content Search for targeted collections.	`folderid:4D6DD7F943C29041A65787E30F02AD1F00000000013A0000` `folderid:2370FB455F82FC44BE31397F47B632A70000000001160000 AND participants:garthf@contoso.com`	The first example returns all items in the specified mailbox folder. The second example returns all items in the specified mailbox folder that were sent or received by garthf@contoso.com.
From	The sender of an email message.[1]	`from:pilarp@contoso.com` `from:contoso.com`	Messages sent by the specified user or sent from a specified domain.
HasAttachment	Indicates whether a message has an attachment. Use the values **true** or **false**.	`from:pilar@contoso.com AND hasattachment:true`	Messages sent by the specified user that have attachments.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLES	SEARCH RESULTS RETURNED BY THE EXAMPLES

Importance	The importance of an email message, which a sender can specify when sending a message. By default, messages are sent with normal importance, unless the sender sets the importance as **high** or **low**.	`importance:high`
		`importance:medium`
		`importance:low`
			Messages that are marked as high importance, medium importance, or low importance.

IsRead	Indicates whether messages have been read. Use the values **true** or **false**.	`isread:true` `isread:false`	The first example returns messages with the IsRead property set to **True**. The second example returns messages with the IsRead property set to **False**.
ItemClass	Use this property to search specific third-party data types that your organization imported to Office 365. Use the following syntax for this property: `itemclass:ipm.externaldata.<third-party data type>*`	`itemclass:ipm.externaldata.Facebook* AND subject:contoso` `itemclass:ipm.externaldata.Twitter* AND from:"Ann Beebe" AND "Northwind Traders"`	The first example returns Facebook items that contain the word "contoso" in the Subject property. The second example returns Twitter items that were posted by Ann Beebe and that contain the keyword phrase "Northwind Traders". For a complete list of values to use for third-party data types for the ItemClass property, see Use Content Search to search third-party data that was imported to Office 365.

Kind	The type of email message to search for. Possible values:
	contacts
	docs
	email
	externaldata
	faxes
	im
	journals
	meetings
	microsoftteams (returns items from chats, meetings, and calls in Microsoft Teams)
	notes
	posts
	rssfeeds
	tasks
	voicemail
		`kind:email`
		`kind:email OR kind:im OR kind:voicemail`
		`kind:externaldata`
			The first example returns email messages that meet the search criteria. The second example returns email messages, instant messaging conversations (including Skype for Business conversations and chats in Microsoft Teams), and voice messages that meet the search criteria. The third example returns items that were imported to mailboxes in Microsoft 365 from third-party data sources, such as Twitter, Facebook, and Cisco Jabber, that meet the search criteria. For more information, see Archiving third-party data in Office 365.

Participants	All the people fields in an email message. These fields are From, To, Cc, and Bcc.[1]	`participants:garthf@contoso.com` `participants:contoso.com`	Messages sent by or sent to garthf@contoso.com. The second example returns all messages sent by or sent to a user in the contoso.com domain.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLES	SEARCH RESULTS RETURNED BY THE EXAMPLES
Received	The date that an email message was received by a recipient.	`received:04/15/2016` `received>=01/01/2016 AND received<=03/31/2016`	Messages that were received on April 15, 2016. The second example returns all messages received between January 1, 2016 and March 31, 2016.
Recipients	All recipient fields in an email message. These fields are To, Cc, and Bcc.[1]	`recipients:garthf@contoso.com` `recipients:contoso.com`	Messages sent to garthf@contoso.com. The second example returns messages sent to any recipient in the contoso.com domain.
Sent	The date that an email message was sent by the sender.	`sent:07/01/2016` `sent>=06/01/2016 AND sent<=07/01/2016`	Messages that were sent on the specified date or sent within the specified date range.
Size	The size of an item, in bytes.	`size>26214400` `size:1..1048567`	Messages larger than 25?? MB. The second example returns messages from 1 through 1,048,567 bytes (1 MB) in size.
Subject	The text in the subject line of an email message. **Note:** When you use the Subject property in a query, the search returns all messages in which the subject line contains the text you're searching for. In other words, the query doesn't return only those messages that have an exact match. For example, if you search for `subject:"Quarterly Financials"`, your results will include messages with the subject "Quarterly Financials 2018".	`subject:"Quarterly Financials"` `subject:northwind`	Messages that contain the phrase "Quarterly Financials" anywhere in the text of the subject line. The second example returns all messages that contain the word northwind in the subject line.

To	The To field of an email message.[1]	`to:annb@contoso.com`
		`to:annb`
		`to:"Ann Beebe"`
			All examples return messages where Ann Beebe is specified in the To: line.

NOTE

[1] For the value of a recipient property, you can use email address (also called *user principal name* or UPN), display name, or alias to specify a user. For example, you can use annb@contoso.com, annb, or "Ann Beebe" to specify the user Ann Beebe.

Recipient expansion

When searching any of the recipient properties (From, To, Cc, Bcc, Participants, and Recipients), Microsoft 365 attempts to expand the identity of each user by looking them up in Azure Active Directory (Azure AD). If the user

is found in Azure AD, the query is expanded to include the user's email address (or UPN), alias, display name, and LegacyExchangeDN. For example, a query such as `participants:ronnie@contoso.com` expands to

```
participants:ronnie@contoso.com OR participants:ronnie OR participants:"Ronald Nelson" OR participants:"<LegacyExchangeDN>"
```

.

To prevent recipient expansion, add a wild card character (asterisk) to the end of the email address and use a reduced domain name; for example, `participants:"ronnie@contoso*"` Be sure to surround the email address with double quotation marks.

However, be aware that preventing recipient expansion in the search query may result in relevant items not being returned in the search results. Email messages in Exchange can be saved with different text formats in the recipient fields. Recipient expansion is intended to help mitigate this fact by returning messages that may contain different text formats. So preventing recipient expansion may result in the search query not returning all items that may be relevant to your investigation.

> **NOTE**
>
> If you need to review or reduce the items returned by a search query due to recipient expansion, consider using Advanced eDiscovery. You can search for messages (taking advantage of recipient expansion), add them to a review set, and then use review set queries or filters to review or narrow the results. For more information, see Collect data for a case and Query the data in a review set.

Searchable site properties

The following table lists some of the SharePoint and OneDrive for Business properties that can be searched by using the Content Search feature in the Security & Compliance Center or by using the **New-ComplianceSearch** or the **Set-ComplianceSearch** cmdlet. The table includes an example of the *property:value* syntax for each property and a description of the search results returned by the examples.

For a complete list of SharePoint properties that can be searched, see Overview of crawled and managed properties in SharePoint. Properties marked with a **Yes** in the **Queryable** column can be searched.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLE	SEARCH RESULTS RETURNED BY THE EXAMPLES
Author	The author field from Office documents, which persists if a document is copied. For example, if a user creates a document and the emails it to someone else who then uploads it to SharePoint, the document will still retain the original author. Be sure to use the user's display name for this property.	`author:"Garth Fort"`	All documents that are authored by Garth Fort.
ContentType	The SharePoint content type of an item, such as Item, Document, or Video.	`contenttype:document`	All documents would be returned.
Created	The date that an item is created.	`created>=06/01/2016`	All items created on or after June 1, 2016.
CreatedBy	The person that created or uploaded an item. Be sure to use the user's display name for this property.	`createdby:"Garth Fort"`	All items created or uploaded by Garth Fort.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLE	SEARCH RESULTS RETURNED BY THE EXAMPLES
DetectedLanguage	The language of an item.	`detectedlanguage:english`	All items in English.
DocumentLink	The path (URL) of a specific folder on a SharePoint or OneDrive for Business site. If you use this property, be sure to search the site that the specified folder is located in. To return items located in subfolders of the folder that you specify for the documentlink property, you have to add /* to the URL of the specified folder; for example, `documentlink: "https://contoso.sharepoint.com/Shared Documents/*"` For more information about searching for the documentlink property and using a script to obtain the documentlink URLs for folders on a specific site, see Use Content Search for targeted collections.	`documentlink:"https://contoso-my.sharepoint.com/personal/garthf_contoso_com/Documents/Priva` `documentlink:"https://contoso-my.sharepoint.com/personal/garthf_contoso_com/Documents/Share with Everyone/*" AND filename:confidential`	The first example returns all items in the specified OneDrive for Business folder. The second example returns documents in the specified site folder (and all subfolders) that contain the word "confidential" in the file name.
FileExtension	The extension of a file; for example, docx, one, pptx, or xlsx.	`fileextension:xlsx`	All Excel files (Excel 2007 and later)
FileName	The name of a file.	`filename:"marketing plan"` `filename:estimate`	The first example returns files with the exact phrase "marketing plan" in the title. The second example returns files with the word "estimate" in the file name.
LastModifiedTime	The date that an item was last changed.	`lastmodifiedtime>=05/01/2016` `lastmodifiedtime>=05/10/2016 AND lastmodifiedtime<=06/1/2016`	The first example returns items that were changed on or after May 1, 2016. The second example returns items changed between May 1, 2016 and June 1, 2016.
ModifiedBy	The person who last changed an item. Be sure to use the user's display name for this property.	`modifiedby:"Garth Fort"`	All items that were last changed by Garth Fort.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLE	SEARCH RESULTS RETURNED BY THE EXAMPLES
Path	The path (URL) of a specific site in a SharePoint or OneDrive for Business site. To return items located in folders in the site that you specify for the path property, you have to add /* to the URL of the specified site; for example, `path: "https://contoso.sharepoint.com/Shared Documents/*"` **Note:** Using the `Path` property to search OneDrive locations won't return media files, such as .png, .tiff, or .wav files, in the search results. Use a different site property in your search query to search for media files in OneDrive folders.	`path:"https://contoso-my.sharepoint.com/personal/garthf_contoso_com/"` `path:"https://contoso-my.sharepoint.com/personal/garthf_contoso_com/*" AND filename:confidential`	The first example returns all items in the specified OneDrive for Business site. The second example returns documents in the specified site (and folders in the site) that contain the word "confidential" in the file name.
SharedWithUsersOWSUser	Documents that have been shared with the specified user and displayed on the **Shared with me** page in the user's OneDrive for Business site. These are documents that have been explicitly shared with the specified user by other people in your organization. When you export documents that match a search query that uses the SharedWithUsersOWSUser property, the documents are exported from the original content location of the person who shared the document with the specified user. For more information, see Searching for site content shared within your organization.	`sharedwithusersowsuser:garthf` `sharedwithusersowsuser:"garthf@contoso.com"`	Both examples return all internal documents that have been explicitly shared with Garth Fort and that appear on the **Shared with me** page in Garth Fort's OneDrive for Business account.
Site	The URL of a site or group of sites in your organization.	`site:"https://contoso-my.sharepoint.com"` `site:"https://contoso.sharepoint.com/sites/teams"`	The first example returns items from the OneDrive for Business sites for all users in the organization. The second example returns items from all team sites.
Size	The size of an item, in bytes.	`size>=1` `size:1..10000`	The first example returns items larger than 1 byte. The second example returns items from 1 through 10,000 bytes in size.

PROPERTY	PROPERTY DESCRIPTION	EXAMPLE	SEARCH RESULTS RETURNED BY THE EXAMPLES
Title	The title of the document. The Title property is metadata that's specified in Microsoft Office documents. It's different from the file name of the document.	`title:"communication plan"`	Any document that contains the phrase "communication plan" in the Title metadata property of an Office document.

Searchable contact properties

The following table lists the contact properties that are indexed and that you can search for using Content Search. These are the properties that are available for users to configure for the contacts (also called personal contacts) that are located in the personal address book of a user's mailbox. To search for contacts, you can select the mailboxes to search and then use one or more contact properties in the keyword query.

> **TIP**
> To search for values that contain spaces or special characters, use double quotation marks (" ") to contain the phrase; for example, `businessaddress:"123 Main Street"`.

PROPERTY	PROPERTY DESCRIPTION		
BusinessAddress	The address in the **Business Address** property. The property is also called the **Work** address on the contact properties page.		
BusinessPhone	The phone number in any of the **Business Phone** number properties.		
CompanyName	The name in the **Company** property.		
Department	The name in the **Department** property.		
DisplayName	The display name of the contact. This is the name in the **Full Name** property of the contact.		
EmailAddress	The address for any email address property for the contact. Users can add multiple email addresses for a contact. Using this property would return contacts that match any of the contact's email addresses.		

PROPERTY	PROPERTY DESCRIPTION		
FileAs	The **File as** property. This property is used to specify how the contact is listed in the user's contact list. For example, a contact could be listed as *FirstName,LastName* or *LastName,FirstName*.		
GivenName	The name in the **First Name** property.		
HomeAddress	The address in any of the **Home** address properties.		
HomePhone	The phone number in any of the **Home** phone number properties.		
IMAddress	The IM address property, which is typically an email address used for instant messaging.		
MiddleName	The name in the **Middle** name property.		
MobilePhone	The phone number in the **Mobile** phone number property.		
Nickname	The name in the **Nickname** property.		
OfficeLocation	The value in **Office** or **Office location** property.		
OtherAddress	The value for the **Other** address property.		
Surname	The name in the **Last** name property.		
Title	The title in the **Job title** property.		

Searchable sensitive data types

You can use eDiscovery search tools in the Microsoft 365 compliance center to search for sensitive data, such as credit card numbers or social security numbers, that is stored in documents on SharePoint and OneDrive for Business sites. You can do this by using the `SensitiveType` property and the name (or ID) of a sensitive information type in a keyword query. For example, the query `SensitiveType:"Credit Card Number"` returns documents that contain a credit card number. The query `SensitiveType:"U.S. Social Security Number (SSN)"` returns documents that contain a U.S. social security number.

To see a list of the sensitive information types that you can search for, go to **Data classifications** > **Sensitive info types** in the Microsoft 365 compliance center. Or you can use the **Get-DlpSensitiveInformationType**

cmdlet in Security & Compliance Center PowerShell to display a list of sensitive information types.

For more information about creating queries using the `SensitiveType` property, see Form a query to find sensitive data stored on sites.

Limitations for searching sensitive data types

- To search for custom sensitive information types, you have to specify the ID of the sensitive information type in the `SensitiveType` property. Using the name of a custom sensitive information type (as shown in the example for built-in sensitive information types in the previous section) will return no results. Use the **Publisher** column on the **Sensitive info types** page in the compliance center (or the **Publisher** property in PowerShell) to differentiate between built-in and custom sensitive information types. Built-in sensitive data types have a value of `Microsoft Corporation` for the **Publisher** property.

  To display the name and ID for the custom sensitive data types in your organization, run the following command in Security & Compliance Center PowerShell:

  ```
  Get-DlpSensitiveInformationType | Where-Object {$_.Publisher -ne "Microsoft Corporation"} | FT
  Name,Id
  ```

  Then you can use the ID in the `SensitiveType` search property to return documents that contain the custom sensitive data type; for example, `SensitiveType:7e13277e-6b04-3b68-94ed-1aeb9d47de37`

- You can't use sensitive information types and the `SensitiveType` search property to search for sensitive data at-rest in Exchange Online mailboxes. However, you can use data loss prevention (DLP) policies to protect sensitive email data in transit. For more information, see Overview of data loss prevention policies and Search for and find personal data.

Search operators

Boolean search operators, such as **AND**, **OR**, and **NOT**, help you define more-precise searches by including or excluding specific words in the search query. Other techniques, such as using property operators (such as `>=` or `..`), quotation marks, parentheses, and wildcards, help you refine a search query. The following table lists the operators that you can use to narrow or broaden search results.

OPERATOR	USAGE	DESCRIPTION	
AND	keyword1 AND keyword2	Returns items that include all of the specified keywords or `property:value` expressions. For example, `from:"Ann Beebe" AND subject:northwind` would return all messages sent by Ann Beebe that contained the word northwind in the subject line. [2]	

OPERATOR	USAGE	DESCRIPTION	
+	keyword1 + keyword2 + keyword3	Returns items that contain *either* `keyword2` or `keyword3` *and* that also contain `keyword1`. Therefore, this example is equivalent to the query `(keyword2 OR keyword3) AND keyword1`. The query `keyword1 + keyword2` (with a space after the **+** symbol) isn't the same as using the **AND** operator. This query would be equivalent to `"keyword1 + keyword2"` and return items with the exact phase `"keyword1 + keyword2"`.	
OR	keyword1 OR keyword2	Returns items that include one or more of the specified keywords or `property:value` expressions. [2]	
NOT	keyword1 NOT keyword2 NOT from:"Ann Beebe" NOT kind:im	Excludes items specified by a keyword or a `property:value` expression. In the second example excludes messages sent by Ann Beebe. The third example excludes any instant messaging conversations, such as Skype for Business conversations that are saved to the Conversation History mailbox folder. [2]	
-	keyword1 -keyword2	The same as the **NOT** operator. So this query returns items that contain `keyword1` and would exclude items that contain `keyword2`.	
NEAR	keyword1 NEAR(n) keyword2	Returns items with words that are near each other, where n equals the number of words apart. For example, `best NEAR(5) worst` returns any item where the word "worst" is within five words of "best". If no number is specified, the default distance is eight words. [2]	

OPERATOR	USAGE	DESCRIPTION	
:	property:value	The colon (:) in the `property:value` syntax specifies that the value of the property being searched for contains the specified value. For example, `recipients:garthf@contoso.com` returns any message sent to garthf@contoso.com.	
=	property=value	The same as the : operator.	
<	property<value	Denotes that the property being searched is less than the specified value. [1]	
>	property>value	Denotes that the property being searched is greater than the specified value.[1]	
<=	property<=value	Denotes that the property being searched is less than or equal to a specific value.[1]	
>=	property>=value	Denotes that the property being searched is greater than or equal to a specific value.[1]	
..	property:value1..value2	Denotes that the property being searched is greater than or equal to value1 and less than or equal to value2.[1]	
" "	"fair value" subject:"Quarterly Financials"	Use double quotation marks (" ") to search for an exact phrase or term in keyword and `property:value` search queries.	

OPERATOR	USAGE	DESCRIPTION	
*	cat* subject:set*	Prefix wildcard searches (where the asterisk is placed at the end of a word) match for zero or more characters in keywords or `property:value` queries. For example, `title:set*` returns documents that contain the word set, setup, and setting (and other words that start with "set") in the document title. **Note:** You can use only prefix wildcard searches; for example, **cat*** or **set***. Suffix searches (***cat**), infix searches (**c*t**), and substring searches (***cat***) are not supported.	
( )	(fair OR free) AND (from:contoso.com) (IPO OR initial) AND (stock OR shares) (quarterly financials)	Parentheses group together Boolean phrases, `property:value` items, and keywords. For example, `(quarterly financials)` returns items that contain the words quarterly and financials.	

NOTE

[1] Use this operator for properties that have date or numeric values.

[2] Boolean search operators must be uppercase; for example, **AND**. If you use a lowercase operator, such as **and**, it will be treated as a keyword in the search query.

Search conditions

You can add conditions to a search query to narrow a search and return a more refined set of results. Each condition adds a clause to the KQL search query that is created and run when you start the search.

Conditions for common properties

Conditions for mail properties

Conditions for document properties

Operators used with conditions

Guidelines for using conditions

Examples of using conditions in search queries

Conditions for common properties

Create a condition using common properties when searching mailboxes and sites in the same search. The following table lists the available properties to use when adding a condition.

CONDITION	DESCRIPTION
Date	For email, the date a message was received by a recipient or sent by the sender. For documents, the date a document was last modified.
Sender/Author	For email, the person who sent a message. For documents, the person cited in the author field from Office documents. You can type more than one name, separated by commas. Two or more values are logically connected by the **OR** operator.
Size (in bytes)	For both email and documents, the size of the item (in bytes).
Subject/Title	For email, the text in the subject line of a message. For documents, the title of the document. As previously explained, the Title property is metadata specified in Microsoft Office documents. You can type the name of more than one subject/title, separated by commas. Two or more values are logically connected by the **OR** operator.
Compliance label	For both email and documents, retention labels that have been assigned to messages and documents automatically by autolabel policies or retention labels that have been manually assigned by users. Retention labels are used to classify email and documents for information governance and enforce retention rules based on the settings defined by the label. You can type part of the retention label name and use a wildcard or type the complete label name. For more information about retention labels, see Learn about retention policies and retention labels.

Conditions for mail properties

Create a condition using mail properties when searching mailboxes or public folders. The following table lists the email properties that you can use for a condition. These properties are a subset of the email properties that were previously described. These descriptions are repeated for your convenience.

CONDITION	DESCRIPTION
Message kind	The message type to search. This is the same property as the Kind email property. Possible values: contacts docs email externaldata faxes im journals meetings microsoftteams notes posts rssfeeds tasks voicemail
Participants	All the people fields in an email message. These fields are From, To, Cc, and Bcc.

CONDITION	DESCRIPTION
Type	The message class property for an email item. This is the same property as the ItemClass email property. It's also a multi-value condition. So to select multiple message classes, hold the **CTRL** key and then click two or more message classes in the drop-down list that you want to add to the condition. Each message class that you select in the list will be logically connected by the **OR** operator in the corresponding search query. For a list of the message classes (and their corresponding message class ID) that are used by Exchange and that you can select in the **Message class** list, see Item Types and Message Classes.
Received	The date that an email message was received by a recipient. This is the same property as the Received email property.
Recipients	All recipient fields in an email message. These fields are To, Cc, and Bcc.
Sender	The sender of an email message.
Sent	The date that an email message was sent by the sender. This is the same property as the Sent email property.
Subject	The text in the subject line of an email message.
To	The recipient of an email message in the To field.

Conditions for document properties

Create a condition using document properties when searching for documents on SharePoint and OneDrive for Business sites. The following table lists the document properties that you can use for a condition. These properties are a subset of the site properties that were previously described. These descriptions are repeated for your convenience.

CONDITION	DESCRIPTION
Author	The author field from Office documents, which persists if a document is copied. For example, if a user creates a document and the emails it to someone else who then uploads it to SharePoint, the document will still retain the original author.
Title	The title of the document. The Title property is metadata that's specified in Office documents. It's different than the file name of the document.
Created	The date that a document is created.
Last modified	The date that a document was last changed.
File type	The extension of a file; for example, docx, one, pptx, or xlsx. This is the same property as the FileExtension site property.

Operators used with conditions

When you add a condition, you can select an operator that is relevant to type of property for the condition. The

following table describes the operators that are used with conditions and lists the equivalent that is used in the search query.

OPERATOR	QUERY EQUIVALENT	DESCRIPTION
After	`property>date`	Used with date conditions. Returns items that were sent, received, or modified after the specified date.
Before	`property<date`	Used with date conditions. Returns items that were sent, received, or modified before the specified date.
Between	`date..date`	Use with date and size conditions. When used with a date condition, returns items there were sent, received, or modified within the specified date range. When used with a size condition, returns items whose size is within the specified range.
Contains any of	`(property:value) OR (property:value)`	Used with conditions for properties that specify a string value. Returns items that contain any part of one or more specified string values.
Doesn't contain any of	`-property:value` `NOT property:value`	Used with conditions for properties that specify a string value. Returns items that don't contain any part of the specified string value.
Doesn't equal any of	`-property=value` `NOT property=value`	Used with conditions for properties that specify a string value. Returns items that don't contain the specific string.
Equals	`size=value`	Returns items that are equal to the specified size.[1]
Equals any of	`(property=value) OR (property=value)`	Used with conditions for properties that specify a string value. Returns items that are an exact match of one or more specified string values.
Greater	`size>value`	Returns items where the specified property is greater than the specified value.[1]
Greater or equal	`size>=value`	Returns items where the specified property is greater than or equal to the specified value.[1]
Less	`size<value`	Returns items that are greater than or equal to the specific value.[1]
Less or equal	`size<=value`	Returns items that are greater than or equal to the specific value.[1]
Not equal	`size<>value`	Returns items that don't equal the specified size.[1]

NOTE

[1] This operator is available only for conditions that use the Size property.

Guidelines for using conditions

Keep the following in mind when using search conditions.

- A condition is logically connected to the keyword query (specified in the keyword box) by the **AND** operator. That means that items have to satisfy both the keyword query and the condition to be included in the results. This is how conditions help to narrow your results.
- If you add two or more unique conditions to a search query (conditions that specify different properties), those conditions are logically connected by the **AND** operator. That means only items that satisfy all the conditions (in addition to any keyword query) are returned.
- If you add more than one condition for the same property, those conditions are logically connected by the **OR** operator. That means items that satisfy the keyword query and any one of the conditions are returned. So, groups of the same conditions are connected to each other by the **OR** operator and then sets of unique conditions are connected by the **AND** operator.
- If you add multiple values (separated by commas or semi-colons) to a single condition, those values are connected by the **OR** operator. That means items are returned if they contain any of the specified values for the property in the condition.
- The search query that is created by using the keywords box and conditions is displayed on the **Search** page, in the details pane for the selected search. In a query, everything to the right of the notation `(c:c)` indicates conditions that are added to the query.
- Conditions only add properties to the search query; the don't add operators. This is why the query displayed in the detail pane doesn't show operators to the right of the `(c:c)` notation. KQL adds the logical operators (according to the previously explained rules) when the executing the query.
- You can use the drag and drop control to resequence the order of conditions. Click on the control for a condition and move it up or down.
- As previously explained, some condition properties allow you to type multiple values. Each value is logically connected by the **OR** operator. This results in the same logic as having multiple instances of the same condition, where each has a single value. The following illustrations show an example of a single condition with multiple values and an example of multiple conditions (for the same property) with a single value. Both examples result in the same query:
  `(filetype:docx) OR (filetype:pptx) OR (filetype:xlsx)`

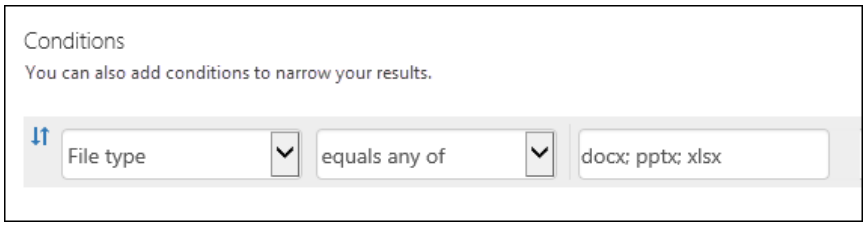

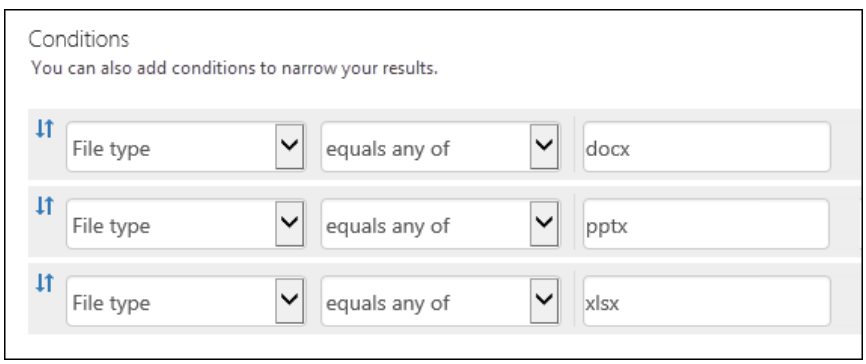

TIP

If a condition accepts multiple values, we recommend that you use a single condition and specify multiple values (separated by commas or semi-colons). This helps ensure the query logic that's applied is what you intend.

Examples of using conditions in search queries

The following examples show the GUI-based version of a search query with conditions, the search query syntax that is displayed in the details pane of the selected search (which is also returned by the **Get-ComplianceSearch** cmdlet), and the logic of the corresponding KQL query.

Example 1

This example returns documents on SharePoint and OneDrive for Business sites that contain a credit card number and were last modified before January 1, 2016.

GUI

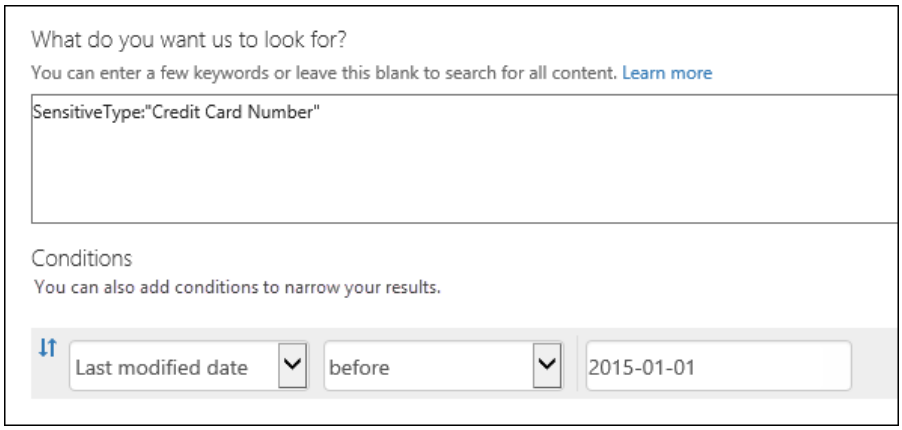

Search query syntax

```
SensitiveType:"Credit Card Number"(c:c)(lastmodifiedtime<2016-01-01)
```

Search query logic

```
SensitiveType:"Credit Card Number" AND (lastmodifiedtime<2016-01-01)
```

Example 2

This example returns email items or documents that contain the keyword "report", that were sent or created before April 1, 2105, and that contain the word "northwind" in the subject field of email messages or in the title property of documents. The query excludes Web pages that meet the other search criteria.

GUI

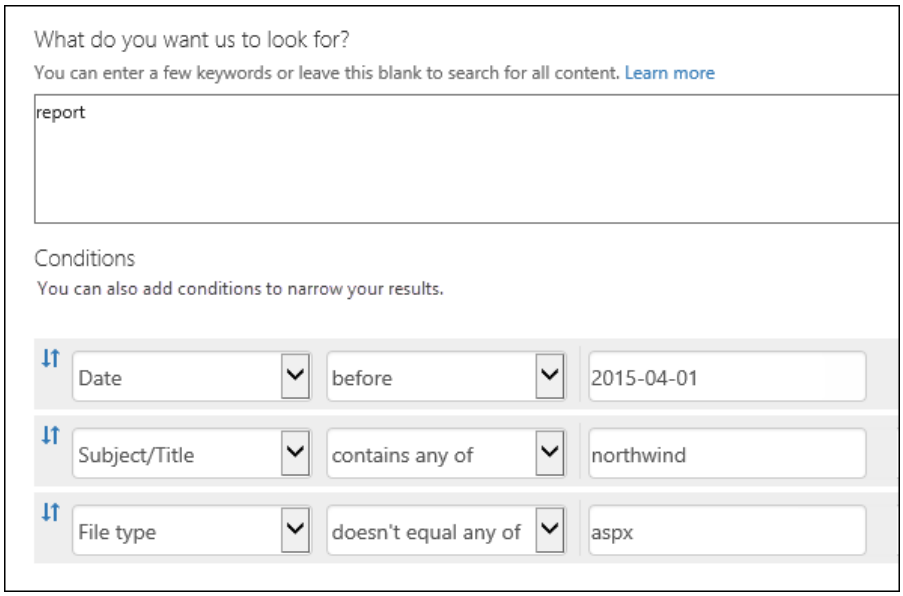

Search query syntax

```
report(c:c)(date<2016-04-01)(subjecttitle:"northwind")(-filetype:aspx)
```

Search query logic

```
report AND (date<2016-04-01) AND (subjecttitle:"northwind") NOT (filetype:aspx)
```

Example 3

This example returns email messages or calendar meetings that were sent between 12/1/2016 and 11/30/2016 and that contain words that start with "phone" or "smartphone".

GUI

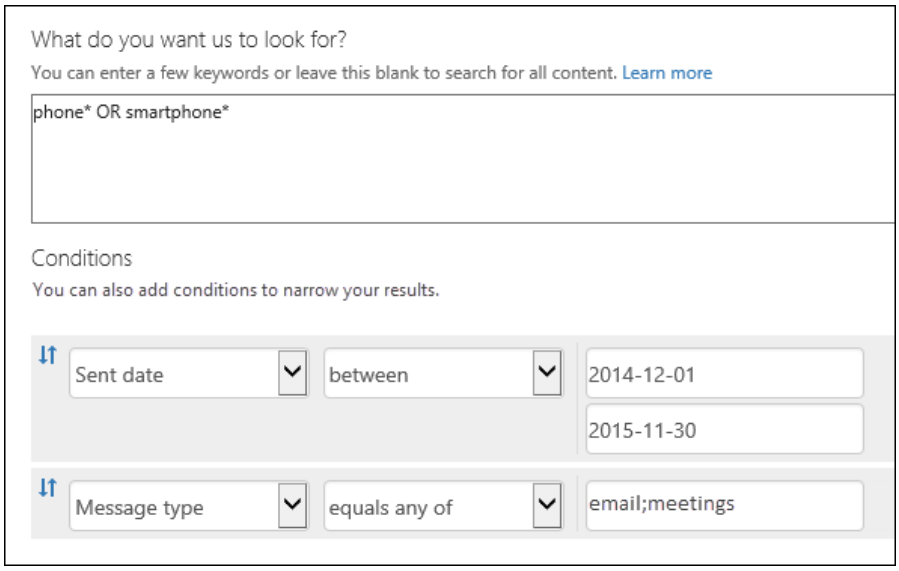

Search query syntax

```
phone* OR smartphone*(c:c)(sent=2016-12-01..2016-11-30)(kind="email")(kind="meetings")
```

Search query logic

```
phone* OR smartphone* AND (sent=2016-12-01..2016-11-30) AND ((kind="email") OR (kind="meetings"))
```

Special characters

Some special characters are not included in the search index and therefore are not searchable. This also includes the special characters that represent search operators in the search query. Here's a list of special characters that are either replaced by a blank space in the actual search query or cause a search error.

```
+ - = : ! @ # % ^ & ; _ / ? ( ) [ ] { }
```

Searching for site content shared with external users

You can also use the Content Search feature in the Security & Compliance Center to search for documents stored on SharePoint and OneDrive for Business sites that have been shared with people outside of your organization. This can help you identify sensitive or proprietary information that's being shared outside your organization. You can do this by using the `ViewableByExternalUsers` property in a keyword query. This property returns documents or sites that have been shared with external users by using one of the following sharing methods:

- A sharing invitation that requires users to sign in to your organization as an authenticated user.
- An anonymous guest link, which allows anyone with this link to access the resource without having to be authenticated.

Here are some examples:

- The query `ViewableByExternalUsers:true AND SensitiveType:"Credit Card Number"` returns all items that have been shared with people outside your organization and contain a credit card number.
- The query
  ```
  ViewableByExternalUsers:true AND ContentType:document AND
  site:"https://contoso.sharepoint.com/Sites/Teams"
  ```
  returns a list of documents on all team sites in the organization that have been shared with external users.

> **TIP**
> A search query such as `ViewableByExternalUsers:true AND ContentType:document` might return a lot of .aspx files in the search results. To eliminate these (or other types of files), you can use the `FileExtension` property to exclude specific file types; for example `ViewableByExternalUsers:true AND ContentType:document NOT FileExtension:aspx`.

What is considered content that is shared with people outside your organization? Documents in your organization's SharePoint and OneDrive for Business sites that are shared by sending a sharing invitation or that are shared in public locations. For example, the following user activities result in content that is viewable by external users:

- A user shares a file or folder with a person outside your organization.
- A user creates and sends a link to a shared file to a person outside your organization. This link allows the external user to view (or edit) the file.
- A user sends a sharing invitation or a guest link to a person outside your organization to view (or edit) a shared file.

Issues using the ViewableByExternalUsers property

While the `ViewableByExternalUsers` property represents the status of whether a document or site is shared with external users, there are some caveats to what this property does and doesn't reflect. In the following scenarios, the value of the `ViewableByExternalUsers` property won't be updated, and the results of a Content Search query that uses this property may be inaccurate.

- Changes to sharing policy, such as turning off external sharing for a site or for the organization. The property will still show previously shared documents as being externally accessible even though external access might have been revoked.
- Changes to group membership, such as adding or removing external users to Microsoft 365 Groups or Microsoft 365 security groups. The property won't automatically be updated for items the group has access to.
- Sending sharing invitations to external users where the recipient hasn't accepted the invitation, and therefore doesn't yet have access to the content.

In these scenarios, the `ViewableByExternalUsers` property won't reflect the current sharing status until the site or document library is recrawled and reindexed.

Searching for site content shared within your organization

As previously explained, you can use the `SharedWithUsersOWSUser` property so search for documents that have been shared between people in your organization. When a person shares a file (or folder) with another user inside your organization, a link to the shared file appears on the **Shared with me** page in the OneDrive for Business account of the person who the file was shared with. For example, to search for the documents that have been shared with Sara Davis, you can use the query `SharedWithUsersOWSUser:"sarad@contoso.com"`. If you export the results of this search, the original documents (located in the content location of the person who shared the documents with Sara) will be downloaded.

Documents must be explicitly shared with a specific user to be returned in search results when using the

`SharedWithUsersOWSUser` property. For example, when a person shares a document in their OneDrive account, they have the option to share it with anyone (inside or outside the organization), share it only with people inside the organization, or share it with a specific person. Here's a screenshot of the **Share** window in OneDrive, that shows the three sharing options.

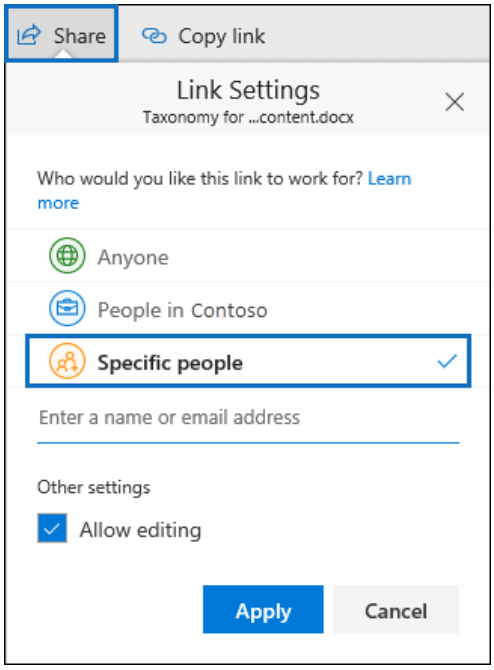

Only documents that are shared by using the third option (shared with **Specific people**) will be returned by a search query that uses the `SharedWithUsersOWSUser` property.

Searching for Skype for Business conversations

You can use the following keyword query to specifically search for content in Skype for Business conversations:

```
kind:im
```

The previous search query also returns chats from Microsoft Teams. To prevent this, you can narrow the search results to include only Skype for Business conversations by using the following keyword query:

```
kind:im AND subject:conversation
```

The previous keyword query excludes chats in Microsoft Teams because Skype for Business conversations are saved as email messages with a Subject line that starts with the word "Conversation".

To search for Skype for Business conversations that occurred within a specific date range, use the following keyword query:

```
kind:im AND subject:conversation AND (received=startdate..enddate)
```

Search tips and tricks

- Keyword searches are not case-sensitive. For example, **cat** and **CAT** return the same results.
- The Boolean operators **AND**, **OR**, **NOT**, and **NEAR** must be uppercase.
- A space between two keywords or two `property:value` expressions is the same as using **AND**. For example, `from:"Sara Davis" subject:reorganization` returns all messages sent by Sara Davis that contain the word reorganization in the subject line.
- Use syntax that matches the `property:value` format. Values are not case-sensitive, and they can't have a

space after the operator. If there is a space, your intended value will be a full-text search. For example `to: pilarp` searches for "pilarp" as a keyword, rather than for messages that were sent to pilarp.

- When searching a recipient property, such as To, From, Cc, or Recipients, you can use an SMTP address, alias, or display name to denote a recipient. For example, you can use pilarp@contoso.com, pilarp, or "Pilar Pinilla".
- You can use only prefix wildcard searches; for example, **cat*** or **set***. Suffix searches (***cat**), infix searches (**c*t**), and substring searches (***cat***) are not supported.
- When searching a property, use double quotation marks (" ") if the search value consists of multiple words. For example `subject:budget Q1` returns messages that contain **budget** in the subject line and that contain **Q1** anywhere in the message or in any of the message properties. Using `subject:"budget Q1"` returns all messages that contain **budget Q1** anywhere in the subject line.
- To exclude content marked with a certain property value from your search results, place a minus sign (-) before the name of the property. For example, `-from:"Sara Davis"` excludes any messages sent by Sara Davis.
- You can export items based on message type. For example, to export Skype conversations and chats in Microsoft Teams, use the syntax `kind:im`. To return only email messages, you would use `kind:email`. To return chats, meetings, and calls in Microsoft Teams, use `kind:microsoftteams`.

Configure permissions filtering for Content Search

2/18/2021 • 18 minutes to read • Edit Online

You can use search permissions filtering to let an eDiscovery manager search only a subset of mailboxes and sites in your organization. You can also use permissions filtering to let that same eDiscovery manager search only for mailbox or site content that meets a specific search criteria. For example, you might let an eDiscovery manager search only the mailboxes of users in a specific location or department. You do this by creating a filter that uses a supported recipient filter to limit which mailboxes a specific user or group of users can search. You can also create a filter that specifies what mailbox content a user can search for. This is done by creating a filter that uses a searchable message property. Similarly, you can let an eDiscovery manager search only specific SharePoint sites in your organization. You do this by creating a filter that limits which site can be searched. You can also create a filter that specifies what site content can be searched. This is done by creating a filter that uses a searchable site property.

You can also use search permissions filtering to create logical boundaries (called *compliance boundaries*) within an organization that control the user content locations (such as mailboxes, SharePoint sites, and OneDrive accounts) that specific eDiscovery managers can search. For more information, see Set up compliance boundaries for eDiscovery investigations in Office 365.

Search permissions filtering is supported by the Content Search feature in the Security & Compliance Center. These four cmdlets let you configure and manage search permissions filters:

New-ComplianceSecurityFilter

Get-ComplianceSecurityFilter

Set-ComplianceSecurityFilter

Remove-ComplianceSecurityFilter

Requirements to configure permissions filtering

- To run the compliance security filter cmdlets, you have to be a member of the Organization Management role group in the Security & Compliance Center. For more information, see Permissions in the Security & Compliance Center.
- You have to connect to both Exchange Online and Security & Compliance Center PowerShell to use the compliance security filter cmdlets. This is necessary because these cmdlets require access to mailbox properties, which is why you have to connect to Exchange Online PowerShell. See the steps in the next section.
- See the More information section for additional information about search permissions filters.
- Search permissions filtering is applicable to inactive mailboxes, which means you can use mailbox and mailbox content filtering to limit who can search an inactive mailbox. See the More information section for additional information about permissions filtering and inactive mailboxes.
- Search permissions filtering can't be used to limit who can search public folders in Exchange.
- There is no limit to the number of search permissions filters that can be created in an organization. But search performance will be impacted when there are more than 100 search permissions filters. To keep the number of search permissions filters in your organization as small as possible, create filters that combine rules for Exchange, SharePoint, and OneDrive in a single filter whenever possible.

Connect to Exchange Online and Security & Compliance Center PowerShell in a single session

Before you can successfully run the script in this section, you have to download and install the Exchange Online PowerShell V2 module. For information, see About the Exchange Online PowerShell V2 module.

1. Save the following text to a Windows PowerShell script file by using a filename suffix of **.ps1**. For example, you could save it to a file named **ConnectEXO-SCC.ps1**.

```
Import-Module ExchangeOnlineManagement
$UserCredential = Get-Credential
Connect-ExchangeOnline -Credential $UserCredential -ShowBanner:$false
Connect-IPPSSession -Credential $UserCredential
$Host.UI.RawUI.WindowTitle = $UserCredential.UserName + " (Exchange Online + Compliance Center)"
```

2. On your local computer, open Windows PowerShell, go to the folder where the script that you created in the previous step is located, and then run the script; for example:

```
.\ConnectEXO-SCC.ps1
```

How do you know if this worked? After you run the script, cmdlets from Exchange Online and Security & Compliance PowerShell are imported to your local Windows PowerShell session. If you don't receive any errors, you connected successfully. A quick test is to run an Exchange Online and Security & Compliance Center cmdlet. For example, you can run and **Get-Mailbox** and **Get-ComplianceSearch**.

For troubleshooting PowerShell connection errors, see:

- Connect to Exchange Online PowerShell
- Connect to Security & Compliance Center PowerShell

New-ComplianceSecurityFilter

The **New-ComplianceSecurityFilter** is used to create a search permissions filter. The following table describes the parameters for this cmdlet. All parameters are required to create a compliance security filter.

PARAMETER	DESCRIPTION
Action	The *Action* parameter specifies that type of search action that the filter is applied to. The possible Content Search actions are: **Export:** The filter is applied when exporting search results. **Preview:** The filter is applied when previewing search results. **Purge:** The filter is applied when purging search results. **Search:** The filter is applied when running a search. **All:** The filter is applied to all search actions.
FilterName	The *FilterName* parameter specifies the name of the permissions filter. This name is used to identity a filter when using the **Get-ComplianceSecurityFilter**, **Set-ComplianceSecurityFilter,** and **Remove-ComplianceSecurityFilter** cmdlets.
Filters	The *Filters* parameter specifies the search criteria for the compliance security filter. You can create three different types

PARAMETER	DESCRIPTION

	of filters:
	Mailbox filtering: This type of filter specifies the mailboxes the assigned users (specified by the *Users* parameter) can search. The syntax for this type of filter is **Mailbox_** *MailboxPropertyName*, where *MailboxPropertyName* specifies a mailbox property used to scope the mailboxes that can be searched. For example, the mailbox filter `"Mailbox_CustomAttribute10 -eq 'OttawaUsers'"` would allow the user assigned this filter to search only the mailboxes that have the value "OttawaUsers" in the CustomAttribute10 property.
	Any supported filterable recipient property can be used for the *MailboxPropertyName* property. For a list of supported properties, see Filterable properties for the -RecipientFilter parameter.
	
	Mailbox content filtering: This type of filter is applied on the content that can be searched. It specifies the mailbox content the assigned users can search for. The syntax for this type of filter is **MailboxContent_** *SearchablePropertyName: value*, where *SearchablePropertyName* specifies a Keyword Query Language (KQL) property that can be specified in a Content Search. For example, the mailbox content filter `MailboxContent_recipients:contoso.com` would allow the user assigned this filter to only search for messages sent to recipients in the contoso.com domain.
	For a list of searchable message properties, see Keyword queries and search conditions for Content Search.
	
	Important: A single search filter can't contain a mailbox filter and a mailbox content filter. To combine these in a single filter, you have to use a filters list. But a filter can contain a more complex query of the same type. For example,
	`"Mailbox_CustomAttribute10 -eq 'FTE' -and Mailbox_MemberOfGroup -eq '$($DG.DistinguishedName)'"`
	
	Site and site content filtering: There are two SharePoint and OneDrive for Business site-related filters that you can use to specify what site or site content the assigned users can search:
	
	- **Site_** *SearchableSiteProperty*
	- **SiteContent_** *SearchableSiteProperty*
	
	These two filters are interchangeable. For example,
	`"Site_Path -like 'https://contoso.sharepoint.com/sites/doctors*'"`
	and
	`"SiteContent_Path -like 'https://contoso.sharepoint.com/sites/doctors*'"`
	return the same results. But to help you identify what a filter does, you can use `Site_` to specify site-related properties (such as a site URL) and `SiteContent_` to specify content-related properties (such as document types. For example, the filter
	`"Site_Path -like 'https://contoso.sharepoint.com/sites/doctors*'"`
	would allow the user assigned this filter to only search for content in the https://contoso.sharepoint.com/sites/doctors site collection. The filter
	`"SiteContent_FileExtension -eq 'docx'"` would allow the user assigned this filter to only search for Word documents (Word 2007 and later).

PARAMETER	DESCRIPTION
	For a list of searchable site properties, see Overview of crawled and managed properties in SharePoint. Properties marked with a **Yes** in the **Queryable** column can be used to create a site or site content filter. **Important:** You have to create a search permissions filter to explicitly prevent users from searching content locations in a specific service (such as preventing a user from searching any Exchange mailbox or any SharePoint site). In other words, creating a search permissions filter that allows a user to search all SharePoint sites in the organization doesn't prevent that user from searching mailboxes. For example, to allow SharePoint admins to only search SharePoint sites, you have to create a filter that prevents them from searching mailboxes. Similarly, to allow Exchange admins to only search mailboxes, you have to create a filter that prevents them from searching sites.
Users	The *Users* parameter specifies the users who get this filter applied to their Content Searches. Identify users by their alias or primary SMTP address. You can specify multiple values separated by commas, or you can assign the filter to all users by using the value **All**. You can also use the *Users* parameter to specify a Security & Compliance Center role group. This lets you create a custom role group and then assign that role group a search permissions filter. For example, let's say you have a custom role group for eDiscovery managers for the U.S. subsidiary of a multi-national corporation. You can use the *Users* parameter to specify this role group (by using the Name property of the role group) and then use the *Filter* parameter to allow only mailboxes in the U.S. to be searched. You can't specify distribution groups with this parameter.

Using a filters list to combine filter types

A *filters list* is a filter that includes a mailbox filter and a site filter separated by a comma. Using a filters list is the only supported method for combining different types of filters. In the following example, notice that a comma separates the **Mailbox** and **Site** filters:

```
-Filters "Mailbox_CustomAttribute10 -eq 'OttawaUsers'", "Site_Path -like
'https://contoso.sharepoint.com/sites/doctors*'"
```

When a filter that contains a filters list is processed during the running of a content search, two search permissions filters are created from the filters list: One for each filter that's separated by a comma. So in the previous example, one mailbox search permissions filter and one site search permissions filter would be created.

An alternative to using a filters list would be to create two separate search permissions filters. So in the previous example, you'd create one filter for the mailbox attribute and one filter for the site attribute. In either case, the results are the same. Using a filters list or creating separate search permissions filters is a matter of preference.

Keep the following things in mind about using a filters list:

- You have to use a filters list to create a filter that includes a **Mailbox** filter and a **MailboxContent** filter.
- As previously suggested, you don't have to use a filters list to include a **Site** and a **SiteContent** filter in a single search permissions filter. For example, you can combine **Site** and a **SiteContent** filters using an **-or** operator.

```
-Filters "Site_ComplianceAttribute -eq 'FourthCoffee' -or Site_Path -like
'https://contoso.sharepoint.com/sites/FourthCoffee*'"
```

- Each component of a filters list can contain a complex filter syntax. For example, the mailbox and site filters can contain multiple filters separated by an **-or** operator:

```
-Filters "Mailbox_Department -eq 'CohoWinery' -or Mailbox_CustomAttribute10 -eq 'CohoUsers'",
"Site_ComplianceAttribute -eq 'CohoWinery' -or Site_Path -like
'https://contoso.sharepoint.com/sites/CohoWinery*'"
```

Examples of creating search permissions filters

Here are examples of using the **New-ComplianceSecurityFilter** cmdlet to create a search permissions filter.

This example allows the user annb@contoso.com to perform all Content Search actions only for mailboxes in Canada. This filter contains the three-digit numeric country code for Canada from ISO 3166-1.

```
New-ComplianceSecurityFilter -FilterName CountryFilter  -Users annb@contoso.com -Filters
"Mailbox_CountryCode  -eq '124'" -Action All
```

This example allows the users donh and suzanf to search only the mailboxes that have the value 'Marketing' for the CustomAttribute1 mailbox property.

```
New-ComplianceSecurityFilter -FilterName MarketingFilter  -Users donh,suzanf -Filters
"Mailbox_CustomAttribute1  -eq 'Marketing'" -Action Search
```

This example allows members of the "US Discovery Managers" role group to perform all Content Search actions only on mailboxes in the United States. This filter contains the three-digit numeric country code for the United States from ISO 3166-1.

```
New-ComplianceSecurityFilter -FilterName USDiscoveryManagers  -Users "US Discovery Managers" -Filters
"Mailbox_CountryCode  -eq '840'" -Action All
```

This example allows members of the eDiscovery Manager role group to search only the mailboxes of members of the Ottawa Users distribution group. The Get-DistributionGroup cmdlet in Exchange Online PowerShell is used to find the members of the Ottawa Users group.

```
$DG = Get-DistributionGroup "Ottawa Users"
```

```
New-ComplianceSecurityFilter -FilterName DGFilter  -Users eDiscoveryManager -Filters "Mailbox_MemberOfGroup
-eq '$($DG.DistinguishedName)'" -Action Search
```

This example prevents any user from deleting content from the mailboxes of members of the Executive Team distribution group. The Get-DistributionGroup cmdlet in Exchange Online PowerShell is used to find the members of the Executive Team group.

```
$DG = Get-DistributionGroup "Executive Team"
```

```
New-ComplianceSecurityFilter -FilterName NoExecutivesPreview  -Users All -Filters "Mailbox_MemberOfGroup -ne
'$($DG.DistinguishedName)'" -Action Purge
```

This example allows members of the OneDrive eDiscovery Managers custom role group to only search for content in OneDrive for Business locations in the organization.

```
New-ComplianceSecurityFilter -FilterName OneDriveOnly  -Users "OneDrive eDiscovery Managers" -Filters
"Site_Path -like 'https://contoso-my.sharepoint.com/personal*'" -Action Search
```

NOTE

To restrict users to searching specific sites, use the filter `Site_Path`, as shown in the previous example. Using `Site_Site` will not work.

This example restricts the user to performing all Content Search actions only on email messages sent during the calendar year 2015.

```
New-ComplianceSecurityFilter -FilterName EmailDateRestrictionFilter -Users donh@contoso.com -Filters
"MailboxContent_Received -ge '01-01-2015' -and MailboxContent_Received -le '12-31-2015'" -Action All
```

Similar to the previous example, this example restricts the user to performing all Content Search actions on documents that were last changed sometime in the calendar year 2015.

```
New-ComplianceSecurityFilter -FilterName DocumentDateRestrictionFilter -Users donh@contoso.com -Filters
"SiteContent_LastModifiedTime -ge '01-01-2015' -and SiteContent_LastModifiedTime -le '12-31-2015'" -Action
All
```

This example prevents members of the "OneDrive Discovery Managers" role group from performing content search actions on any mailbox in the organization.

```
New-ComplianceSecurityFilter -FilterName NoEXO -Users "OneDrive Discovery Managers" -Filters "Mailbox_Alias
-notlike '*'"  -Action All
```

This example prevents anyone in the organization from searching for email messages that were sent or received by janets or sarad.

```
New-ComplianceSecurityFilter -FilterName NoSaraJanet -Users All -Filters "MailboxContent_Participants -
notlike 'janets@contoso.onmicrosoft.com' -and MailboxContent_Participants -notlike
'sarad@contoso.onmicrosoft.com'" -Action Search
```

This example uses a filters list to combine mailbox and site filters.

```
New-ComplianceSecurityFilter -FilterName "Coho Winery Security Filter" -Users "Coho Winery eDiscovery
Managers", "Coho Winery Investigators" -Filters "Mailbox_Department -eq 'CohoWinery'",
"Site_ComplianceAttribute -eq 'CohoWinery' -or Site_Path -like
'https://contoso.sharepoint.com/sites/CohoWinery*'" -Action ALL
```

Get-ComplianceSecurityFilter

The **Get-ComplianceSecurityFilter** is used to return a list of search permissions filters. Use the *FilterName*

parameter to return information for a specific search filter.

Set-ComplianceSecurityFilter

The **Set-ComplianceSecurityFilter** is used to modify an existing search permissions filter. The only required parameter is *FilterName*.

PARAMETER	DESCRIPTION
Action	The *Action* parameter specifies that type of search action that the filter is applied to. The possible Content Search actions are: **Export:** The filter is applied when exporting search results. **Preview:** The filter is applied when previewing search results. **Purge:** The filter is applied when purging search results. **Search:** The filter is applied when running a search. **All:** The filter is applied to all search actions.
FilterName	The *FilterName* parameter specifies the name of the permissions filter.
Filters	The *Filters* parameter specifies the search criteria for the compliance security filter. You can create two different types of filters: **Mailbox filtering:** This type of filter specifies the mailboxes the assigned users (specified by the *Users* parameter) can search. The syntax for this type of filter is **Mailbox_** *MailboxPropertyName*, where *MailboxPropertyName* specifies a mailbox property used to scope the mailboxes that can be searched. For example, the mailbox filter `"Mailbox_CustomAttribute10 -eq 'OttawaUsers'"` would allow the user assigned this filter to search only the mailboxes that have the value "OttawaUsers" in the CustomAttribute10 property. Any supported filterable recipient property can be used for the *MailboxPropertyName* property. For a list of supported properties, see Filterable properties for the -RecipientFilter parameter. **Mailbox content filtering:** This type of filter is applied on the content that can be searched. It specifies the mailbox content the assigned users can search for. The syntax for this type of filter is **MailboxContent_** *SearchablePropertyName:value*, where *SearchablePropertyName* specifies a Keyword Query Language (KQL) property that can be specified in a Content Search. For example, the mailbox content filter `MailboxContent_recipients:contoso.com` would allow the user assigned this filter to only search for messages sent to recipients in the contoso.com domain. For a list of searchable message properties, see Keyword queries for Content Search. **Site and site content filtering:** There are two SharePoint and OneDrive for Business site-related filters that you can use to specify what site or site content the assigned users can search: - **Site_** *SearchableSiteProperty* - **SiteContent_** *SearchableSiteProperty*

PARAMETER	DESCRIPTION
	- SiteContent_*SearchableSiteProperty* These two filters are interchangeable. For example, `"Site_Path -like 'https://contoso.spoppe.com/sites/doctors*'"` and `"SiteContent_Path -like 'https://contoso.spoppe.com/sites/doctors*'"` returns the same results. But to help you identify what a filter does, you can use `Site_` to specify site-related properties (such as a site URL) and `SiteContent_` to specify content-related properties (such as document types. For example, the filter `"Site_Path -like 'https://contoso.spoppe.com/sites/doctors*'"` would allow the user assigned this filter to only search for content in the https://contoso.spoppe.com/sites/doctors site collection. The filter `"SiteContent_FileExtension -eq 'docx'"` would allow the user assigned this filter to only search for Word documents (Word 2007 and later). For a list of searchable site properties, see Overview of crawled and managed properties in SharePoint. Properties marked with a **Yes** in the **Queryable** column can be used to create a site or site content filter.
Users	The *Users* parameter specifies the users who get this filter applied to their Content Searches. Because this is a multi-value property, specifying a user or group of users with this parameter overwrite the existing list of users. See the following examples for the syntax to add and remove selected users. You can also use the *Users* parameter to specify a Security & Compliance Center role group. This lets you create a custom role group and then assign that role group a search permissions filter. For example, let's say you have a custom role group for eDiscovery managers for the U.S. subsidiary of a multi-national corporation. You can use the *Users* parameter to specify this role group (by using the Name property of the role group) and then use the *Filter* parameter to allow only mailboxes in the U.S. to be searched. You can't specify distribution groups with this parameter.

Examples of changing search permissions filters

These examples show how to use the **Get-ComplianceSecurityFilter** and **Set-ComplianceSecurityFilter** cmdlets to add or remove a user to the existing list of users that the filter is assigned to. When you add or remove users from a filter, specify the user by using their SMTP address.

This example adds a user to the filter.

```
$filterusers = Get-ComplianceSecurityFilter -FilterName OttawaUsersFilter
```

```
$filterusers.users.add("pilarp@contoso.com")
```

```
Set-ComplianceSecurityFilter -FilterName OttawaUsersFilter -Users $filterusers.users
```

This example removes a user from the filter.

```
$filterusers = Get-ComplianceSecurityFilter -FilterName OttawaUsersFilter
```

```
$filterusers.users.remove("annb@contoso.com")
```

```
Set-ComplianceSecurityFilter -FilterName OttawaUsersFilter -Users $filterusers.users
```

Remove-ComplianceSecurityFilter

The **Remove-ComplianceSecurityFilter** is used to delete a search filter. Use the *FilterName* parameter to specify the filter you want to delete.

More information

- **How does search permissions filtering work?** The permissions filter is added to the search query when a Content Search is run. The permissions filter is joined to the search query by the **AND** Boolean operator. For example, you have a permissions filter that allows Bob to perform all search actions on the mailboxes of members of the Workers distribution group. Then Bob runs a Content Search on all mailboxes in the organization with the search query `sender:jerry@adatum.com`. Because the permissions filter and the search query are logically combined by an **AND** operator, the search returns any message sent by jerry@adatum.com to any member of the Workers distribution group.

- **What happens if you have multiple search permissions filters?** 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. Let's take the previous example, where a search filter allows Bob to search only the mailboxes of the members of the Workers distribution group. Then we create another filter that prevents Bob from searching Phil's mailbox ("Mailbox_Alias -ne 'Phil'"). And let's also assume that Phil is a member of the Workers group. When Bob runs a Content Search (from the previous example) on all mailboxes in the organization, search results are returned for Phil's mailbox even though you applied filter to prevent Bob from searching Phil's mailbox. This is because the first filter, which allows Bob to search the Workers group, is true. And because Phil is a member of the Workers group, Bob can search Phil's mailbox.

- **Does search permissions filtering work for inactive mailboxes?** Yes, you can use mailbox and mailbox content filters to limit who can search inactive mailboxes in your organization. Like a regular mailbox, an inactive mailbox has to be configured with the recipient property that's used to create a permissions filter. If necessary, you can use the **Get-Mailbox -InactiveMailboxOnly** command to display the properties of inactive mailboxes. For more information, see Create and manage inactive mailboxes in Office 365.

- **Does search permissions filtering work for public folders?** No. As previously explained, search permissions filtering can't be used to limit who can search public folders in Exchange. For example, items in public folder locations can't be excluded from the search results by a permissions filter.

- **Does allowing a user to search all content locations in a specific service also prevent them from searching content locations in a different service?** No. As previously explained, you have to create a search permissions filter to explicitly prevent users from searching content locations in a specific

service (such as preventing a user from searching any Exchange mailbox or any SharePoint site). In other words, creating a search permissions filter that allows a user to search all SharePoint sites in the organization doesn't prevent that user from searching mailboxes. For example, to allow SharePoint admins to only search SharePoint sites, you have to create a filter that prevents them from searching mailboxes. Similarly, to allow Exchange admins to only search mailboxes, you have to create a filter that prevents them from searching sites.

Set up compliance boundaries for eDiscovery investigations

2/18/2021 • 20 minutes to read • Edit Online

The guidance in this article can be applied when using either Core eDiscovery or Advanced eDiscovery to manage investigations.

Compliance boundaries create logical boundaries within an organization that control the user content locations (such as mailboxes, OneDrive accounts, and SharePoint sites) that eDiscovery managers can search. Also, compliance boundaries control who can access eDiscovery cases used to manage the legal, human resources, or other investigations within your organization. The need for compliance boundaries is often necessary for multi-national corporations that have to respect geographical boarders and regulations and for governments, which are often divided into different agencies. In Microsoft 365, compliance boundaries help you meet these requirements when performing content searches and managing investigations with eDiscovery cases.

We use the example in the following illustration to explain how compliance boundaries work.

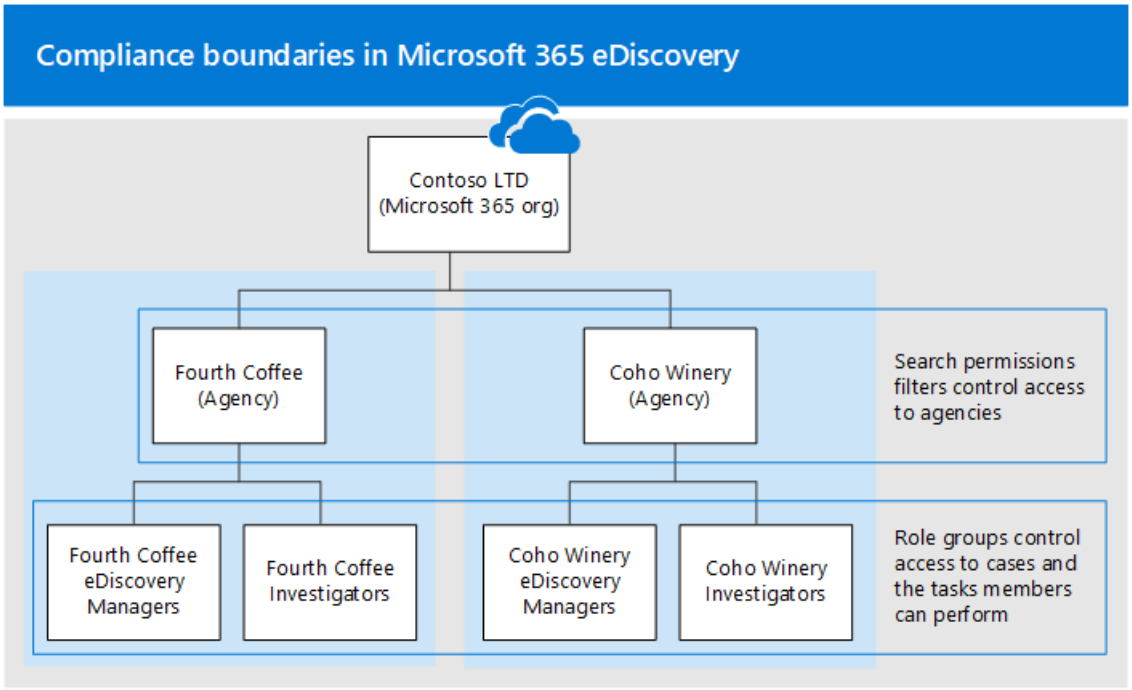

In this example, Contoso LTD is an organization that consists of two subsidiaries, Fourth Coffee and Coho Winery. The business requires that eDiscovery mangers and investigators can only search the Exchange mailboxes, OneDrive accounts, and SharePoint sites in their agency. Also, eDiscovery managers and investigators can only see eDiscovery cases in their agency, and they can only access the cases that they're a member of. Additionally in this scenario, investigators cannot place content locations on hold or export content from a case. Here's how compliance boundaries meet these requirements.

- The search permissions filtering functionality in Content search controls the content locations that eDiscovery managers and investigators can search. This means eDiscovery managers and investigators in the Fourth Coffee agency can only search content locations in the Fourth Coffee subsidiary. The same restriction applies to the Coho Winery subsidiary.
- Role groups provide the following functions for compliance boundaries:
    - Control who can see the eDiscovery cases in the Security & Compliance Center. This means that

eDiscovery managers and investigators can only see the eDiscovery cases in their agency.

- Control who can assign members to an eDiscovery case. This means eDiscovery managers and investigators can only assign members to cases that they themselves are a member of.
- Control the eDiscovery-related tasks that members can perform by adding or removing roles that assign specific permissions.

Here's the process for setting up compliance boundaries:

Step 1: Identify a user attribute to define your agencies

Step 2: File a request with Microsoft Support to synchronize the user attribute to OneDrive accounts

Step 3: Create a role group for each agency

Step 4: Create a search permissions filter to enforce the compliance boundary

Step 5: Create an eDiscovery case for an intra-agency investigations

Before you set up compliance boundaries

You have to meet the following prerequisites before the Azure Active Directory (Azure AD) attribute that you identity (in Step 1) can be successfully synched to a user's OneDrive account (in Step 2):

- Users must be assigned an Exchange Online license and a SharePoint Online license.
- User mailboxes must be at least 10 MB in size. If a user's mailbox is less than 10 MB, the attribute used to define your agencies won't be synched to the user's OneDrive account.
- Compliance boundaries and the attributes used to create search permissions filters require that Azure Active Directory (Azure AD) attributes are synchronized to user mailboxes. To verify that the attributes that you want to use have been synchronized, run the Get-User cmdlet in Exchange Online PowerShell. The output of this cmdlet displays the Azure AD attributes synchronized to Exchange Online.

Step 1: Identify a user attribute to define your agencies

The first step is to choose an Azure AD attribute to use that will define your agencies. This attribute is used to create the search permissions filter that limits an eDiscovery manager to search only the content locations of users who are assigned a specific value for this attribute. For example, let's say Contoso decides to use the **Department** attribute. The value for this attribute for users in the Fourth Coffee subsidiary would be `FourthCoffee` and the value for users in Coho Winery subsidiary would be `CohoWinery`. In Step 4, you use this `attribute:value` pair (for example, *Department:FourthCoffee*) to limit the user content locations that eDiscovery managers can search.

Here's a list of Azure AD user attributes that you can use for compliance boundaries:

- Company
- CustomAttribute1 - CustomAttribute15
- Department
- Office
- C (Two-letter country code) *

> **NOTE**
>
> * This attribute maps to the CountryOrRegion property that is returned by running the **Get-User** cmdlet in Exchange Online PowerShell. The cmdlet returns the localized country name, which is translated from the two-letter country code. For more information, see the CountryOrRegion parameter description in the Set-User cmdlet reference article.

Although more user attributes are available, particularly for Exchange mailboxes, the attributes listed above are the only ones currently supported by OneDrive.

Step 2: File a request with Microsoft Support to synchronize the user attribute to OneDrive accounts

The next step is to file a request with Microsoft Support to synchronize the Azure AD attribute that you chose in Step 1 to all OneDrive accounts in your organization. After this synchronization occurs, the attribute (and its value) that you chose in Step 1 will be mapped to a hidden managed property named `ComplianceAttribute`. You use this attribute to create the search permissions filter for OneDrive in Step 4.

Include the following information when you submit the request to Microsoft support:

- The default domain name of your organization
- The name of the Azure AD attribute (from Step 1)
- The following title or description of the purpose of the support request: "Enable OneDrive for Business Synchronization with Azure AD for Compliance Security Filters". This helps route the request to the eDiscovery engineering team who implements the request.

After the engineering change is made and the attribute is synchronized to OneDrive, Microsoft Support will send you the build number that the change was made in and an estimated deployment date. The deployment process usually takes 4–6 weeks after you submit the support request.

> **IMPORTANT**
>
> You can complete Step 3 through Step 5 before this attribute change is deployed. But running content searches won't return documents from OneDrive accounts that are specified in a search permissions filter until after the attribute synch is deployed.

Step 3: Create a role group for each agency

The next step is to create the role groups in the Security & Compliance Center that will align with your agencies. We recommend that you create a role group by copying the built-in eDiscovery Managers group, adding the appropriate members, and removing roles that may not be applicable to your needs. For more information about eDiscovery-related roles, see Assign eDiscovery permissions in the Office 365 Security & Compliance Center.

To create the role groups, go to the **Permissions** page in the Security & Compliance Center and create a role group for each team in each agency that will use compliance boundaries and eDiscovery cases to manage investigations.

Using the Contoso compliance boundaries scenario, four role groups need to be created and the appropriate members added to each one.

- Fourth Coffee eDiscovery Managers

- Fourth Coffee Investigators
- Coho Winery eDiscovery Managers
- Coho Winery Investigators

To meet the requirements of the Contoso compliance boundaries scenario, you would also remove the **Hold** and **Export** roles from the investigators role groups to prevent investigators from placing holds on content locations and exporting content from a case.

Step 4: Create a search permissions filter to enforce the compliance boundary

After you've created role groups for each agency, the next step is to create the search permissions filters that associate each role group to its specific agency and defines the compliance boundary itself. You need to create one search permissions filter for each agency. For more information about creating security permissions filters, see Configure permissions filtering for Content Search.

Here's the syntax that's used to create a search permissions filter used for compliance boundaries.

```
New-ComplianceSecurityFilter -FilterName <name of filter> -Users <role groups> -Filters
"Mailbox_<ComplianceAttribute>  -eq '<AttributeVale> '", "Site_<ComplianceAttribute>  -eq '<AttributeValue>'
-or Site_Path -like '<SharePointURL>*'" -Action <Action >
```

Here's a description of each parameter in the command:

- `FilterName`: Specifies the name of the filter. Use a name that describes or identifies the agency that the filter is used in.
- `Users`: Specifies the users or groups who get this filter applied to the Content Search actions they perform. For compliance boundaries, this parameter specifies the role groups (that you created in Step 3) in the agency that you're creating the filter for. Note this is a multi-value parameter so you can include one or more role groups, separated by commas.
- `Filters`: Specifies the search criteria for the filter. For the compliance boundaries, you define the following filters. Each one applies to a content location.
  - `Mailbox`: Specifies the mailboxes that the role groups defined in the `Users` parameter can search. For compliance boundaries, *ComplianceAttribute* is the same attribute that you identified in Step 1 and *AttributeValue* specifies the agency. This filter allows members of the role group to search only the mailboxes in a specific agency; for example, `"Mailbox_Department -eq 'FourthCoffee'"`.
  - `Site`: Specifies the OneDrive accounts that the role groups defined in the `Users` parameter can search. For the OneDrive filter, use the actual string `ComplianceAttribute`. This maps to the same attribute that you identified in Step 1 and that's synchronized to OneDrive accounts as a result of the support request that you submitted in Step 2; *AttributeValue* specifies the agency. This filter allows members of the role group to search only the OneDrive accounts in a specific agency; for example, `"Site_ComplianceAttribute -eq 'FourthCoffee'"`.
  - `Site_Path`: Specifies the SharePoint sites that the role groups defined in the `Users` parameter can search. The *SharePointURL* specifies the sites in the agency that members of the role group can search. For example, `"Site_Path -like 'https://contoso.sharepoint.com/sites/FourthCoffee*'"`. Notice the `Site` and `Site_Path` filters are connected by an **-or** operator.

> **NOTE**
>
> The syntax for the `Filters` parameter includes a *filters list*. A filters list is a filter that includes a mailbox filter and a site filter separated by a comma. In the previous example, notice that a comma separates **Mailbox_ComplianceAttribute** and **Site_ComplianceAttribute**:
> `-Filters "Mailbox_<ComplianceAttribute> -eq '<AttributeVale> '", "Site_ComplianceAttribute -eq '<AttributeValue>' -or Site_Path -like '<SharePointURL>*'"`
> . When this filter is processed during the running of a content search, two search permissions filters are created from the filters list: one mailbox filter and one site filter. An alternative to using a filters list would be to create two separate search permissions filters for each agency: one search permissions filter for the mailbox attribute and one filter for the site attributes. In either case, the results will be the same. Using a filters list or creating separate search permissions filters is a matter of preference.

- `Action`: Specifies the type of Compliance Search action that the filter is applied to. For example, `-Action Search` would only apply the filter when members of the role group defined in the `Users` parameter run a content search. In this case, the filter wouldn't be applied when exporting search results. For compliance boundaries, use `-Action All` so the filter applies to all search actions.

  For a list of the Content Search actions, see the "New-ComplianceSecurityFilter" section in Configure permissions filtering for Content Search.

Here are examples of the two search permissions filters that would be created to support the Contoso compliance boundaries scenario. Both of these examples include a comma-separated filters list, in which the mailbox and site filters are included in the same search permissions filter and are separated by a comma.

Fourth Coffee

```
New-ComplianceSecurityFilter -FilterName "Fourth Coffee Security Filter" -Users "Fourth Coffee eDiscovery
Managers", "Fourth Coffee Investigators" -Filters "Mailbox_Department -eq 'FourthCoffee'",
"Site_ComplianceAttribute -eq 'FourthCoffee' -or Site_Path -like
'https://contoso.sharepoint.com/sites/FourthCoffee*'" -Action ALL
```

Coho Winery

```
New-ComplianceSecurityFilter -FilterName "Coho Winery Security Filter" -Users "Coho Winery eDiscovery
Managers", "Coho Winery Investigators" -Filters "Mailbox_Department -eq 'CohoWinery'",
"Site_ComplianceAttribute -eq 'CohoWinery' -or Site_Path -like
'https://contoso.sharepoint.com/sites/CohoWinery*'" -Action ALL
```

Step 5: Create an eDiscovery case for intra-agency investigations

The final step is to create a Core eDiscovery case or Advanced eDiscovery case in the Microsoft 365 compliance center and then add the role group that you created in Step 3 as a member of the case. This results in two important characteristics of using compliance boundaries:

- Only members of the role group added to the case will be able to see and access the case in the Security & Compliance Center. For example, if the Fourth Coffee Investigators role group is the only member of a case, then members of the Fourth Coffee eDiscovery Managers role group (or members of any other role group) won't be able to see or access the case.
- When a member of the role group assigned to a case runs a search associated with the case, they will only be able to search the content locations within their agency (which is defined by the search permissions filter that you created in Step 4.)

To create a case and assign members:

1. Go to the **Core eDiscovery** or **Advanced eDiscovery** page in the Microsoft 365 compliance center and create a case.
2. In the list of cases, click the name of the case you created.
3. In the **Manage this case** flyout page, under **Manage role groups**, click **+ Add**.

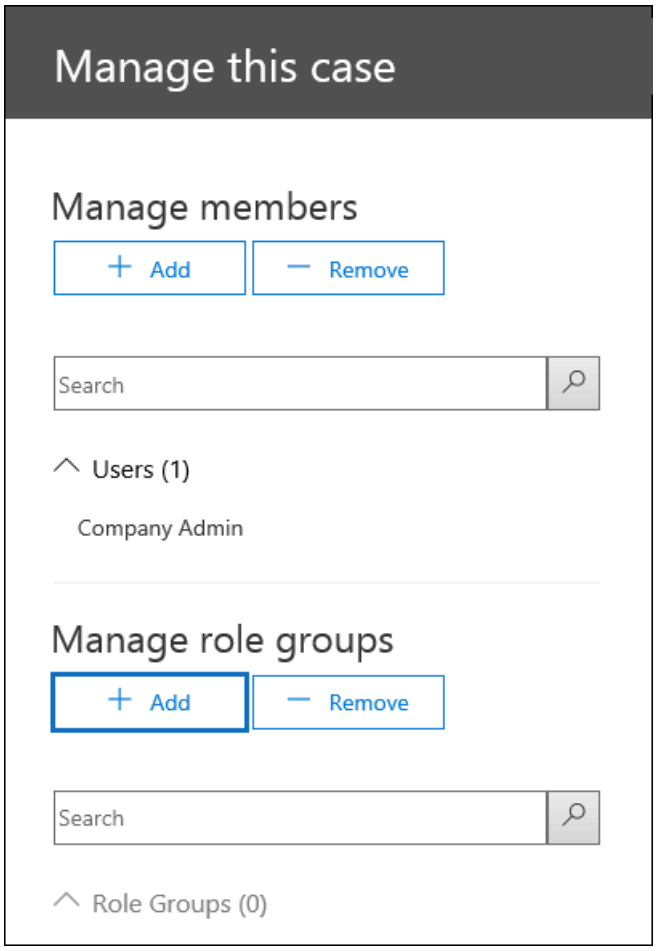

4. In the list of role groups, select one of the role groups that you created in Step 3, and click **Add**.
5. Click **Save** on the **Manage this case** flyout to save the change.

> **NOTE**

When adding a role group to a case, you can only add the role groups that you are a member of.

Searching and exporting content in Multi-Geo environments

Search permissions filters also let you control where content is routed for export and which datacenter can be searched when searching content locations in a SharePoint Multi-Geo environment.

- **Export search results:** You can export the search results from Exchange mailboxes, SharePoint sites, and OneDrive accounts from a specific datacenter. This means that you can specify the datacenter location that search results will be exported from.

  Use the **Region** parameter for **New-ComplianceSecurityFilter** or **Set-ComplianceSecurityFilter** cmdlets to create or change which datacenter the export will be routed through.

PARAMETER VALUE	DATACENTER LOCATION
NAM	North American (datacenters are in the US)
EUR	Europe
APC	Asia Pacific
CAN	Canada

- **Route content searches:** You can route the content searches of SharePoint sites and OneDrive accounts to a satellite datacenter. This means you can specify the datacenter location where searches will be run.

  Use one of the following values for the **Region** parameter to control the datacenter location that searches will run in when searching SharePoint sites and OneDrive accounts.

PARAMETER VALUE	DATACENTER ROUTING LOCATIONS FOR SHAREPOINT
NAM	US
EUR	Europe
APC	Asia Pacific
CAN	US
AUS	Asia Pacific
KOR	The organization's default datacenter
GBR	Europe
JPN	Asia Pacific
IND	Asia Pacific
LAM	US
NOR	Europe
BRA	North American datacenters

  If you don't specify the **Region** parameter for a search permissions filter, the organization's primary SharePoint region will be searched. Search results are exported to the closest datacenter.

  To simplify the concept, the **Region** parameter controls the datacenter that is used to search for content in SharePoint and OneDrive. This doesn't apply to searching for content in Exchange because Exchange content searches aren't bound by the geographic location of datacenters. Also, the same **Region** parameter value may also dictate the datacenter that exports are routed through. This is often necessary to control the movement of data across geographic boarders.

NOTE

If you're using Advanced eDiscovery, the **Region** parameter doesn't control the region that data is exported from. Data is exported from the organization's primary datacenter. Also, searching for content in SharePoint and OneDrive isn't bound by the geographic location of datacenters. All datacenters are searched. For more information about Advanced eDiscovery, see Overview of the Advanced eDiscovery solution in Microsoft 365.

Here are examples of using the **Region** parameter when creating search permission filters for compliance boundaries. This assumes that the Fourth Coffee subsidiary is located in North America and that Coho Winery is in Europe.

```
New-ComplianceSecurityFilter -FilterName "Fourth Coffee Security Filter" -Users "Fourth Coffee eDiscovery Managers", "Fourth Coffee Investigators" -Filters "Mailbox_Department -eq 'FourthCoffee'", "Site_Department -eq 'FourthCoffee' -or Site_Path -like 'https://contoso.sharepoint.com/sites/FourthCoffee*'" -Action ALL -Region NAM
```

```
New-ComplianceSecurityFilter -FilterName "Coho Winery Security Filter" -Users "Coho Winery eDiscovery Managers", "Coho Winery Investigators" -Filters "Mailbox_Department -eq 'CohoWinery'", "Site_Department -eq 'CohoWinery' -or Site_Path -like 'https://contoso.sharepoint.com/sites/CohoWinery*'" -Action ALL -Region EUR
```

Keep the following things in mind when searching and exporting content in multi-geo environments.

- The **Region** parameter doesn't control searches of Exchange mailboxes. All datacenters will be searched when you search mailboxes. To limit the scope of which Exchange mailboxes are searched, use the **Filters** parameter when creating or changing a search permissions filter.
- If it's necessary for an eDiscovery Manager to search across multiple SharePoint regions, you need to create a different user account for that eDiscovery manager to use in the search permissions filter to specify the region where the SharePoint sites or OneDrive accounts are located. For more information about setting this up, see the "Searching for content in a SharePoint Multi-Geo environment" section in Content Search.
- When searching for content in SharePoint and OneDrive, the **Region** parameter directs searches to either the primary or satellite location where the eDiscovery manager will conduct eDiscovery investigations. If an eDiscovery manager searches SharePoint and OneDrive sites outside of the region that's specified in the search permissions filter, no search results are returned.
- When exporting search results, content from all content locations (including Exchange, Skype for Business, SharePoint, OneDrive, and other services that you can search by using the Content Search tool) are uploaded to the Azure Storage location in the datacenter that's specified by the **Region** parameter. This helps organizations stay within compliance by not allowing content to be exported across controlled borders. If no region is specified in the search permissions filter, content is uploaded to the organization's primary datacenter.
- You can edit an existing search permissions filter to add or change the region by running the following command:

  ```
  Set-ComplianceSecurityFilter -FilterName <Filter name>  -Region <Region>
  ```

Using compliance boundaries for SharePoint hub sites

SharePoint hub sites often align with the same geographical or agency boundaries that eDiscovery compliance boundaries follow. That means you can use the site ID property of the hub site to create a compliance boundary.

To do this, use the Get-SPOHubSite cmdlet in SharePoint Online PowerShell to obtain the SiteId for the hub site and then use this value for the department ID property to create a search permissions filter.

Use the following syntax to create a search permissions filter for a SharePoint hub site:

```
New-ComplianceSecurityFilter -FilterName <Filter Name> -Users <User or Group> -Filters "Site_Departmentid -eq '{SiteId of hub site}'" -Action ALL
```

Here's an example of creating a search permissions filter for a hub site for the Coho Winery agency:

```
New-ComplianceSecurityFilter -FilterName "Coho Winery Hub Site Security Filter" -Users "Coho Winery eDiscovery Managers", "Coho Winery Investigators" -Filters "Site_Departmentid -eq '44252d09-62c4-4913-9eb0-a2a8b8d7f863'" -Action ALL
```

Compliance boundary limitations

Keep the following limitations in mind when managing eDiscovery cases and investigations that use of compliance boundaries.

- When creating and running a search, you can select content locations that are outside of your agency. However, because of the search permissions filter, content from those locations isn't included in the search results.
- Compliance boundaries don't apply to holds in eDiscovery cases. That means an eDiscovery manager in one agency can place a user in a different agency on hold. However, the compliance boundary will be enforced if the eDiscovery manager searches the content locations of the user who was placed on hold. That means the eDiscovery manager won't be able search the user's content locations, even though they were able to place the user on hold.

  Also, hold statistics will only apply to content locations in the agency.
- Search permissions filters aren't applied to Exchange public folders.

More information

- If a mailbox is de-licensed or soft-deleted, Azure AD attributes are no longer synchronized to the mailbox. If a hold was placed on the mailbox when it was deleted, the content preserved in the mailbox is still subject to a compliance boundary or search permissions filter based on the last time the Azure AD attributes were synchronized before the mailbox was deleted.

  Additionally, the synchronization between the user's mailbox and OneDrive account will cease if the mailbox is de-licensed or soft-deleted. The last stamped value of the compliance attribute for the OneDrive account will remain in effect.
- The compliance attribute is synchronized from a user's Exchange mailbox to their OneDrive account every seven days. As previously stated, this synchronization only occurs when the user is assigned both an Exchange Online and SharePoint Online license and the user's mailbox is at least 10 MB.
- If compliance boundaries and search permissions filters are implemented for both a user's mailbox and OneDrive account, then we recommend that you don't delete a user's mailbox and not their OneDrive account. In other words, if you delete a user's mailbox, you should also remove the user's OneDrive account.
- There are situations (such as a returning employee) where a user might have two or more OneDrive accounts. In these cases, only the primary OneDrive account associated with the user in Azure AD will be synchronized.

- Compliance boundaries and search permissions filters depend on attributes being stamped on content in Exchange, OneDrive, and SharePoint and the subsequent indexing of this stamped content.
- We don't recommend using exclusion filters (such as using `-not()` in a search permissions filter) for a content-based compliance boundary. Using an exclusion filter can have unexpected results if content with recently updated attributes hasn't been indexed.

Frequently asked questions

Who can create and manage search permissions filters (using New-ComplianceSecurityFilter and Set-ComplianceSecurityFilter cmdlets)?

To create, view, and modify search permissions filters, you have to be a member of the Organization Management role group in the Security & Compliance Center.

If an eDiscovery manager is assigned to more than one role group that spans multiple agencies, how do they search for content in one agency or the other?

The eDiscovery manager can add parameters to their search query that restrict the search to a specific agency. For example, if an organization has specified the **CustomAttribute10** property to differentiate agencies, they can append the following to their search query to search mailboxes and OneDrive accounts in a specific agency: `CustomAttribute10:<value> AND Site_ComplianceAttribute:<value>`.

What happens if the value of the attribute that's used as the compliance attribute in a search permissions filter is changed?

It takes up to three days for a search permissions filter to enforce the compliance boundary if the value of the attribute that's used in the filter is changed. For example, in the Contoso scenario let's say that a user in the Fourth Coffee agency is transferred to the Coho Winery agency. As a result, the value of the **Department** attribute on the user object is changed from *FourthCoffee* to *CohoWinery*. In this situation, Fourth Coffee eDiscovery and investors will get search results for that user for up three days after the attribute is changed. Similarly, it takes up to three days before Coho Winery eDiscovery managers and investigators get search results for the user.

Can an eDiscovery manager see content from two separate compliance boundaries?

Yes, this can be done when searching Exchange mailboxes by adding the eDiscovery manager to role groups that have visibility to both agencies. However when searching SharePoint sites and OneDrive accounts, an eDiscovery manager can search for content in different compliance boundaries only if the agencies are in the same region or geo location. **Note:** This limitation for sites doesn't apply in Advanced eDiscovery because searching for content in SharePoint and OneDrive isn't bound by geographic location.

Do search permissions filters work for eDiscovery case holds, Microsoft 365 retention policies, or DLP?

No, not at this time.

If I specify a region to control where content is exported, but I don't have a SharePoint organization in that region, can I still search SharePoint?

If the region specified in the search permissions filter doesn't exist in your organization, the default region will be searched.

What is the maximum number of search permissions filters that can be created in an organization?

There is no limit to the number of search permissions filters that can be created in an organization. However, search performance will be impacted when there are more than 100 search permissions filters. To keep the

number of search permissions filters in your organization as small as possible, create filters that combine rules for Exchange, SharePoint, and OneDrive into a single search permissions filter whenever possible.

Decryption in Microsoft 365 eDiscovery tools

2/18/2021 • 3 minutes to read • Edit Online

Encryption is an important part of your file protection and information protection strategy. Organizations of all types use encryption technology to protect sensitive content within their organization and ensure that only the right people have access to that content.

To execute common eDiscovery tasks on encrypted content, eDiscovery managers were required to decrypt email message content as it was exported from content searches, Core eDiscovery cases, and Advanced eDiscovery cases. Content encrypted with Microsoft encryption technologies wasn't available for review until after it was exported.

To make it easier to manage encrypted content in the eDiscovery workflow, Microsoft 365 eDiscovery tools now incorporate decryption of encrypted files that are attached to email messages and sent in Exchange Online. Additionally, encrypted documents stored in SharePoint Online and OneDrive for Business are decrypted in Advanced eDiscovery.

Prior to this new capability, only the content of an email message protected by rights management (and not attached files) were decrypted. Encrypted documents in SharePoint and OneDrive couldn't be decrypted during the eDiscovery workflow. Now, if a file that's encrypted with a Microsoft encryption technology is attached to an email message or located on a SharePoint or OneDrive account, those encrypted items are decrypted when the search results are prepared for preview, added to a review set in Advanced eDiscovery, and exported. This allows eDiscovery managers to view the content of encrypted email attachments and site documents when previewing search results, and review them after they have been added to a review set in Advanced eDiscovery.

Supported encryption technologies

Microsoft eDiscovery tools support items encrypted with Microsoft encryption technologies. These technologies include Office Message Encryption, Azure Rights Management, and Microsoft Information Protection (specifically sensitivity labels). For more information about Microsoft encryption technologies, see Encryption. Content encrypted by third-party encryption technologies isn't supported. For example, previewing or exporting content encrypted with non-Microsoft technologies isn't supported.

eDiscovery activities that support encrypted items

The following table identifies the supported tasks that can be performed in Microsoft 365 eDiscovery tools on encrypted files attached to email massages and encrypted documents in SharePoint and OneDrive. These supported tasks can be performed on encrypted files that match the criteria of a search. A value of `N/A` indicates the functionality isn't available in the corresponding eDiscovery tool.

EDISCOVERY TASK	CONTENT SEARCH	CORE EDISCOVERY	ADVANCED EDISCOVERY
Search for content in encrypted files in email and sites	Yes	Yes	Yes
Preview encrypted files attached to email	Yes	Yes	Yes

EDISCOVERY TASK	CONTENT SEARCH	CORE EDISCOVERY	ADVANCED EDISCOVERY
Preview encrypted documents in SharePoint and OneDrive	No	No	Yes
Review encrypted files in a review set	N/A	N/A	Yes
Export encrypted files attached to email	Yes	Yes	Yes
Export encrypted documents in SharePoint and OneDrive	No	No	Yes

Note: eDiscovery doesn't support encrypted files in SharePoint and OneDrive when a sensitivity label that applied the encryption is configured with either of the following settings:

- Users can assign permissions when they manually apply the label to a document. This is sometimes referred to as *user-defined permissions*.
- User access to the document has an expiration setting that is set to a value other than **Never**.

For more information about these settings, see the "Configure encryption settings" section in Restrict access to content by using sensitivity labels to apply encryption.

Documents encrypted with the previous settings can still be returned by an eDiscovery search. This may happen when a document property (such as the title, author, or modified date) matches the search criteria. Although these documents might be included in search results, they can't be previewed or reviewed. These documents will also remain encrypted when they're exported in Advanced eDiscovery.

Requirements for decryption in eDiscovery

You have to be assigned the RMS Decrypt role to preview, review, and export files encrypted with Microsoft encryption technologies. You also have to be assigned this role to review and query encrypted files that are added to a review set in Advanced eDiscovery.

This role is assigned by default to the eDiscovery Manager role group on the **Permissions** page in the Office 365 Security & Compliance Center. For more information about the RMS Decrypt role, see Assign eDiscovery permissions.

Collect eDiscovery diagnostic information

2/18/2021 • 3 minutes to read • Edit Online

Occasionally Microsoft Support engineers require specific information about your issue when you open a support case related to Core eDiscovery or Advanced eDiscovery. This article provides guidance on how to collect diagnostic information to help support engineers investigate and resolve issues. Typically, you don't need to collect this information until asked to do so by a Microsoft Support engineer.

> **IMPORTANT**
>
> The output from the cmdlets and diagnostic information described in this article may include sensitive information about litigation or internal investigations in your organization. Before sending the raw diagnostic information to Microsoft Support, you should review the information and redact any sensitive information (such as names or other information about parties to litigation or investigation) by replacing it with `xxxxxxx`. Using this method will also indicate to the Microsoft Support engineer that information was redacted.

Collect diagnostic information for Core eDiscovery

Collecting diagnostic information for Core eDiscovery is cmdlet-based, so you'll have to use Security & Compliance Center PowerShell. The following PowerShell examples will run cmdlets and then save the output to a specified text file. In most support cases, you should only have to run one of these commands.

To run the following cmdlets, connect to Security & Compliance Center PowerShell. After you're connected, run one or more of the following commands and be sure to replace placeholders with the actual object names.

After reviewing the generated text file and redacting sensitive information, send it to the Microsoft Support engineer working on your case.

> **NOTE**
>
> You can also run the commands in this section to collect diagnostic information for the searches and exports listed on the **Content search** page in the Microsoft 365 compliance center.

Collect information about searches

The following command collects information that's helpful when investigating issues with a Content search or a search associated with a Core eDiscovery case.

```
Get-ComplianceSearch "<Search name>" | FL > "ComplianceSearch.txt"
```

Collect information about search actions

The following command collects information to investigate problems with previewing, exporting, or purging the results of a Content search or a search associated with a Core eDiscovery case. You can identify the name of the search action by clicking an export that's listed on the **Exports** tab. To identify the names of preview and purge actions, you can run the **Get-ComplianceSearchAction** cmdlet to display a list of all actions. The format for the search action name is constructed by appending `_Preview`, `_Export`, or `_Purge` to the name of the corresponding search.

```
Get-ComplianceSearchAction "<Search action name>" | FL > "ComplianceSearchAction.txt"
```

Collect information about eDiscovery holds

When an eDiscovery hold associated with a Core eDiscovery case isn't functioning as expected, run the following command to collect information about the Case Hold Policy and associated Case Hold Rule for the eDiscovery hold. The *Case hold policy name* in the following command is the same as the name of the eDiscovery hold. You can identify this name on the **Holds** tabs in the Core eDiscovery case.

```
Get-CaseHoldPolicy "<Case hold policy name>" | %{"--CaseHoldPolicy--";$_|FL;"--CaseHoldRule--";Get-
CaseHoldRule -Policy $_.Name | FL} > "eDiscoveryCaseHold.txt"
```

Collect all case information

Sometimes, it's not apparent what information is required by Microsoft Support to investigate your issue. In this situation, you can collect all of the diagnostics information for a Core eDiscovery case. The *Core eDiscovery case name* in the following command is the same as the name of a case that's displayed on the **Core eDiscovery** page in the Microsoft 365 compliance center.

```
Get-ComplianceCase "<Core eDiscovery case name>"| %{"$($_.Name)";"`t==Searches==";Get-ComplianceSearch -Case
$_.Name | FL;"`t==Search Actions==";Get-ComplianceSearchAction -Case $_.Name |FL;"`t==Holds==";Get-
CaseHoldPolicy -Case $_.Name | %{$_|FL;"`t`t ==$($_.Name) Rules==";Get-CaseHoldRule -Policy $_.Name | FL}} >
"eDiscoveryCase.txt"
```

Collect diagnostic information for Advanced eDiscovery

The **Settings** tab in an Advanced eDiscovery case lets you quickly copy the diagnostic information for the case. The diagnostic information is saved to the clipboard so you can paste it to a text file and send to Microsoft Support.

1. Go to https://compliance.microsoft.com and then click **Show all** > **eDiscovery** > **Advanced**.
2. Select a case and then click the **Settings** tab.
3. Under **Case Information**, click **Select**.
4. On the flyout page, click **Copy diagnostic information** to copy the info to the clipboard.
5. Open a text file (in Notepad) and then paste the information in the text file.
6. Save the text file and name it something like `AeD Diagnostic Info YYYY.MM.DD` (for example, `AeD Diagnostic Info 2020.11.03`).

After reviewing the file and redacting sensitive information, send it to the Microsoft Support engineer working on your case.

Search for eDiscovery activities in the audit log

11/2/2020 • 18 minutes to read • Edit Online

Content Search and eDiscovery-related activities (for Core eDiscovery and Advanced eDiscovery) that are performed in Security & Compliance Center or by running the corresponding PowerShell cmdlets are logged in the audit log. Events are logged when administrators or eDiscovery managers (or any user assigned eDiscovery permissions) perform the following Content Search and Core eDiscovery tasks in the Security & Compliance Center:

- Creating and managing Core and Advanced eDiscovery cases
- Creating, starting, and editing Content Searches
- Performing Content Search actions, such as previewing, exporting, and deleting search results
- Managing custodians and review sets in Advanced eDiscovery
- Configuring permissions filtering for Content Search
- Managing the eDiscovery Administrator role

> **IMPORTANT**
> The activities described in this article are only the result of eDiscovery tasks performed by using the Security & Compliance Center. eDiscovery tasks that were performed by using the In-Place eDiscovery tool in Exchange Online or the eDiscovery Center in SharePoint Online aren't included.

For more information about searching the audit log, the permissions that are required, and exporting search results, see Search the audit log in the Security & Compliance Center.

How to search for and view eDiscovery activities

Currently, you have to do a few specific things to view eDiscovery activities in the audit log. Here's how.

1. Go to https://protection.office.com.
2. Sign in using your work or school account.
3. In the left pane, click **Search**, and then click **Audit log search**.
4. In the **Activities** drop-down list, under **eDiscovery activities** or **Advanced eDiscovery activities**, click one or more activities to search for.

   > **NOTE**
   > The **Activities** drop-down list also includes a group of activities named **eDiscovery cmdlet activities** that will return records from the cmdlet audit log.

5. Select a date and time range to display eDiscovery events that occurred within that period.
6. In the **Users** box, select one or more users to display search results for. Leave this box blank to return entries for all users.
7. Click **Search** to run the search using your search criteria.

8. After the search results are displayed, you can click **Filter results** to filter or sort the resulting activity records. Unfortunately, you can't use filtering to explicitly exclude certain activities.

9. To view details about an activity, click the activity record in the list of search results.

   A **Details** fly out page is displayed that contains the detailed properties from the event record. To display additional details, click **More information**. For a description of these properties, see the Detailed properties for eDiscovery activities section.

10. If desired, you can export the audit log search results to a CSV file, and then use the Excel Power Query feature to format and filter these records. For more information, see Export, configure, and view audit log records.

eDiscovery activities

The following table describes the Content Search and Core eDiscovery activities that are logged when an administrator or eDiscovery manager performs an eDiscovery-related activity using the Security & Compliance Center or running the corresponding cmdlet in Security & Compliance Center PowerShell. Note also that some activities performed in Advanced will be returned when you search for activities in this list.

> **NOTE**
> The eDiscovery activities described in this section provide similar information to the eDiscovery cmdlet activities described in the next section. We recommend that you use the eDiscovery activities described in this section because they will appear in the audit log search results within 30 minutes. It takes up to 24 hours for the eDiscovery cmdlet activities to appear in audit log search results.

FRIENDLY NAME	OPERATION	CORRESPONDING CMDLET	DESCRIPTION
Added member to eDiscovery case	CaseMemberAdded	Add-ComplianceCaseMember	A user was added as a member of an eDiscovery case. As a member of a case, a user can perform various case-related tasks depending on whether they have been assigned the necessary permissions.
Changed content search	SearchUpdated	Set-ComplianceSearch	An existing content search was changed. Changes can include adding or removing content locations or editing the search query.
Changed eDiscovery administrator membership	CaseAdminUpdated	Update-eDiscoveryCaseAdmin	The list of eDiscovery Administrators in your organization was changed. This activity is logged when the list of eDiscovery Administrators is replaced with a group of new users. If a single user is added or removed, the CaseAdminAdded operation is logged.

FRIENDLY NAME	OPERATION	CORRESPONDING CMDLET	DESCRIPTION
Changed eDiscovery case	CaseUpdated	Set-ComplianceCase	An eDiscovery case was changed. Changes include closing an open case or reopening a closed case.
Changed eDiscovery case membership	CaseMemberUpdated	Update-ComplianceCaseMember	The membership list of an eDiscovery case was changed. This activity is logged when all members are replaced with a group of new users. If a single member is added or removed, CaseMemberAdded or CaseMemberRemoved operation is logged.
Changed search permissions filter	SearchPermissionUpdated	Set-ComplianceSecurityFilter	A search permissions filter was changed.
Changed search query for eDiscovery case hold	HoldUpdated	Set-CaseHoldRule	A query-based hold associated with an eDiscovery case was changed. Possible changes include editing the query or date range for a query-based hold.
Content search preview item downloaded	PreviewItemDownloaded	N/A	A user downloaded an item to their local computer (by clicking the **Download original item** link) when previewing search results.
Content search preview item listed	PreviewItemListed	N/A	A user clicked **Preview search results** to display the preview search results page, which lists up to 1000 items from the results of a Content Search.
Content search preview item viewed	PreviewItemRendered	N/A	An eDiscovery manager viewed an item by clicking it when previewing search results.
Created content search	SearchCreated	New-ComplianceSearch	A new content search was created.
Created eDiscovery administrator	CaseAdminAdded	Add-eDiscoveryCaseAdmin	A user was added as an eDiscovery Administrator in the organization.

FRIENDLY NAME	OPERATION	CORRESPONDING CMDLET	DESCRIPTION
Created eDiscovery case	CaseAdded	New-ComplianceCase	An eDiscovery case was created. When a case is created, you only have to give it a name. Other case-related tasks such as adding members, creating holds, and creating content searches associated with the case result in additional events being logged.
Created search permissions filter	SearchPermissionCreated	New-ComplianceSecurityFilter	A search permissions filter was created.
Created search query for eDiscovery case hold	HoldCreated	New-CaseHoldRule	A query-based hold associated with an eDiscovery case was created.
Deleted content search	SearchRemoved	Remove-ComplianceSearch	An existing content search was deleted.
Deleted eDiscovery administrator	CaseAdminRemoved	Remove-eDiscoveryCaseAdmin	An eDiscovery Administrator was deleted from your organization.
Deleted eDiscovery case	CaseRemoved	Remove-ComplianceCase	An eDiscovery case was deleted. Any hold associated with the case has to be removed before the case can be deleted.
Deleted search permissions filter	SearchPermissionRemoved	Remove-ComplianceSecurityFilter	A search permissions filter was deleted.
Deleted search query for eDiscovery case hold	HoldRemoved	Remove-CaseHoldRule	A query-based hold associated with an eDiscovery case was deleted. Removing the query from the hold is often the result of deleting a hold. When a hold or a hold query is deleted, the content locations that were on hold are released.
Downloaded export of content search	SearchExportDownloaded	N/A	A user downloaded the results of a content search to their local computer. A **Started export of content search** activity has to be initiated before search results can be downloaded.
Previewed results of content search	SearchPreviewed	N/A	A user previewed the results of a content search.

FRIENDLY NAME	OPERATION	CORRESPONDING CMDLET	DESCRIPTION
Purged results of content search	SearchResultsPurged	New-ComplianceSearchAction	A user purged the results of a Content Search by running the **New-ComplianceSearchAction -Purge** command.
Removed analysis of content search	RemovedSearchResultsSentToZoom	Remove-ComplianceSearchAction	A content search prepare action (to prepare search results for Advanced eDiscovery) was deleted. If the preparation action was less than two weeks old, the search results that were prepared for Advanced eDiscovery were deleted from the Microsoft Azure storage area. If the preparation action was older than 2 weeks, then this event indicates that only the corresponding preparation action was deleted.
Removed export of content search	RemovedSearchExported	Remove-ComplianceSearchAction	A content search export action was deleted. If the export action was less than two weeks old, the search results that were uploaded to the Microsoft Azure storage area were deleted. If the export action was older than 2 weeks, then this event indicates that only the corresponding export action was deleted.
Removed member from eDiscovery case	CaseMemberRemoved	Remove-ComplianceCaseMember	A user was removed as a member of an eDiscovery case.
Removed preview results of content search	RemovedSearchPreviewed	Remove-ComplianceSearchAction	A content search preview action was deleted.
Removed purge action performed on content search	RemovedSearchResultsPurged	Remove-ComplianceSearchAction	A content search purge action was deleted.
Removed search report	SearchReportRemoved	Remove-ComplianceSearchAction	A content search export report action was deleted.
Started analysis of content search	SearchResultsSentToZoom	New-ComplianceSearchAction	The results of a content search were prepared for analysis in Advanced eDiscovery.

FRIENDLY NAME	OPERATION	CORRESPONDING CMDLET	DESCRIPTION
Started content search	SearchStarted	Start-ComplianceSearch	A content search was started. When you create or change a content search by using the Security & Compliance Center GUI, the search is automatically started. If you create or change a search by using the **New-ComplianceSearch** or **Set-ComplianceSearch** cmdlet, you have to run the **Start-ComplianceSearch** cmdlet to start the search.
Started export of content search	SearchExported	New-ComplianceSearchAction	A user exported the results of a content search.
Started export report	SearchReport	New-ComplianceSearchAction	A user exported a content search report.
Stopped content search	SearchStopped	Stop-ComplianceSearch	A user stopped a content search.
(none)	CaseViewed	Get-ComplianceCase	A user viewed the list of cases on the **eDiscovery** page in the security and compliance center or by running the cmdlet.
(none)	SearchViewed	Get-ComplianceSearch	A user viewed the list on content searches (listed on the **Searches** tab) in the security and compliance center or by running the cmdlet. This activity is also logged when a user views the list of content searches associated with an eDiscovery case (by clicking the **Searches** tab in a case) or by running the **Get-ComplianceSearch -Case** command.
(none)	ViewedSearchExported	Get-ComplianceSearchAction -Export	A user viewed the list of content search export jobs (listed on the **Exports** tab) in the security and compliance center or by running the cmdlet. This activity is also logged when a user views the list of export jobs in an eDiscovery case (listed on the **Exports** tab in a case) or by running the **Get-ComplianceSearchAction -Case -Export** command.

FRIENDLY NAME	OPERATION	CORRESPONDING CMDLET	DESCRIPTION
(none)	ViewedSearchPreviewed	Get-ComplianceSearchAction -Preview	A user previews the results of a content search in the security and compliance center or by running the cmdlet.

Advanced eDiscovery activities

The following table describes the Advanced eDiscovery activities logged in the audit log. These activities (in addition to relevant eDiscovery activities can be used to help you track the progression of activity in an Advanced eDiscovery case.

FRIENDLY NAME	OPERATION	DESCRIPTION
Added data to another review set	AddWorkingSetQueryToWorkingSet	User added documents from one review set to a different review set.
Added data to review set	AddQueryToWorkingSet	User added the search results from a content search associated with an Advanced eDiscovery case to a review set.
Added non-Microsoft 365 data to review set	AddNonOffice365DataToWorkingSet	User added non-Microsoft 365 data to a review set.
Added remediated documents to review set	AddRemediatedData	User uploads documents that had indexing errors that were fixed to a review set.
Analyzed data in review set	RunAlgo	User ran analytics on the documents in a review set.
Annotated document in review set	AnnotateDocument	User annotated a document in a review set. Annotation includes redacting content in a document.
Compared load sets	LoadComparisonJob	User compared two different load sets in a review set. A load set is when data from a content search that associated with the case is added to a review set.
Converted redacted documents to PDF	BurnJob	User converted all the redacted documents in a review set to PDF files.
Created review set	CreateWorkingSet	User created a review set.
Created review set search	CreateWorkingSetSearch	User created a search query that searches the documents in a review set.

FRIENDLY NAME	OPERATION	DESCRIPTION
Created tag	CreateTag	User created a tag group in a review set. A tag group can contain one or more child tags. These tags are then used to tag documents in the review set.
Deleted review set search	DeleteWorkingSetSearch	User deleted a search query in a review set.
Deleted tag	DeleteTag	User deleted a tag or a tag group in a review set.
Downloaded document	DownloadDocument	User downloaded a document from a review set.
Edited tag	UpdateTag	User changed a tag in a review set.
Exported documents from review set	ExportJob	User exported documents from a review set.
Modified case setting	UpdateCaseSettings	User modified the settings for a case. Case settings include case information, access permissions, and settings that control search and analytics behavior.
Modified review set search	UpdateWorkingSetSearch	User edited a search query in a review set.
Previewed review set search	PreviewWorkingSetSearch	User previewed the results of a search query in a review set.
Remediated error documents	ErrorRemediationJob	User fixes files that contained indexing errors.
Tagged document	TagFiles	User tags a document in a review set.
Tagged results of a query	TagJob	User tags all of the documents that match the criteria of search query in a review set.
Viewed document in review set	ViewDocument	User viewed a document in a review set.

eDiscovery cmdlet activities

The following table lists the cmdlet audit log records that are logged when an administrator or user performs an eDiscovery-related activity by using the Security & Compliance Center or by running the corresponding cmdlet in remote PowerShell that's connected to your organization's Security & Compliance Center. The detailed information in the audit log record is different for the cmdlet activities listed in this table and the eDiscovery activities described in the previous section.

As previously stated, it takes up to 24 hours for eDiscovery cmdlet activities to appear in the audit log search results.

TIP

The cmdlets in the **Operation** column in the following table are linked to the corresponding cmdlet help topic on TechNet. Go to the cmdlet help topic for a description of the available parameters for each cmdlet. The parameter and the parameter value that were used with a cmdlet are included in the audit log entry for each eDiscovery cmdlet activity that's logged.

FRIENDLY NAME	OPERATION (CMDLET)	DESCRIPTION
Created hold in eDiscovery case	New-CaseHoldPolicy	A hold was created for an eDiscovery case. A hold can be created with or without specifying a content source. If content sources are specified, they'll be identified in the audit log entry.
Deleted hold from eDiscovery case	Remove-CaseHoldPolicy	A hold that is associated with an eDiscovery case was deleted. Deleting a hold releases all of the content locations from the hold. Deleting the hold also results in deleting the case hold rules associated with the hold (see **Remove-CaseHoldRule** below).
Changed hold in eDiscovery case	Set-CaseHoldPolicy	A hold that is associated with an eDiscovery was changed. Possible changes include adding or removing content locations or turning off (disabling) the hold.
Created search query for eDiscovery case hold	New-CaseHoldRule	A query-based hold associated with an eDiscovery case was created.
Deleted search query for eDiscovery case hold	Remove-CaseHoldRule	A query-based hold associated with an eDiscovery case was deleted. Removing the query from the hold is often the result of deleting a hold. When a hold or a hold query is deleted, the content locations that were on hold are released.
Changed search query for eDiscovery case hold	Set-CaseHoldRule	A query-based hold associated with an eDiscovery case was changed. Possible changes include editing the query or date range for a query-based hold.
Created eDiscovery case	New-ComplianceCase	An eDiscovery case was created. When a case is created, you only have to give it a name. Other case-related tasks such as adding members, creating holds, and creating content searches associated with the case result in additional events being logged.
Deleted eDiscovery case	Remove-ComplianceCase	An eDiscovery case was deleted. Any hold associated with the case has to be removed before the case can be deleted.

FRIENDLY NAME	OPERATION (CMDLET)	DESCRIPTION
Changed eDiscovery case	Set-ComplianceCase	An eDiscovery case was changed. Changes include closing an open case or reopening a closed case.
Added member to eDiscovery case	Add-ComplianceCaseMember	A user was added as a member of an eDiscovery case. As a member of a case, a user can perform various case-related tasks depending on whether they have been assigned the necessary permissions.
Removed member from eDiscovery case	Remove-ComplianceCaseMember	A user was removed as a member of an eDiscovery case.
Changed eDiscovery case membership	Update-ComplianceCaseMember	The membership list of an eDiscovery case was changed. This activity is logged when all members are replaced with a group of new users. If a single member is added or removed, the **Add-ComplianceCaseMember** or **Remove-ComplianceCaseMember** operation is logged.
Created content search	New-ComplianceSearch	A new content search was created.
Deleted content search	Remove-ComplianceSearch	An existing content search was deleted.
Changed content search	Set-ComplianceSearch	An existing content search was changed. Changes can include adding or removing content locations that are searched and editing the search query.
Started content search	Start-ComplianceSearch	A content search was started. When you create or change a content search by using the Security & Compliance Center GUI, the search is automatically started. If you create or change a search by using the **New-ComplianceSearch** or **Set-ComplianceSearch** cmdlet, you have to run the **Start-ComplianceSearch** cmdlet to start the search.
Stopped content search	Stop-ComplianceSearch	A content search that was running was stopped.
Created content search action	New-ComplianceSearchAction	A content search action was created. Content search actions include previewing search results, exporting search results, preparing search results for analysis in Advanced eDiscovery, and permanently deleting items that match the search criteria of a content search.
Deleted content search action	Remove-ComplianceSearchAction	A content search action was deleted.

FRIENDLY NAME	OPERATION (CMDLET)	DESCRIPTION
Created search permissions filter	New-ComplianceSecurityFilter	A search permissions filter was created.
Deleted search permissions filter	Remove-ComplianceSecurityFilter	A search permissions filter was deleted.
Changed search permissions filter	Set-ComplianceSecurityFilter	A search permissions filter was changed.
Created eDiscovery administrator	Add-eDiscoveryCaseAdmin	A user was added as an eDiscovery Administrator in your organization.
Deleted eDiscovery administrator	Remove-eDiscoveryCaseAdmin	An eDiscovery Administrator was deleted from your organization.
Changed eDiscovery administrator membership	Update-eDiscoveryCaseAdmin	The list of eDiscovery Administrators in your organization was changed. This activity is logged when the list of eDiscovery Administrators is replaced with a group of new users. If a single user is added or removed, the **Add-eDiscoveryCaseAdmin** or **Remove-eDiscoveryCaseAdmin** operation is logged.

Detailed properties for eDiscovery activities

The following table describes the properties that are included when you click **More information** on the **Details** page for an eDiscovery activity listed in the search results. These properties are also included in the CSV file when you export the audit log search results. An audit log record for an eDiscovery activity won't include every detailed property listed below.

> **TIP**
> When you export the search results, the CSV file contains a column named **Detail**, which contains the detailed properties described in the following table in a multi-value property. You can use the Power Query feature in Excel to split this column into multiple columns so that each property will have its own column. This will let you sort and filter on one or more of these properties. For more information, see the "Export the search results to a file" section in Search the audit log.

PROPERTY	DESCRIPTION
Case	The identity (GUID) of the eDiscovery case that was created, changed, or deleted.
ClientApplication	eDiscovery cmdlet activities have a value of **EMC** for this property. This indicates the activity was performed by using the Security & Compliance Center GUI or running the cmdlet in PowerShell.
ClientIP	The IP address of the device that was used when the activity was logged. The IP address is displayed in either an IPv4 or IPv6 address format.
ClientRequestId	For eDiscovery activities, this property is typically blank.

PROPERTY	DESCRIPTION
CmdletVersion	The build number for the version of the Security & Compliance Center running in your organization.
CreationTime	The date and time in Coordinated Universal Time (UTC) when the eDiscovery activity was completed.
EffectiveOrganization	The name of the Microsoft 365 organization.
ExchangeLocations	The Exchange Online mailboxes that are included in a content search or placed on hold in an eDiscovery case.
Exclusions	Mailbox or site locations that are excluded from a content search or a hold in an eDiscovery case.
ExtendedProperties	Additional properties from a content search, a content search action, or hold in an eDiscovery case, such as the object GUID and the corresponding cmdlet and cmdlet parameters that were used when the activity was performed.
Id	The ID of the report entry. The ID uniquely identifies the audit log entry.
NonPIIParameters	A list of the parameters (without any values) that were used with the cmdlet identified in the Operation property. The parameters listed in this property are the same as those listed in the Parameters property.
ObjectId	The GUID or name of the object (for example, a Content Search or an eDiscovery case) that was created, changed, or deleted by the activity listed in the Operation property. This object is also identified in the Item column in the audit log search results.
ObjectType	The type of eDiscovery object that the user created, deleted, or modified; for example, a content search action (preview, export, or purge), an eDiscovery case, or a content search.
Operation	The name of the operation that corresponds to the eDiscovery activity that was performed.
OrganizationId	The GUID for your Microsoft 365 organization.
Parameters	The name and value for the parameters that were used with the corresponding cmdlet.
PublicFolderLocations	The public folder locations in Exchange Online that are included in a content search or placed on hold in an eDiscovery case.
Query	The search query associated with the activity, such as a content search or a query-based hold.

PROPERTY	DESCRIPTION
RecordType	The type of operation indicated by the record. The value of **18** indicates an event related to an activity listed in the eDiscovery cmdlet activities section. A value of **24** indicates an event related to an activity listed in the How to search for and view eDiscovery activities section.
ResultStatus	Indicates whether the action (specified in the Operation property) was successful or not.
SecurityComplianceCenterEventType	Indicates that the activity was a Security & Compliance Center event. All eDiscovery activities will have a value of **0** for this property.
SharepointLocations	The SharePoint Online sites that are included in a content search or placed on hold in an eDiscovery case.
StartTime	The date and time in Coordinated Universal Time (UTC) when the eDiscovery activity was started.
UserId	The user who performed the activity (specified in the Operation property) that resulted in the record being logged. Records for eDiscovery activity performed by system accounts (such as NT AUTHORITY\SYSTEM) are also included in the audit log.
UserKey	An alternative ID for the user identified in the UserId property. For eDiscovery activities, the value for this property is typically the same as the UserId property.
UserServicePlan	The subscription used by your organization. For eDiscovery activities, this property is typically blank.
UserType	The type of user that performed the operation. The following values indicate the user type. 0 A regular user. 2 An administrator in your organization. 3 A Microsoft datacenter administrator or datacenter system account. 4 A system account. 5 An application. 6 A service principal.
Version	Indicates the version number of the activity (identified by the Operation property) that's logged.
Workload	Theservice where the activity occurred. For eDiscovery activities, the value is **SecurityComplianceCenter**.

Troubleshoot AzCopy in Advanced eDiscovery

11/2/2020 • 2 minutes to read • Edit Online

When loading non-Microsoft 365 data or documents for error remediation in Advanced eDiscovery, the user interface supplies an Azure AzCopy command that contains parameters with the location of where the files that you want to upload are stored and the Azure storage location that the files will be uploaded to. To upload your documents, you copy this command and then run it in a Command Prompt on your local computer. The follow screenshot shows an example of an AzCopy command:

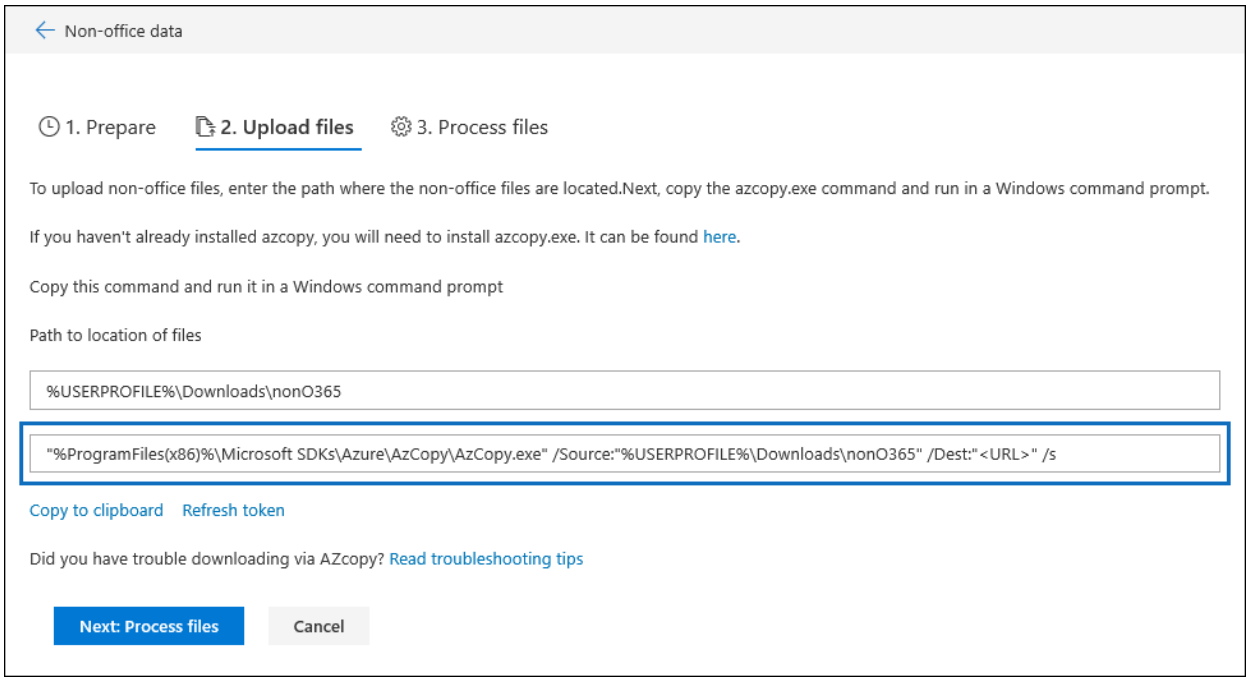

Usually the command that's provided works when you run it. However, there may be cases when the command that's displayed will not run successfully. Here's a few possible reasons.

The supported version of AzCopy isn't installed on the local computer

At this time, you must use AzCopy v8.1 to load non-Microsoft 365 data in Advanced eDiscovery. The AzCopy command that's displayed on the **Upload files** page shown in the previous screenshot returns an error if you're not using AzCopy v8.1. To install this version, see Transfer data with the AzCopy v8.1 on Windows.

AzCopy isn't installed on the local computer or it's not installed in the default location

If AzCopy isn't installed or it's installed in a location other than the default install location (which is `%ProgramFiles(x86)%`), you may receive the following error when you run the AzCopy command:

> The system cannot find the path specified.

If AzCopy isn't installed on the local computer, you can find installation information in Transfer data with the AzCopy v8.1 on Windows. Be sure to install it in the default location.

If AzCopy is installed, but it's installed in a location different than the default location, you can copy the command, paste it to a text file, and then change the path to the location where AzCopy is installed. For example, if Azcopy is located in `%ProgramFiles%`, then you can change the first part of the command from

`%ProgramFiles(x86)%\Microsoft SDKs\Azure\AzCopy.exe` to `%ProgramFiles%\Microsoft SDKs\Azure\AzCopy`. After you make this change, copy it from the text file and then run it a Command Prompt.

> **TIP**
> If AzCopy is installed in a location other then the default install location, consider uninstalling it and then re-installing it in the default location. This will help prevent this issue in the future.

Retirement of legacy eDiscovery tools

2/18/2021 • 16 minutes to read • Edit Online

IMPORTANT

Microsoft has been evaluating the public health situation, and we understand the impact this is having on our customers. We want to be strong partners and responsible global citizens. To ease one of the many burdens you are facing, we are going to delay the scheduled retirement for the legacy eDiscovery tools described in this article by three months. **The updated retirement dates are reflected below.**

Over the years, Microsoft has provided eDiscovery tools that let you search, preview, and export email content from Exchange Online. However, these tools no longer offer an effective way to search for non-Exchange content in other Microsoft 365 services, such as SharePoint Online and Microsoft 365 Groups. To address this, Microsoft offers other eDiscovery tools that help you search for a wide variety of Microsoft 365 content. And we've been working hard to incorporate the most current and powerful eDiscovery functionality in the Microsoft 365 compliance center. This allows organizations to respond to legal, internal, and other document requests for content across many Microsoft 365 services, including Exchange Online.

As a result of this new and improved eDiscovery functionality in the Microsoft 365 compliance center, we're retiring the following eDiscovery-related features and functionality related to searching for email content in Exchange Online and Microsoft 365:

- In-Place eDiscovery and In-Place Holds in the Exchange admin center.
- The Exchange Online PowerShell cmdlets that support In-Place eDiscovery and In-Place Holds (these cmdlets are collectively identified as **-MailboxSearch* cmdlets). This includes the following cmdlets:
  - New-MailboxSearch
  - Start-MailboxSearch
  - Stop-MailboxSearch
  - Set-MailboxSearch

  NOTE

  The Get-MailboxSearch and Remove-MailboxSearch cmdlets will be available after the other **-MailboxSearch* cmdlets are retired so that you can use them to help in your transition to other eDiscovery and hold tools. However, after a certain date (cited below) Microsoft Support will no longer supports these two cmdlets.
- The Search-Mailbox cmdlet in Exchange Online PowerShell.
- The following operations in the Exchange Web Services API:
  - GetSearchableMailboxes
  - SearchMailboxes
  - SetHoldOnMailboxes
  - GetHoldOnMailboxes
- Office 365 Advanced eDiscovery v1.0, which is the first version of Advanced eDiscovery that's accessed

through a Core eDiscovery case in the Office 365 Security & Compliance Center. The retirement of Advanced eDiscovery v1.0 doesn't impact your ability to create and manage Core eDiscovery cases.

> **NOTE**
> The eDiscovery functionality being retired only applies to cloud-based versions of Microsoft 365 and Office 365. eDiscovery functionality in on-premises versions of Exchange and SharePoint will still be supported until further notice.

The following sections in this article provide guidance about each feature being retired. This information including timelines and alternative tools that you can use instead of the retired tool.

In-Place eDiscovery and In-Place Holds in the Exchange admin center

As per the original announcement on July 1, 2017, the In-Place eDiscovery & Hold functionality in the Exchange admin center (EAC) is being retired. The In-Place eDiscovery & Holds page in the EAC allowed you to search, hold, and export content from Exchange Online. In-Place eDiscovery also let you copy search results to a discovery mailbox so that you or other eDiscovery managers could review content and make it available for legal, regulatory, and public requests.

Because all of these capabilities (except for copying search results to a discovery mailbox) are now available in the content search, eDiscovery and Advanced eDiscovery tools in the Microsoft 365 compliance center (with improved functionality, reliability, and support for a wide range of Microsoft 365 services), we recommend that you start using these tools as soon as possible. To help you in the transition to these other eDiscovery tools, the table below lists the tools you can use instead of In-Place eDiscovery and In-Place Hold.

Scope of affected organizations

- Office 365 and Microsoft 365 Enterprise organizations
- Office 365 and Microsoft 365 Education organizations
- Office 365 and Microsoft 365 Government organizations; this includes GCC, GCC High, and DoD
- Office 365 Germany

Timeline for retirement

- July 1, 2020: You won't be able to create new searches and holds, but you can still run, edit, and delete existing searches at your own risk. Microsoft Support will no longer In-Place eDiscovery & Holds in the EAC.
- October 1, 2020: The In-Place eDiscovery & Holds functionality in the EAC will be placed in a read-only mode. This means you'll only be able to remove existing searches and holds.

Alternative tools

The following table describes other tools that you can use to replace the existing functionality that's being retired.

FUNCTIONALITY	ALTERNATIVE TOOL	COMMENTS

FUNCTIONALITY	ALTERNATIVE TOOL	COMMENTS
Search, export, and hold for legal purposes	Core eDiscovery cases in the Microsoft 365 compliance center	Using the capabilities of core eDiscovery cases provide the functional parity to In-Place eDiscovery and In-Place Holds. This includes the following: • Search scales to millions of locations • Higher reliability for searching, exporting, and placing content on hold • Searching for content in for Exchange Online, SharePoint Online, OneDrive for Business, Skype for Business, Microsoft Teams, Yammer Groups, Microsoft 365 Groups, and other content stored in Office 365 applications
Hold for retention purposes	Retention policies in the Microsoft 365 compliance center	You can use Retention policies to retain content and, if desired, delete it after the retention period expires. Other capabilities include: • Applying policies to your entire organization • Applying policies to specific content locations such as Exchange Online, SharePoint Online, OneDrive for Business, Skype for Business, Microsoft Teams, and Office 365 Groups • Applying policies to specific users For more information, see Learn about retention policies and retention labels.

FUNCTIONALITY	ALTERNATIVE TOOL	COMMENTS
Copy email search results to a discovery mailbox for review	Review sets in Advanced eDiscovery v2.0	Reviewing content in Microsoft 365 has never been easier. Review sets have many great capabilities to help reduce the amount of time and data needed to review, including: • Perform fast search queries and filter content in a review set • Analyze content in a review set; this includes email threading, near-duplicate detection, Themes analysis, and Predictive coding • Tag documents in a review set • Tagging suggestions to help identify attorney client privilege content For more information, see Overview of the Advanced eDiscovery solution in Microsoft 365. Alternatively, you can export search results to PST files and then use Microsoft 365 Import service to import the PSTs to a discovery mailbox. For step-by-step instruction, see Use network upload to import PST files to Office 365.
Copy messages from one mailbox to a different mailbox	Assign permissions to a mailbox	To give a person access to another user's email (such as when an employee leaves your organization and you need to give another person access to the former employee's email), we recommended that you assign that person permissions to access the former employee's mailbox. So instead of copying mailbox items to another user mailbox or a shared mailbox, just assign a user permissions to access the source mailbox.
Restore items from the Recoverable Items folder	Restore-RecoverableItems	You can restore permanently deleted items (also known as *soft-deleted* items) in mailboxes, as long as the deleted item retention period for an item hasn't expired. For more information, see Recoverable Items folder in Exchange Online.

FAQs about In-Place eDiscovery and In-Place Holds

I use the copy search results functionality of In-Place eDiscovery & Holds in the EAC to copy

search results to a discovery mailbox for review by attorneys. What options do I have now?

There are two ways to replicate this functionality today. The first is to use review sets in Advanced eDiscovery v2.0. Review sets have many of the same capabilities you see in a traditional review tool like fast search of documents, tagging, email threading, near duplicate grouping, themes analysis, and predictive coding. If you still want to use discovery mailboxes for review, the second option is to export search results to PST files and then import the PST files to a discovery mailbox by using the PST import feature in the Microsoft compliance center.

How do I control which content locations (such as mailboxes or sites) that can an eDiscovery manager can search using the new tools?

The Microsoft 365 compliance center also uses compliance boundaries to control which content locations an eDiscovery Manager can search. Compliance boundaries are useful in government entities that need to stay within agency boundaries or multi-national corporations required to respect geographical boarders.

How can I move my current searches and holds to the Microsoft 365 compliance center?

It's possible to migrate In-Place eDiscovery searches and holds from the EAC by using PowerShell. For instructions, see Migrate searches and holds from the EAC to the Microsoft 365 compliance center.

*-MailboxSearch cmdlets

As per the original notice announced on July 1, 2017 in the Exchange admin center, the In-Place eDiscovery & Hold functionality and the corresponding ***-MailboxSearch** cmdlets are being retired. These cmdlets provide users the ability to search, hold, and export mailbox content for legal, regulatory, and public requests.

Because these capabilities are now available in the Microsoft 365 compliance center and Office 365 Security & Compliance Center PowerShell with improved performance and scalability, you should using these improved cmdlets. These cmdlets include *-ComplianceCase, *-ComplianceSearch, *-CaseHoldPolicy, *-CaseHoldRule, and *-ComplianceSearchAction.

Scope of affected organizations

- Office 365 and Microsoft 365 Enterprise organizations
- Office 365 and Microsoft 365 Education organizations
- Office 365 and Microsoft 365 Government organizations; this includes GCC, GCC High, and DoD
- Office 365 Germany

Timeline

- July 1, 2020: You won't be able to use **New-MailboxSearch** to create new In-Place eDiscovery searches and In-Place Holds, but you can still use cmdlets to run, edit, and delete existing searches and holds at your own risk. Microsoft Support will no longer provide assistance for these types of searches and holds.
- October 1, 2020: As previously stated, The In-Place eDiscovery & Holds functionality in the EAC will be placed in a read-only mode. That also means that you won't be able to use the **New-MailboxSearch**, **Start-MailboxSearch**, or **Set-MailboxSearch** cmdlets. You'll only be able to get and remove existing searches and holds.

Alternative tools

The following table describes other tools that you can use to replace the existing functionality that's being retired.

FUNCTIONALITY	ALTERNATIVE TOOLS	COMMENTS

FUNCTIONALITY	ALTERNATIVE TOOLS	COMMENTS
Search and export	*-ComplianceSearch *-ComplianceSearchAction *-ComplianceCase	The ComplianceSearch and ComplianceSearchAction cmdlets work together to help you search and export content. You can create a new search and view the search estimate by using the **New-**, **Get-**, and **Start-ComplianceSearch** cmdlets. Then you can use the **New-ComplianceSearchAction** cmdlet to export the search results. You'll still have to use the core eDiscovery tool in the Microsoft 365 compliance center to download those search results to your local computer. **Note:** If you use these cmdlets to create searches that aren't associated with a core eDiscovery case, these searches will be located on the **Content search** page in the Microsoft 365 compliance center.
Hold content in a mailbox	*-CaseHoldPolicy *-CaseHoldRule *-ComplianceCase	Holds in the Microsoft 365 compliance center must be associated with a ComplianceCase. First, create the compliance case, and then create a CaseHoldPolicy and a CaseHoldRule. **Note:** Creating a CaseHoldPolicy without a creating CaseHoldRule will render the hold inoperable until the CaseHoldRule is created and associated to the CaseHoldPolicy. See the cmdlet documentation for more information.
Copy search results to a discovery mailbox	None	There's no direct replacement for this functionality because it does not provide access to all Microsoft 365 services. See the following FAQ below for alternative solutions.
Copy messages from one mailbox to a different mailbox	Assign permissions to a mailbox	To give a person access to another user's email (such as when an employee leaves your organization and you need to give another person access to the former employee's email), we recommended that you assign that person permissions to access the former employee's mailbox. So instead of copying mailbox items to another user mailbox or a shared mailbox, just assign a user permissions to access the source mailbox.

FAQs about *-MailboxSearch cmdlets

We use Copy Search to export email messages or instant Messages for purposes other eDiscovery and legal investigations. What other options do we have after these cmdlets are retired?

The Microsoft Graph APIs provide a number of methods for extracting data for analysis and other purposes that are far more resilient and scalable than the using the ***-MailboxSearch** cmdlets.

How can I migrate my searches and holds to the Microsoft 365 compliance center?

It's possible to migrate In-Place eDiscovery searches and holds from the Exchange admin center by using a PowerShell script. For more information, see Migrate legacy eDiscovery searches and holds to the Microsoft 365 compliance center.

After the cmdlets are retired, will I still be able to remove or retrieve searches?

Yes, although we're removing the ability to create and modify searches, you'll still be able to use **Get-MailboxSearch** and **Remove-MailboxSearch** until further notice. However, the use of these cmdlets will be at your own risk after the retirement dates and Microsoft Support will no longer be able to provide assistance.

Search-Mailbox cmdlet

The **Search-Mailbox** cmdlet in Exchange Online PowerShell is being retired as originally announced in a warning in the cmdlet output starting back in 2018. The **Search-Mailbox** cmdlet was originally used to search a user's mailbox and purge malicious content. We recommend that you start using the **New-ComplianceSearch** and **New-ComplianceSearchAction** cmdlets in Office 365 Security & Compliance Center PowerShell to search for and purge content. For a built-in security experience, the Microsoft 365 security features provide robust threat protection for email and many other Microsoft services.

Scope of affected organizations

- Office 365 and Microsoft 365 Enterprise organizations
- Office 365 and Microsoft 365 Education organizations
- Office 365 and Microsoft 365 Government organizations; this includes GCC, GCC High, and DoD
- Office 365 Germany

Timeline

- July 1, 2020: The **Search-Mailbox** cmdlet will no longer be available and Microsoft Support will no longer provide assistance.

Alternative tools

The following table describes other tools that you can use to replace the existing functionality that's being retired.

FUNCTIONALITY	ALTERNATIVE TOOLS	COMMENTS

FUNCTIONALITY	ALTERNATIVE TOOLS	COMMENTS
Search a mailbox	*-ComplianceSearch *-ComplianceSearchAction	The ComplianceSearch and ComplianceSearchAction cmdlets work together to help you search and export content. You can create a new search and view the search estimate by using the **New-**, **Get-**, and **Start-ComplianceSearch** cmdlets. Then you can use the **New-ComplianceSearchAction -Export** command to export the search results. You'll still have to use the core eDiscovery tool in the Microsoft 365 compliance center to download those search results to your local computer.
Delete bulk email from a mailbox	Set up an archive and deletion policy for mailboxes	Admins can create an archiving and deletion policy that automatically moves items to a user's archive mailbox and automatically deletes items from the mailbox.
Copy search results to a discovery mailbox		There's no direct replacement for this functionality because it does not provide access to all Microsoft 365 services. See the FAQs in the ***-MailboxSearch cmdlets** section for alternative solutions.
Copy messages from one mailbox to a different mailbox	Assign permissions to a mailbox	To give a person access to another user's email (such as when an employee leaves your organization and you need to give another person access to the former employee's email), we recommended that you assign that person permissions to access the former employee's mailbox. So instead of copying mailbox items to another user mailbox or a shared mailbox, just assign a user permissions to access the source mailbox.
Purge messages from a mailbox	*-ComplianceSearch *-ComplianceSearchAction	The ComplianceSearch and ComplianceSearchAction cmdlets work together to help you search and purge content. You can create and run a search with **New-ComplianceSearch** and **New-ComplianceSearch** cmdlets, and then you can purge the content by using **New-ComplianceSearchAction -Purge -PurgeType** command. For more information, see Search for and delete messages.

FUNCTIONALITY	ALTERNATIVE TOOLS	COMMENTS
Purge messages from a mailbox	Assign permissions to a mailbox	To purge messages from a mailbox, assign an administrator permissions to access the employee's mailbox. Messages can be deleted and recycled as needed taking advantage of the built in search and viewing capabilities in Outlook.

Exchange Web Services API operations

These operations in the Exchange Web Services API are used by the In-Place eDiscovery & Holds feature in the Exchange admin center and the corresponding ***-MailboxSearch** cmdlets in Exchange Online PowerShell. They will also be retired as part of retiring the other legacy eDiscovery tools.

Scope of affected organizations

- Office 365 and Microsoft 365 Enterprise organizations
- Office 365 and Microsoft 365 Education organizations
- Office 365 and Microsoft 365 Government organizations; this includes GCC, GCC High, and DoD
- Office 365 Germany

Timeline

- July 1, 2020: The GetSearchableMailboxes, SearchMailboxes, SetHoldOnMailboxes, and GetHoldOnMailboxes operations will no longer be available, and Microsoft Support will no longer provide assistance.

Advanced eDiscovery v1.0

Advanced eDiscovery v1.0, which is the version of Advanced eDiscovery available in a core eDiscovery case by clicking **Switch to Advanced eDiscovery**, is being retired. Its functionality has been replaced by the new Advanced eDiscovery solution in the Microsoft 365 compliance center.

To determine if your organization is using Advanced eDiscovery v1.0:

1. Go to the Office 365 Security & Compliance Center.
2. In the left navigation pane of the Security & Compliance Center, click **eDiscovery > eDiscovery**, and open a Core eDiscovery case.
3. If you see the **Switch to Advanced eDiscovery** button, then clicking it will take you to the 1.0 version of Advanced eDiscovery, which is being retired. The ability to create and manage cases in Core eDiscovery won't be affected. Only the ability to add and analyze case data in Advanced eDiscovery v1.0 (by clicking **Switch to Advanced eDiscovery**) is being retired.

The new Advanced eDiscovery solution in Microsoft 365 (also known as *Advanced eDiscovery v2.0*) provides all of the capabilities of the original solution, but now includes a custodian-based approach of identifying content in other Microsoft 365 services, collecting that content, and then adding it to a review set where reviewers can take advantage of fast search queries, tagging, and analytics features to help cull relevant documents. Advanced eDiscovery now includes improved processing and native viewers for both Microsoft and non-Microsoft file types, a full list of file types is here and supported metadata fields are here. Also, the new Advanced eDiscovery solution provides a powerful custodian holds management feature that lets you apply holds to content in different services, notify users of the holds, and track custodian responses, all within an Advanced eDiscovery case.

To access Advanced eDiscovery v2.0:

1. Go to the Microsoft 365 compliance center.
2. In the left navigation pane of the Microsoft 365 compliance center, click **Show all**, and then click **eDiscovery > Advanced**.

At this time, we recommend that you begin to transition your eDiscovery workflow to the new Advanced eDiscovery functionality. If required, you can archive your Advanced eDiscovery 1.0 cases by exporting the content and storing it offline. Although you'll still be able to access Advanced eDiscovery v1.0 in existing cases until December 31, 2020, Microsoft Support won't provide support after October 1, 2020. See the following timeline for more details.

Scope of affected organizations

- Office 365 and Microsoft 365 Enterprise organizations
- Office 365 and Microsoft 365 Education organizations
- Office 365 and Microsoft 365 Government organizations; this includes GCC, GCC High, and DoD
- Office 365 Germany

Timeline

- July 1, 2020: You won't be able to create new Advanced eDiscovery v1.0 cases.
- October 1, 2020: You won't be able to add new data (Prepare search results for Advanced eDiscovery) to any cases. You'll be able to continue working with data in existing cases at your own risk. Microsoft Support will no longer provide assistance.
- December 31, 2020: You won't be able to access Advanced eDiscovery v1.0 cases.

Alternative tools

The Advanced eDiscovery solution in the Microsoft 365 compliance center.

How to identify the type of hold placed on an Exchange Online mailbox

11/2/2020 • 14 minutes to read • Edit Online

This article explains how to identify holds placed on Exchange Online mailboxes in Microsoft 365.

Microsoft 365 offers several ways that your organization can prevent mailbox content from being permanently deleted. This allows your organization to retain content to meet compliance regulations or during legal and other types of investigations. Here's a list of the retention features (also called *holds*) in Office 365:

- **Litigation Hold**: Holds that are applied to user mailboxes in Exchange Online.
- **eDiscovery hold**: Holds that are associated with a Core eDiscovery case in the security and compliance center. eDiscovery holds can be applied to user mailboxes and to the corresponding mailbox for Microsoft 365 Groups and Microsoft Teams.
- **In-Place Hold**: Holds that are applied to user mailboxes by using the In-Place eDiscovery & Hold tool in the Exchange admin center in Exchange Online.
- **Microsoft 365 retention policies**: Can be configured to retain (or retain and then delete) content in user mailboxes in Exchange Online and in the corresponding mailbox for Microsoft 365 Groups and Microsoft Teams. You can also create a retention policy to retain Skype for Business Conversations, which are stored in user mailboxes.

  There are two types of Microsoft 365 retention policies that can be assigned to mailboxes.

  - **Specific location retention policies**: These are policies that are assigned to the content locations of specific users. You use the **Get-Mailbox** cmdlet in Exchange Online PowerShell to get information about retention policies assigned to specific mailboxes. For more information about this type of retention policy, see the section A policy with specific inclusions or exclusions from the retention policy documentation.
  - **Organization-wide retention policies**: These are policies that are assigned to all content locations in your organization. You use the **Get-OrganizationConfig** cmdlet in Exchange Online PowerShell to get information about organization-wide retention policies. For more information about this type of retention policy, see the section A policy that applies to entire locations from the retention policy documentation.
- **Microsoft 365 retention labels**: If a user applies a Microsoft 365 retention label (one that's configured to retain content or retain and then delete content) to *any* folder or item in their mailbox, a hold is placed on the mailbox as if the mailbox was placed on Litigation Hold or assigned to a Microsoft 365 retention policy. For more information, see the Identifying mailboxes on hold because a retention label has been applied to a folder or item section in this article.

To manage mailboxes on hold, you may have to identify the type of hold that's placed on a mailbox so that you can perform tasks such as changing the hold duration, temporarily or permanently removing the hold, or excluding a mailbox from a Microsoft 365 retention policy. In these cases, the first step is to identify the type of hold placed on the mailbox. And because multiple holds (and different types of holds) can be placed on a single mailbox, you have to identify all holds placed on a mailbox if you want to remove or change a hold.

Step 1: Obtain the GUID for holds placed on a mailbox

You can run the following two cmdlets in Exchange Online PowerShell to get the GUID of the holds that are placed on a mailbox. After you obtain a GUID, you use it to identify the specific hold in Step 2. A Litigation Hold isn't identified by a GUID. Litigation Holds are either enabled or disabled for a mailbox.

- **Get-Mailbox:** Use this cmdlet to determine whether Litigation Hold is enabled for a mailbox and to get the GUIDs for eDiscovery holds, In-Place Holds, and Microsoft 365 retention policies that are specifically assigned to a mailbox. The output of this cmdlet will also indicate if a mailbox has been explicitly excluded from an organization-wide retention policy.
- **Get-OrganizationConfig:** Use this cmdlet to get the GUIDs for organization-wide retention policies.

To connect to Exchange Online PowerShell, see Connect to Exchange Online PowerShell.

Get-Mailbox

Run the following command to get information about the holds and Microsoft 365 retention policies applied to a mailbox.

```
Get-Mailbox <username> | FL LitigationHoldEnabled,InPlaceHolds
```

TIP

If there are too many values in the InPlaceHolds property and not all of them are displayed, you can run the `Get-Mailbox <username> | Select-Object -ExpandProperty InPlaceHolds` command to display each GUID on a separate line.

The following table describes how to identify different types of holds based on the values in the *InPlaceHolds* property when you run the **Get-Mailbox** cmdlet.

HOLD TYPE	EXAMPLE VALUE	HOW TO IDENTIFY THE HOLD
Litigation Hold	`True`	Litigation Hold is enabled for a mailbox when the *LitigationHoldEnabled* property is set to `True`.
eDiscovery hold	`UniH7d895d48-7e23-4a8d-8346-533c3beac15d`	The *InPlaceHolds property* contains the GUID of any hold associated with an eDiscovery case in the security and compliance center. You can tell this is an eDiscovery hold because the GUID starts with the `UniH` prefix (which denotes a Unified Hold).
In-Place Hold	`c0ba3ce811b6432a8751430937152491` or `cld9c0a984ca74b457fbe4504bf7d3e00de`	The *InPlaceHolds* property contains the GUID of the In-Place Hold that's placed on the mailbox. You can tell this is an In-Place Hold because the GUID either doesn't start with a prefix or it starts with the `cld` prefix.

HOLD TYPE	EXAMPLE VALUE	HOW TO IDENTIFY THE HOLD
Microsoft 365 retention policy specifically applied to the mailbox	`mbxcdbbb86ce60342489bff371876e7f224:1` or `skp127d7cf1076947929bf136b7a2a8c36f:3`	The InPlaceHolds property contains GUIDs of any specific location retention policy that's applied to the mailbox. You can identify retention policies because the GUID starts with the `mbx` or the `skp` prefix. The `skp` prefix indicates that the retention policy is applied to Skype for Business conversations in the user's mailbox.
Excluded from an organization-wide Microsoft 365 retention policy	`-mbxe9b52bf7ab3b46a286308ecb29624696`	If a mailbox is excluded from an organization-wide Microsoft 365 retention policy, the GUID for the retention policy that the mailbox is excluded from is displayed in the InPlaceHolds property and is identified by the `-mbx` prefix.

Get-OrganizationConfig

If the *InPlaceHolds* property is empty when you run the **Get-Mailbox** cmdlet, there still may be one or more organization-wide Microsoft 365 retention policies applied to the mailbox. Run the following command in Exchange Online PowerShell to get a list of GUIDs for organization-wide Microsoft 365 retention policies.

```
Get-OrganizationConfig | FL InPlaceHolds
```

> **TIP**
> If there are too many values in the InPlaceHolds property and not all of them are displayed, you can run the `Get-OrganizationConfig | Select-Object -ExpandProperty InPlaceHolds` command to display each GUID on a separate line.

The following table describes the different types of organization-wide holds and how to identify each type based on the GUIDs contained in *InPlaceHolds* property when you run the **Get-OrganizationConfig** cmdlet.

HOLD TYPE	EXAMPLE VALUE	DESCRIPTION
Microsoft 365 retention policies applied to Exchange mailboxes, Exchange public folders, and Teams chats	`mbx7cfb30345d454ac0a989ab3041051209:2`	Organization-wide retention policies applied to Exchange mailboxes, Exchange public folders, and 1xN chats in Microsoft Teams are identified by GUIDs that start with the `mbx` prefix. Note 1xN chats are stored in the mailbox of the individual chat participants.
Microsoft 365 retention policy applied to Microsoft 365 Groups and Teams channel messages	`grp1a0a132ee8944501a4bb6a452ec31171:3`	Organization-wide retention policies applied to Microsoft 365 groups and channel messages in Microsoft Teams are identified by GUIDs that start with the `grp` prefix. Note channel messages are stored in the group mailbox that is associated with a Microsoft Team.

For more information about retention policies applied to Microsoft Teams, see Learn about retention policies for Microsoft Teams.

Understanding the format of the InPlaceHolds value for retention policies

In addition to the prefix (mbx, skp, or grp) that identifies an item in the InPlaceHolds property as a Microsoft 365 retention policy, the value also contains a suffix that identifies the type of retention action that's configured for the policy. For example, the action suffix is highlighted in bold type in the following examples:

`skp127d7cf1076947929bf136b7a2a8c36f`**:1**

`mbx7cfb30345d454ac0a989ab3041051209`**:2**

`grp1a0a132ee8944501a4bb6a452ec31171`**:3**

The following table defines the three possible retention actions:

VALUE	DESCRIPTION
1	Indicates that the retention policy is configured to delete items. The policy doesn't retain items.
2	Indicates that the retention policy is configured to hold items. The policy doesn't delete items after the retention period expires.
3	Indicates that the retention policy is configured to hold items and then delete them after the retention period expires.

For more information about retention actions, see the Retaining content for a specific period of time section.

Step 2: Use the GUID to identify the hold

After you obtain the GUID for a hold that is applied to a mailbox, the next step is to use that GUID to identify the hold. The following sections show how to identify the name of the hold (and other information) by using the hold GUID.

eDiscovery holds

Run the following commands in Security & Compliance Center PowerShell to identify an eDiscovery hold that's applied to the mailbox. Use the GUID (not including the UniH prefix) for the eDiscovery hold that you identified in Step 1. The first command creates a variable that contains information about the hold. This variable is used in the other commands. The second command displays the name of the eDiscovery case the hold is associated with. The third command displays the name of the hold and a list of the mailboxes the hold applies to.

```
$CaseHold = Get-CaseHoldPolicy <hold GUID without prefix>
```

```
Get-ComplianceCase $CaseHold.CaseId | FL Name
```

```
$CaseHold | FL Name,ExchangeLocation
```

To connect to Security & Compliance Center PowerShell, see Connect to Security & Compliance Center PowerShell.

In-Place Holds

Run the following command in Exchange Online PowerShell to identify the In-Place Hold that's applied to the mailbox. Use the GUID for the In-Place Hold that you identified in Step 1. The command displays the name of the hold and a list of the mailboxes the hold applies to.

```
Get-MailboxSearch -InPlaceHoldIdentity <hold GUID> | FL Name,SourceMailboxes
```

If the GUID for the In-Place Hold starts with the `cld` prefix, be sure to include the prefix when running the previous command.

> **IMPORTANT**
> As we continue to invest in different ways to preserve mailbox content, we're announcing the retirement of In-Place Holds in the Exchange admin center (EAC). Starting July 1, 2020 you won't be able to create new In-Place Holds in Exchange Online. But you'll still be able to manage In-Place Holds in the EAC or by using the **Set-MailboxSearch** cmdlet in Exchange Online PowerShell. However, starting October 1, 2020, you won't be able to manage In-Place Holds. You'll only be remove them in the EAC or by using the **Remove-MailboxSearch** cmdlet. For more information about the retirement of In-Place Holds, see Retirement of legacy eDiscovery tools.

Microsoft 365 retention policies

Run the following command in Security & Compliance Center PowerShell to identity the Microsoft 365 retention policy (organization-wide or specific location) that's applied to the mailbox. Use the GUID (not including the mbx, skp, or grp prefix or the action suffix) that you identified in Step 1.

```
Get-RetentionCompliancePolicy <hold GUID without prefix or suffix> -DistributionDetail  | FL Name,*Location
```

Identifying mailboxes on hold because a retention label has been applied to a folder or item

Whenever a user applies a retention label that's configured to retain content or retain and then delete content to any folder or item in their mailbox, the *ComplianceTagHoldApplied* mailbox property is set to **True**. When this happens, the mailbox is considered to be on hold, as if it was placed on Litigation Hold or assigned to a Microsoft 365 retention policy. When the *ComplianceTagHoldApplied* property is set to **True**, the following things may occur:

- If the mailbox or the user's user account is deleted, the mailbox becomes an inactive mailbox.
- You aren't able to disable the mailbox (either the primary mailbox or the archive mailbox, if it's enabled).
- Items in the mailbox may be retained longer than expected. This is because the mailbox is on hold and therefore no items are permanently deleted (purged).

To view the value of the *ComplianceTagHoldApplied* property, run the following command in Exchange Online PowerShell:

```
Get-Mailbox <username> |FL ComplianceTagHoldApplied
```

For more information about retention labels, see retention labels.

Managing mailboxes on delay hold

After any type of hold is removed from a mailbox, a *delay hold* is applied. This means that the actual removal of the hold is delayed for 30 days to prevent data from being permanently deleted (purged) from the mailbox. This gives admins an opportunity to search for or recover mailbox items that will be purged after a hold is removed.

A delay hold is placed on a mailbox the next time the Managed Folder Assistant processes the mailbox and detects that a hold was removed. Specifically, a delay hold is applied to a mailbox when the Managed Folder Assistant sets one of the following mailbox properties to **True**:

- **DelayHoldApplied:** This property applies to email-related content (generated by people using Outlook and Outlook on the web) that's stored in a user's mailbox.
- **DelayReleaseHoldApplied:** This property applies to cloud-based content (generated by non-Outlook apps such as Microsoft Teams, Microsoft Forms, and Microsoft Yammer) that's stored in a user's mailbox. Cloud data generated by a Microsoft app is typically stored in a hidden folder in a user's mailbox.

When a delay hold is placed on the mailbox (when either of the previous properties is set to **True**), the mailbox is still considered to be on hold for an unlimited hold duration, as if the mailbox was on Litigation Hold. After 30 days, the delay hold expires, and Microsoft 365 will automatically attempt to remove the delay hold (by setting the DelayHoldApplied or DelayReleaseHoldApplied property to **False**) so that the hold is removed. After either of these properties are set to **False**, the corresponding items that are marked for removal are purged the next time the mailbox is processed by the Managed Folder Assistant.

To view the values for the DelayHoldApplied and DelayReleaseHoldApplied properties for a mailbox, run the following command in Exchange Online PowerShell.

```
Get-Mailbox <username> | FL *HoldApplied*
```

To remove the delay hold before it expires, you can run one (or both) the following commands in Exchange Online PowerShell, depending on which property you want to change:

```
Set-Mailbox <username> -RemoveDelayHoldApplied
```

Or

```
Set-Mailbox <username> -RemoveDelayReleaseHoldApplied
```

You must be assigned the Legal Hold role in Exchange Online to use the *RemoveDelayHoldApplied* or *RemoveDelayReleaseHoldApplied* parameters.

To remove the delay hold on an inactive mailbox, run one of the following commands in Exchange Online PowerShell:

```
Set-Mailbox <DN or Exchange GUID> -InactiveMailbox -RemoveDelayHoldApplied
```

Or

```
Set-Mailbox <DN or Exchange GUID> -InactiveMailbox -RemoveDelayReleaseHoldApplied
```

> **TIP**
> The best way to specify an inactive mailbox in the previous command is to use its Distinguished Name or Exchange GUID value. Using one of these values helps prevent accidentally specifying the wrong mailbox.

For more information about using these parameters for managing delay holds, see Set-Mailbox.

Keep the following things in mind when managing a mailbox on delay hold:

- If either the DelayHoldApplied or DelayReleaseHoldApplied property is set to **True** and a mailbox (or the corresponding user account) is deleted, the mailbox becomes an inactive mailbox. That's because a mailbox is considered to be on hold if either property is set to **True**, and deleting a mailbox on hold results in an inactive mailbox. To delete a mailbox and not make it an inactive mailbox, you have to set both properties to **False**.
- As previous stated, a mailbox is considered to be on hold for an unlimited hold duration if either the DelayHoldApplied or DelayReleaseHoldApplied property is set to **True**. However, that doesn't mean that *all* content in the mailbox is preserved. It depends on the value that's set to each property. For example, let's say both properties are set to **True** because holds are removed from the mailbox. Then you remove only the delay hold that's applied to non-Outlook cloud data (by using the *RemoveDelayReleaseHoldApplied* parameter). The next time the Managed Folder Assistant processes the mailbox, the non-Outlook items marked for removal are purged. Any Outlook items marked for removal won't be purged because the DelayHoldApplied property is still set to **True**. The opposite would also be true: if DelayHoldApplied is set to **False** and DelayReleaseHoldApplied is set to **True**, then only Outlook items marked for removal would be purged.

Next steps

After you identify the holds that are applied to a mailbox, you can perform tasks such as changing the duration of the hold, temporarily or permanently removing the hold, or excluding an inactive mailbox from a Microsoft 365 retention policy. For more information about performing tasks related to holds, see one of the following topics:

- Run the Set-RetentionCompliancePolicy -Identity <Policy Name> -AddExchangeLocationException <user mailbox> command in Security & Compliance Center PowerShell to exclude a mailbox from an organization-wide Microsoft 365 retention policy. This command can only be used for retention policies where the value for the *ExchangeLocation* property equals `All`.
- Run the Set-Mailbox -ExcludeFromOrgHolds <hold GUID without prefix or suffix> command in Exchange Online PowerShell to exclude an inactive mailbox from an organization-wide Microsoft 365 retention policy.
- Change the hold duration for an inactive mailbox
- Delete an inactive mailbox
- Delete items in the Recoverable Items folder of cloud-based mailboxes on hold

Create a Litigation Hold

11/2/2020 • 4 minutes to read • Edit Online

You can place a mailbox on Litigation Hold to retain all mailbox content, including deleted items and the original versions of modified items. When you place a user mailbox on Litigation Hold, content in the user's archive mailbox (if it's enabled) is also retained. When you create a hold, you can specify a hold duration (also called a *time-based hold*) so that deleted and modified items are retained for a specified period and then permanently deleted from the mailbox. Or you can just retain content indefinitely (called an *infinite hold*) or until the Litigation Hold is removed. If you do specify a hold duration period, it's calculated from the date a message is received or a mailbox item is created.

Here's what happens when you create a Litigation Hold.

- Items that are permanently deleted by the user are retained in the Recoverable Items folder in the user's mailbox for the duration of the hold.
- Items that are purged from the Recoverable Items folder by the user are retained for the duration of the hold.
- The storage quota for the Recoverable Items folder is increased from 30 GB to 110 GB.
- Items in the user's primary and the archive mailboxes are retained

Assign an Exchange Online Plan 2 license

- To place an Exchange Online mailbox on Litigation Hold, it must be assigned an Exchange Online Plan 2 license. If a mailbox is assigned an Exchange Online Plan 1 license, you would have to assign it a separate Exchange Online Archiving license to place it on hold.

Place a mailbox on Litigation Hold

Here are the steps to place a mailbox on Litigation Hold using the Exchange admin center.

1. Go to https://outlook.office.com/ecp and sign in using your global administrator account.
2. Click **Recipients > Mailboxes** in the left navigation pane.
3. Select the mailbox that you want to place on Litigation Hold, and then click **Edit**.
4. On the mailbox properties page, click **Mailbox features**.
5. Under **Litigation hold: Disabled**, click **Enable** to place the mailbox on Litigation Hold.
6. On the **Litigation hold** page, enter the following optional information:
    - **Litigation hold duration (days)** - Use this box to create a time-based hold and specify how long mailbox items are held when the mailbox is placed on Litigation Hold. The duration is calculated from the date a mailbox item is received or created. When the hold duration expires for a specific item, that item will no longer be preserved. If you leave this box blank, items are preserved indefinitely or until the hold is removed. Use days to specify the duration.
    - **Note** - Use this box to inform the user their mailbox is on Litigation Hold. The note will appear on the Account Information page in the user's mailbox if they're using Outlook 2010 or later. To access this page, users can click **File** in Outlook.

- **URL** - Use this box to direct the user to a website for more information about Litigation Hold. This URL appears on the Account Information page in the user's mailbox if they are using Outlook 2010 or later. To access this page, users can click **File** in Outlook..

7. Click **Save** on the **Litigation hold** page, and then click **Save** on the mailbox properties page.

Create a Litigation Hold using PowerShell

You can also create a Litigation Hold by running the following command in Exchange Online PowerShell:

```
Set-Mailbox <username> -LitigationHoldEnabled $true
```

The previous command preserves items indefinitely because the hold duration isn't specified. To created a time-based hold, using the following command:

```
Set-Mailbox <username> -LitigationHoldEnabled $true -LitigationHoldDuration <number of days>
```

For more information, see Set-Mailbox.

How does Litigation Hold work?

In the normal deleted item workflow, a mailbox item is moved to the Deletions subfolder in the Recoverable Items folder when a user permanently deletes it (Shift + Delete) or deletes it from the Deleted Items folder. A deletion policy (which is a retention tag configured with a Delete retention action) also moves items to the Deletions subfolder when the retention period expires. When a user purges an item in the Recoverable Items folder or when the deleted item retention period expires for an item, it's moved to the Purges subfolder in the Recoverable Items folder and marked for permanent deletion. It will be purged from Exchange the next time the mailbox is processed by the Managed Folder Assistant (MFA).

When a mailbox is placed on Litigation Hold, items in the Purges subfolder are preserved for the hold duration specified by the Litigation Hold. The hold duration is calculated from the original date an item was received or created, and defines how long items in the Purges subfolder are held. When the hold duration expires for an item in the Purges subfolder, the item is marked for permanent deletion and will be purged from Exchange the next time the mailbox is processed by the MFA. If an indefinite hold is placed on a mailbox, items will never be purged from the Purges subfolder.

The following illustration shows the subfolders in the Recoverable Items folders and the hold workflow process.

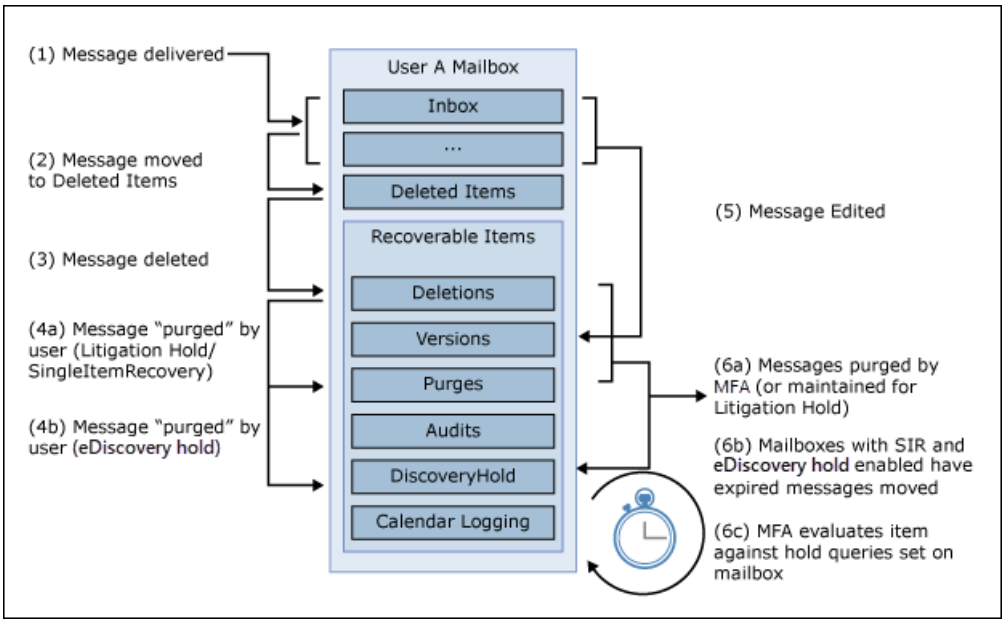

NOTE

If a hold associated with an eDiscovery case is placed on a mailbox, purged items are moved from the Deletions subfolder to the DiscoveryHolds subfolder and are preserved until the mailbox is released from the eDiscovery hold.

Delete items in the Recoverable Items folder of cloud-based mailboxes on hold

2/18/2021 • 23 minutes to read • Edit Online

The Recoverable Items folder for an Exchange Online mailbox exists to protect from accidental or malicious deletions. It's also used to store items that are retained and accessed by compliance features, such as holds and eDiscovery searches. However, in some situations organizations might have data that's been unintentionally retained in the Recoverable Items folder that they must delete. For example, a user might unknowingly send or forward an email message that contains sensitive information or information that may have serious business consequences. Even if the message is permanently deleted, it might be retained indefinitely because a legal hold has been placed on the mailbox. This scenario is known as *data spillage* because data has been unintentionally *spilled* into Office 365. In these situations, you can delete items in a user's Recoverable Items folder for an Exchange Online mailbox, even if that mailbox is placed on hold with one of the different hold features in Office 365. These types of holds include Litigation Holds, In-Place Holds, eDiscovery holds, and retention policies created in the security and compliance center in Office 365 or Microsoft 365.

This article explains how admins can delete items from the Recoverable Items folder for cloud-based mailboxes that are on hold. This procedure involves disabling access to the mailbox and disabling single item recovery, disabling the Managed Folder Assistant from processing the mailbox, temporarily removing the hold, deleting items from the Recoverable Items folder, and then reverting the mailbox to its previous configuration. Here's the process:

Caution

The procedures outlined in this article will result in data being permanently deleted (purged) from an Exchange Online mailbox. That means messages that you delete from the Recoverable Items folder can't be recovered and won't be available for legal discovery or other compliance purposes. If you want to delete messages from a mailbox that's placed on hold as part of a Litigation Hold, In-Place Hold, eDiscovery hold, or retention policy created in the security and compliance center, check with your records management or legal departments before removing the hold. Your organization might have a policy that defines whether a mailbox on hold or a data spillage incident takes priority.

Before you delete items

- To create and run a Content Search, you have to be a member of the eDiscovery Manager role group or be assigned the Compliance Search management role. To delete messages, you have to be a member of the Organization Management role group or be assigned the Search And Purge management role. For information about adding users to a role group, see Assign eDiscovery permissions in the Security & Compliance Center.
- The procedure described in this article isn't supported for inactive mailboxes. That's because you can't

reapply a hold (or retention policy) to an inactive mailbox after you remove it. When you remove a hold from an inactive mailbox, it's changed to a normal soft-deleted mailbox and will be permanently deleted from your organization after it's processed by the Managed Folder Assistant.

- You can't perform this procedure for a mailbox that has been assigned retention settings with a policy that's locked by using Preservation Lock. That's because this lock prevents you from removing or excluding the mailbox from the policy and from disabling the Managed Folder Assistant on the mailbox. For more information about locking policies for retention,see Use Preservation Lock to restrict changes to retention policies and retention label policies.
- If a mailbox isn't placed on hold (or doesn't have single item recovery enabled), you can delete the items from the Recoverable Items folder. For more information about how to do this, see Search for and delete email messages in your organization.

Step 1: Collect information about the mailbox

This first step is to collect selected properties from the target mailbox that will affect this procedure. Be sure to write down these settings or save them to a text file because you'll change some of these properties and then revert back to the original values in Step 6, after you delete items from the Recoverable Items folder. Here's a list of the mailbox properties you need to collect.

- *SingleItemRecoveryEnabled* and *RetainDeletedItemsFor*. If necessary, you'll disable single recovery and increase the deleted items retention period in Step 3.
- *LitigationHoldEnabled* and *InPlaceHolds*. You need to identify all the holds placed on the mailbox so that you can temporarily remove them in Step 3. See the More information section for tips about how to identify the type hold that might be placed on a mailbox.

Additionally, you need to get the mailbox client access settings so you can temporarily disable them so the owner (or other users) can't access the mailbox during this procedure. Finally, you can get the current size and number of items in the Recoverable Items folder. After you delete items in the Recoverable Items folder in Step 5, you'll use this information to verify that items were removed.

1. Connect to Exchange Online PowerShell. Be sure to use a user name and password for an administrator account that's been assigned the appropriate management roles in Exchange Online.
2. Run the following command to get information about single item recovery and the deleted item retention period.

   ```
   Get-Mailbox <username> | FL SingleItemRecoveryEnabled,RetainDeletedItemsFor
   ```

   If single item recovery is enabled, you'll have to disable it in Step 2. If the deleted item retention period isn't set for 30 days (the maximum value in Exchange Online), then you can increase it in Step 2.

3. Run the following command to get the mailbox access settings for the mailbox.

   ```
   Get-CASMailbox <username> | FL
   EwsEnabled,ActiveSyncEnabled,MAPIEnabled,OWAEnabled,ImapEnabled,PopEnabled
   ```

   You'll disable all of these access methods in Step 2.

4. Run the following command to get information about the holds and retention policies applied to the mailbox.

   ```
   Get-Mailbox <username> | FL LitigationHoldEnabled,InPlaceHolds
   ```

TIP

If there are too many values in the *InPlaceHolds* property and not all of them are displayed, you can run the `Get-Mailbox <username> | Select-Object -ExpandProperty InPlaceHolds` command to display each value on a separate line.

5. Run the following command to get information about any organization-wide retention policies.

```
Get-OrganizationConfig | FL InPlaceHolds
```

If your organization has any organization-wide retention policies, you'll have to exclude the mailbox from these policies in Step 3.

TIP

If there are too many values in the *InPlaceHolds* property and not all of them are displayed, you can run the `Get-OrganizationConfig | Select-Object -ExpandProperty InPlaceHolds` command to display each value on a separate line.

6. Run the following command to get the current size and total number of items in folders and subfolders in the Recoverable Items folder in the user's primary mailbox.

```
Get-MailboxFolderStatistics <username> -FolderScope RecoverableItems | FL
Name,FolderAndSubfolderSize,ItemsInFolderAndSubfolders
```

If the user's archive mailbox is enabled, run the following command to get the size and total number of items in folders and subfolders in the Recoverable Items folder in their archive mailbox.

```
Get-MailboxFolderStatistics <username> -FolderScope RecoverableItems -Archive | FL
Name,FolderAndSubfolderSize,ItemsInFolderAndSubfolders
```

When you delete items in Step 5, you can choose to delete or not delete items in the Recoverable Items folder in the user's primary archive mailbox. If auto-expanding archiving is enabled for the mailbox, items in an auxiliary archive mailbox won't be deleted.

Step 2: Prepare the mailbox

After collecting and saving information about the mailbox, the next step is to prepare the mailbox by performing the following tasks:

- **Disable client access to mailbox** so that the mailbox owner can't access their mailbox and make any changes to the mailbox data during this procedure.
- **Increase the deleted item retention period** to 30 days (the maximum value in Exchange Online) so that items aren't purged from the Recoverable Items folder before you can delete them in Step 5.
- **Disable single Item recovery** so that items won't be retained (for the duration of the deleted item retention period) after you delete them from the Recoverable Items folder in Step 5.
- **Disable the Managed Folder Assistant** so that it doesn't process the mailbox and retain the items that you delete in Step 5.

Perform the following steps in Exchange Online PowerShell.

1. Run the following command to disable all client access to the mailbox. The command syntax assumes that all client access methods were enabled on the mailbox.

```
Set-CASMailbox <username> -EwsEnabled $false -ActiveSyncEnabled $false -MAPIEnabled $false -
OWAEnabled $false -ImapEnabled $false -PopEnabled $false
```

> **NOTE**
> It might take up to 60 minutes to disable all client access methods to the mailbox. Note that disabling these access methods won't disconnect the mailbox owner they're currently signed in. If the owner isn't signed in, then they won't be able to access their mailbox after these access methods are disabled.

2. Run the following command to increase the deleted item retention period the maximum of 30 days. This assumes that the current setting is less than 30 days.

```
Set-Mailbox <username> -RetainDeletedItemsFor 30
```

3. Run the following command to disable single item recovery.

```
Set-Mailbox <username> -SingleItemRecoveryEnabled $false
```

> **NOTE**
> It might take up to 60 minutes to disable single item recovery. Don't delete items in the Recoverable Items folder until this period has elapsed.

4. Run the following command to prevent the Managed Folder Assistant from processing the mailbox. As previously explained, you can disable the Managed Folder Assistant only if a retention policy with a Preservation Lock is not applied to the mailbox.

```
Set-Mailbox <username> -ElcProcessingDisabled $true
```

Step 3: Remove all holds from the mailbox

The last step before you can delete items from the Recoverable Items folder is to remove all holds (that you identified in Step 1) placed on the mailbox. All holds must be removed so that items won't be retained after you delete them from the Recoverable Items folder. The following sections contain information about removing different types of holds on a mailbox. See the More information section for tips about how to identify the type hold that might be placed on a mailbox. For more information, see How to identify the type of hold placed on an Exchange Online mailbox.

Caution

As previously stated, check with your records management or legal departments before removing a hold from a mailbox.

Litigation Hold

Run the following command in Exchange Online PowerShell to remove a Litigation Hold from the mailbox.

```
Set-Mailbox <username> -LitigationHoldEnabled $false
```

> **NOTE**
> Similar to disabling the client access methods and single item recovery, it might take up to 60 minutes to remove the Litigation Hold. Don't delete items from the Recoverable Items folder until this period has elapsed.

In-Place Hold

Run the following command in Exchange Online PowerShell to identify the In-Place Hold that's placed on the mailbox. Use the GUID for the In-Place Hold that you identified in Step 1.

```
Get-MailboxSearch -InPlaceHoldIdentity <hold GUID> | FL Name
```

After you identify the In-Place Hold, you can use the Exchange admin center (EAC) or Exchange Online PowerShell to remove the mailbox from the hold. For more information, see Create or remove an In-Place Hold.

Retention policies applied to specific mailboxes

Run the following command in Security & Compliance Center PowerShell to identify the retention policy that is applied to the mailbox. This command will also return any Teams conversation retention policies applied to a mailbox. Use the GUID (not including the `mbx` or `skp` prefix) for the retention policy that you identified in Step 1.

```
Get-RetentionCompliancePolicy <retention policy GUID without prefix> | FL Name
```

After you identify the retention policy, go to the **Information governance** > **Retention** page in the Security & Compliance Center, edit the retention policy that you identified in the previous step, and remove the mailbox from the list of recipients that are included in the retention policy.

Organization-wide retention policies

Organization-wide, Exchange-wide, and Teams-wide retention policies are applied to every mailbox in the organization. They are applied at the organization level (not the mailbox level) and are returned when you run the **Get-OrganizationConfig** cmdlet in Step 1. Run the following command in Security & Compliance Center PowerShell to identify the organization-wide retention policies. Use the GUID (not including the `mbx` prefix) for the organization-wide retention policies that you identified in Step 1.

```
Get-RetentionCompliancePolicy <retention policy GUID without prefix> | FL Name
```

After you identify the organization-wide retention policies, go to the **Information governance** > **Retention** page in the Security & Compliance Center, edit each organization-wide retention policy that you identified in the previous step, and add the mailbox to the list of excluded recipients. Doing this will remove the user's mailbox from the retention policy.

Retention labels

Whenever a user applies a label that's configured to retain content or retain and then delete content to any folder or item in their mailbox, the *ComplianceTagHoldApplied* mailbox property is set to **True**. When this happens, the mailbox is considered to be on hold, as if it was placed on Litigation Hold or assigned to a retention policy.

To view the value of the *ComplianceTagHoldApplied* property, run the following command in Exchange Online PowerShell:

```
Get-Mailbox <username> |FL ComplianceTagHoldApplied
```

After you've identified that a mailbox is on hold because a retention label is applied to a folder or item, you can use the Content Search tool in the security and compliance center to search for labeled items by using the ComplianceTag search condition. For more information, see the "Search conditions" section in Keyword queries and search conditions for Content Search.

For more information about labels, see Learn about retention policies and retention labels.

eDiscovery holds

Run the following commands in Security & Compliance Center PowerShell to identify the hold associated with an eDiscovery case (called *eDiscovery holds*) that's applied to the mailbox. Use the GUID (not including the `UniH` prefix) for the eDiscovery hold that you identified in Step 1. The second command displays the name of the eDiscovery case the hold is associated with; the third command displays the name of the hold.

```
$CaseHold = Get-CaseHoldPolicy <hold GUID without prefix>
```

```
Get-ComplianceCase $CaseHold.CaseId | FL Name
```

```
$CaseHold.Name
```

After you've identified the name of the eDiscovery case and the hold, go to the **eDiscovery** > **eDiscovery** page in the compliance center, open the case, and remove the mailbox from the hold. For more information about identifying eDiscovery holds, see the "eDiscovery holds" section in How to identify the type of hold placed on an Exchange Online mailbox.

Step 4: Remove the delay hold from the mailbox

After any type of hold is removed from a mailbox, the value of the *DelayHoldApplied* or *DelayReleaseHoldApplied* mailbox property is set to **True**. This occurs the next time the Managed Folder Assistant processes the mailbox and detects that a hold has been removed. This is called a *delay hold* and means the actual removal of the hold is delayed for 30 days to prevent data from being permanently deleted from the mailbox. (The purpose of a delay hold is to give admins an opportunity to search for or recover mailbox items that will be purged after a hold is removed.) When a delay hold is placed on the mailbox, the mailbox is still considered to be on hold for an unlimited duration, as if the mailbox was on Litigation Hold. After 30 days, the delay hold expires, and Microsoft 365 will automatically attempt to remove the delay hold (by setting the *DelayHoldApplied* or *DelayReleaseHoldApplied* property to **False**) so that the hold is removed. For more information about a delay hold, see the "Managing mailboxes on delay hold" section in How to identify the type of hold placed on an Exchange Online mailbox.

Before you can delete items in Step 5, you have to remove a delay hold from the mailbox. First, determine if the delay hold is applied to the mailbox by running the following command in Exchange Online PowerShell:

```
Get-Mailbox <username> | FL DelayHoldApplied,DelayReleaseHoldApplied
```

If the value of either the *DelayHoldApplied* or *DelayReleaseHoldApplied* property is set to **False**, a delay hold has not been placed on the mailbox. You can go to Step 5 and delete items in the Recoverable Items folder.

If the value of the *DelayHoldApplied* or *DelayReleaseHoldApplied* property is set to **True**, run one of the following commands to remove the delay hold:

```
Set-Mailbox <username> -RemoveDelayHoldApplied
```

Or

```
Set-Mailbox <username> -RemoveDelayReleaseHoldApplied
```

You must be assigned the Legal Hold role in Exchange Online to use the *RemoveDelayHoldApplied* or *RemoveDelayReleaseHoldApplied* parameter.

Step 5: Delete items in the Recoverable Items folder

Now you're ready to actually delete items in the Recoverable Items folder by using the New-ComplianceSearch and New-ComplianceSearchAction cmdlets in Security & Compliance Center PowerShell.

To search for items that are located in the Recoverable Items folder, we recommend that you perform a *targeted collection*. This means you narrow the scope of your search only to items located in the Recoverable Items folder. You can do this by running the script in the Use Content Search for targeted collections article. This script returns the value of the folder ID property for all the subfolders in the target Recoverable Items folder. Then you use the folder ID in a search query to return items located in that folder.

Here's an overview of the process to search for and delete items in a user's Recoverable Items folder:

1. Run the targeted collection script that returns the folder IDs for all folders in the target user's mailbox. The script connects to Exchange Online PowerShell and Security & Compliance PowerShell in the same PowerShell session. For more information, see Run the script to get a list of folders for a mailbox.

2. Copy the folder IDs for all subfolders in the Recoverable Items folder. Alternatively, you can redirect the output of the script to a text file.

   Here's a list and description of the subfolders in the Recoverable Items folder that you can search and delete items from:

   - **Deletions**: Contains soft-deleted items whose deleted item retention period has not expired. Users can recover soft-deleted items from this subfolder using the Recover Deleted Items tool in Outlook.
   - **Purges**: Contains hard-deleted items whose deleted item retention period has expired. Users can also hard-delete items by purging items from their Recoverable Items folder. If the mailbox is on hold, hard-deleted items are preserved. This subfolder isn't visible to end users.
   - **DiscoveryHolds**: Contains hard-deleted items that have been preserved by an eDiscovery hold or a retention policy. This subfolder isn't visible to end users.
   - **SubstrateHolds**: Contains hard-deleted items from Teams and other cloud-based apps that have been preserved by a retention policy or other type of hold. This subfolder isn't visible to end users.

3. Use the **New-ComplianceSearch** cmdlet (in Security & Compliance Center PowerShell) or use the Content search tool in the compliance center to create a content search that returns items from the target user's Recoverable Items folder. You can do this by including the FolderId in the search query for all subfolders that you want to search. For example, the following query returns all messages in the Purges and eDiscoveryHolds subfolders:

   ```
   folderid:<folder ID of Purges subfolder> OR folderid:<folder ID of DiscoveryHolds subfolder>
   ```

   For more information and examples about running content searches that use the folder ID property, see Use a folder ID or to perform a targeted collection.

> **NOTE**
> If you use the **New-ComplianceSearch** cmdlet to search the Recoverable Items folder, be sure to use **Start-ComplianceSearch** cmdlet to run the search.

4. After you've created a content search and validated that it returns the items that you wan to delete, use the `New-ComplianceSearchAction -Purge -PurgeType HardDelete` command (in Security & Compliance Center PowerShell) to permanently delete the items returned by the content search that you created in the previous step. For example, you can run a command similar to the following command:

   ```
   New-ComplianceSearchAction -SearchName "RecoverableItems" -Purge -PurgeType HardDelete
   ```

5. A maximum of 10 items per mailbox are deleted when you run the previous command. That means you may have to run the `New-ComplianceSearchAction -Purge` command multiple times to delete all the items that you want to delete in the Recoverable Items folder. To delete additional items, you first have to remove the previous compliance search purge action. You do this by running the `Remove-ComplianceSearchAction` cmdlet. For example, to delete the purge action that was run in the previous step, run the following command:

   ```
   Remove-ComplianceSearchAction "RecoverableItems_Purge"
   ```

   After you do this, you can create a new compliance search purge action to delete more items. You'll have to delete each purge action before creating a new one.

   To get a list of the compliance search actions, you can run the `Get-ComplianceSearchAction` cmdlet. Purge actions are identified by `_Purge` appended to the search name.

Verify that items were deleted

To verify that you've successfully deleted items from the Recoverable Items folder of a mailbox, use **Get-MailboxFolderStatistics** cmdlet in Exchange Online PowerShell to check the size and number of items in Recoverable Items folder. You can compare these statistics with the ones you collected in Step 1.

Run the following command in to get the current size and total number of items in folders and subfolders in the Recoverable Items folder in the user's primary mailbox.

```
Get-MailboxFolderStatistics <username> -FolderScope RecoverableItems | FL
Name,FolderAndSubfolderSize,ItemsInFolderAndSubfolders
```

Run the following command to get the size and total number of items in folders and subfolders in the Recoverable Items folder in the user's archive mailbox.

```
Get-MailboxFolderStatistics <username> -FolderScope RecoverableItems -Archive | FL
Name,FolderAndSubfolderSize,ItemsInFolderAndSubfolders
```

Step 6: Revert the mailbox to its previous state

The final step is to revert the mailbox back to its previous configuration. This means resetting the properties that you changed in Step 2 and reapplying the holds that you removed in Step 3. This includes:

- Changing the deleted item retention period back to its previous value. Alternatively, you can just leave this set to 30 days, the maximum value in Exchange Online.

- Re-enabling single Item recovery.
- Re-enabling the client access methods so that the owner can access their mailbox.
- Reapplying the holds and retention policies that you removed.
- Re-enabling the Managed Folder Assistant to process the mailbox.

> **IMPORTANT**
> We recommend that you wait 24 hours after re-applying a hold or retention policy (and verifying that it's in place) before you re-enable the Managed Folder Assistant to process the mailbox.

Perform the following steps (in the specified sequence) in Exchange Online PowerShell.

1. Run the following command to change the deleted item retention period back to its original value. This assumes that the previous setting is less than 30 days; for example, 14 days.

   ```
   Set-Mailbox <username> -RetainDeletedItemsFor 14
   ```

2. Run the following command to re-enable single item recovery.

   ```
   Set-Mailbox <username> -SingleItemRecoveryEnabled $true
   ```

3. Run the following command to re-enable all client access methods to the mailbox.

   ```
   Set-CASMailbox <username> -EwsEnabled $true -ActiveSyncEnabled $true -MAPIEnabled $true -OWAEnabled
   $true -ImapEnabled $true -PopEnabled $true
   ```

4. Reapply the holds that you removed in Step 3. Depending on the type of hold, use one of the following procedures.

   ### Litigation Hold

   Run the following command to re-enable a Litigation Hold for the mailbox.

   ```
   Set-Mailbox <username> -LitigationHoldEnabled $true
   ```

   ### In-Place Hold

   Use the EAC (or Exchange Online PowerShell) to add the mailbox back to the In-Place Hold.

   ### Retention policies applied to specific mailboxes

   Use the Security & Compliance Center to add the mailbox back to the retention policy. Go to the **Information governance** > **Retention** page in the Security & Compliance Center, edit the retention policy, and add the mailbox back to the list of recipients that the retention policy is applied to.

   ### Organization-wide Retention policies

   If you removed an organization-wide or Exchange-wide retention policy by excluding it from the policy, then use the Security & Compliance Center to remove the mailbox from the list of excluded users. Go to the **Information governance** > **Retention** page in the Security & Compliance Center, edit the organization-wide retention policy, and remove the mailbox from the list of excluded recipients. Doing this will reapply the retention policy to the user's mailbox.

eDiscovery case holds

Use the Security & Compliance Center to add the mailbox back the hold that's associated with an eDiscovery case. Go to the **eDiscovery** > **eDiscovery** page, open the case, and add the mailbox back to the hold.

5. Run the following command to allow the Managed Folder Assistant to process the mailbox again. As previously stated, we recommend that you wait 24 hours after reapplying a hold or retention policy (and verifying that it's in place) before you re-enable the Managed Folder Assistant.

```
Set-Mailbox <username> -ElcProcessingDisabled $false
```

6. To verify that the mailbox has been reverted back to its previous configuration, you can run the following commands and then compare the settings to the ones that you collected in Step 1.

```
Get-Mailbox <username> | FL
ElcProcessingDisabled,InPlaceHolds,LitigationHoldEnabled,RetainDeletedItemsFor,SingleItemRecoveryEnab
led
```

```
Get-CASMailbox <username> | FL
EwsEnabled,ActiveSyncEnabled,MAPIEnabled,OWAEnabled,ImapEnabled,PopEnabled
```

More information

Here's a table that describes how to identify different types of holds based on the values in the *InPlaceHolds* property when you run the **Get-Mailbox** or **Get-OrganizationConfig** cmdlets. For more detailed information, see How to identify the type of hold placed on an Exchange Online mailbox.

As previously explained, you have to remove all holds and retention policies from a mailbox before you can successfully delete items in the Recoverable Items folder.

HOLD TYPE	EXAMPLE VALUE	HOW TO IDENTIFY THE HOLD
Litigation Hold	`True`	The *LitigationHoldEnabled* property is set to `True`.
In-Place Hold	`c0ba3ce811b6432a8751430937152491`	The *InPlaceHolds* property contains the GUID of the In-Place Hold that's placed on the mailbox. You can tell this is an In-Place Hold because the GUID doesn't start with a prefix. You can use the `Get-MailboxSearch -InPlaceHoldIdentity <hold GUID> | FL` command in Exchange Online PowerShell to get information about the In-Place Hold on the mailbox.

HOLD TYPE	EXAMPLE VALUE	HOW TO IDENTIFY THE HOLD
Retention policies in the Security & Compliance Center applied to specific mailboxes	`mbxcdbbb86ce60342489bff371876e7f224` or `skp127d7cf1076947929bf136b7a2a8c36f`	When you run the **Get-Mailbox** cmdlet, the *InPlaceHolds* property also contains GUIDs of retention policies applied to the mailbox. You can identify retention policies because the GUID starts with the `mbx` prefix. If the GUID of the retention policy starts with the `skp` prefix, that indicates that the retention policy is applied to Skype for Business conversations. To identity the retention policy that's applied to the mailbox, run the following command in Security & Compliance Center PowerShell: `Get-RetentionCompliancePolicy <retention policy GUID without prefix> | FL Name` Be sure to remove the `mbx` or `skp` prefix when you run this command.
Organization-wide retention policies in the Security & Compliance Center	No value or `-mbxe9b52bf7ab3b46a286308ecb29624696` (indicates that the mailbox is excluded from an organization-wide policy)	Even if the *InPlaceHolds* property is empty when you run the **Get-Mailbox** cmdlet, there still might be one or more organization-wide retention policies applied to the mailbox. To verify this, you can run the `Get-OrganizationConfig | FL InPlaceHolds` command in Exchange Online PowerShell to get a list of the GUIDs for organization-wide retention policies. The GUID for organization-wide retention policies applied to Exchange mailboxes starts with the `mbx` prefix; for example, `mbxa3056bb15562480fadb46ce523ff7b02`. To identity the organization-wide retention policy that's applied to the mailbox, run the following command in Security & Compliance Center PowerShell: `Get-RetentionCompliancePolicy <retention policy GUID without prefix> | FL Name` If a mailbox is excluded from an organization-wide retention policy, the GUID for the retention policy is displayed in the *InPlaceHolds* property of the user's mailbox when you run the **Get-Mailbox** cmdlet; it's identified by the prefix `-mbx`; for example, `-mbxe9b52bf7ab3b46a286308ecb29624696`

HOLD TYPE	EXAMPLE VALUE	HOW TO IDENTIFY THE HOLD

eDiscovery case hold in the Security & Compliance Center	`UniH7d895d48-7e23-4a8d-8346-533c3beac15d`	The *InPlaceHolds* property also contains the GUID of any hold associated with an eDiscovery case in the Security & Compliance Center that might be placed on the mailbox. You can tell this is an eDiscovery case hold because the GUID starts with the `UniH` prefix.
		You can use the `Get-CaseHoldPolicy` cmdlet in Security & Compliance Center PowerShell to get information about the eDiscovery case that the hold on the mailbox is associated with. For example, you can run the command `Get-CaseHoldPolicy <hold GUID without prefix> | FL Name` to display the name of the case hold that's on the mailbox. Be sure to remove the `UniH` prefix when you run this command.
		
		To identity the eDiscovery case that the hold on the mailbox is associated with, run the following commands:
		
		`$CaseHold = Get-CaseHoldPolicy <hold GUID without prefix>`
		
		`Get-ComplianceCase $CaseHold.CaseId | FL Name`

Increase the Recoverable Items quota for mailboxes on hold

11/2/2020 • 10 minutes to read • Edit Online

The default Exchange retention policy—named *Default MRM Policy*—that is automatically applied to new mailboxes in Exchange Online contains a retention tag named Recoverable Items 14 days move to archive. This retention tag moves items from the Recoverable Items folder in the user's primary mailbox to the Recoverable Items folder in the user's archive mailbox after the 14-day retention period expires for an item. For this to happen, the user's archive mailbox must be enabled. If the archive mailbox isn't enabled, no action is taken, which means that items in the Recoverable Items folder for a mailbox on hold aren't moved to the archive mailbox after the 14-day retention period expires. Because nothing is deleted from a mailbox on hold, it's possible that the storage quota for the Recoverable Items folder might be exceeded, especially if the user's archive mailbox isn't enabled.

To help reduce the chance of exceeding this limit, the storage quota for the Recoverable Items folder is automatically increased from 30 GB to 100 GB when a hold is placed on a mailbox in Exchange Online. If the archive mailbox is enabled, the storage quota for the Recoverable Items folder in the archive mailbox is also increased from 30 GB to 100 GB. If the auto-expanding archiving feature in Exchange Online is enabled, the storage quota for the Recoverable Items folder in the user's archive will be unlimited.

The following table summarizes the storage quota for the Recoverable Items folder.

LOCATION OF RECOVERABLE ITEMS FOLDER	MAILBOXES NOT ON HOLD	MAILBOXES ON HOLD
Primary mailbox	30 GB	100 GB
Archive mailbox*	Unlimited	Unlimited
Total storage quota for the Recoverable Items folder	Unlimited	Unlimited

NOTE

* The initial storage quota for the archive mailbox is 100 GB for users with an Exchange Online (Plan 2) license. However, when auto-expanding archiving is turned on for mailboxes on hold, the storage quota for both the archive mailbox and the Recoverable Items folder is increased to 110 GB. Additional archive storage space will be provisioned when necessary which results in an unlimited amount of archive storage. For more information about auto-expanding archiving, see Overview of unlimited archiving in Office 365.

When the storage quota for the Recoverable Items folder in the primary mailbox of a mailbox on hold is close to reaching its limit, you can do the following things:

- **Enable the archive mailbox and turn on auto-expanding archiving.** You can enable an unlimited storage capacity for the Recoverable Items folder simply by enabling the archive mailbox and then turning on the auto-expanding archiving feature in Exchange Online. This results in 110 GB for the Recoverable Items folder in the primary mailbox and an unlimited amount of storage capacity for the Recoverable Items folder in the user's archive. See how: Enable archive mailboxes in the Security & Compliance Center and Enable unlimited archiving in Office 365.

> **NOTE**
>
> After you enable the archive for a mailbox that's close to exceeding the storage quota for the Recoverable Items folder, you might want to run the Managed Folder Assistant to manually trigger the assistant to process the mailbox so that expired items are moved to the Recoverable Items folder in the archive mailbox. See Step 4 for instructions. Note that other items in the user's mailbox might be moved to the new archive mailbox. Consider telling the user that this may happen after you enable the archive mailbox.

- **Create a custom Exchange retention policy for mailboxes on hold.** In addition to enabling the archive mailbox and auto-expanding archiving for mailboxes on Litigation Hold or In-Place Hold, you might also want to create a custom Exchange retention policy for mailboxes on hold. This lets you apply a retention policy to mailboxes on hold that's different from the Default MRM Policy that's applied to mailboxes that aren't on hold, and lets you apply retention tags that are designed for mailboxes on hold. This includes creating a new retention tag for the Recoverable Items folder.

The remainder of this topic describes the step-by-step procedures to create a custom Exchange retention policy for mailboxes on hold.

Step 1: Create a custom retention tag for the Recoverable Items folder

Step 2: Create a new Exchange retention policy for mailboxes on hold

Step 3: Apply the new Exchange retention policy to mailboxes on hold

(Optional) Step 4: Run the Managed Folder Assistant to apply the new retention settings

Step 1: Create a custom retention tag for the Recoverable Items folder

The first step is to create a custom retention tag (called a retention policy tag or RPT) for the Recoverable Items folder. As previously explained, this RPT moves items from the Recoverable Items folder in the user's primary mailbox to the Recoverable Items folder in the user's archive mailbox. You have to use PowerShell to create an RPT for the Recoverable Items folder. You can't use the Exchange admin center (EAC).

1. Connect to Exchange Online using remote PowerShell
2. Run the following command to create a new RPT for the Recoverable Items folder:

```
New-RetentionPolicyTag -Name <Name of RPT> -Type RecoverableItems -AgeLimitForRetention <Number of
days> -RetentionAction MoveToArchive
```

   For example, the following command creates an RPT for the Recoverable Items folder named "Recoverable Items 30 days for mailboxes on hold", with a retention period of 30 days. This means that after an item has been in the Recoverable Items folder for 30 days, it will be moved to the Recoverable Items folder in the user's archive mailbox.

```
New-RetentionPolicyTag -Name "Recoverable Items 30 days for mailboxes on hold" -Type RecoverableItems
-AgeLimitForRetention 30 -RetentionAction MoveToArchive
```

TIP

We recommend that the retention period (defined by the *AgeLimitForRetention* parameter) for the Recoverable Items RPT is the same as the deleted item retention period for the mailboxes that the RPT will be applied to. This allows a user the entire deleted item retention period to recover deleted items before they are moved to the archive mailbox. In the previous example, the retention period was set to 30 days based on the assumption that the deleted item retention period for mailboxes is also 30 days. An Exchange Online mailbox is configured to retain deleted items for 14 days, by default. But you can change this setting to a maximum of 30 days. For more information, see Change the deleted item retention period for a mailbox in Exchange Online.

Step 2: Create a new Exchange retention policy for mailboxes on hold

The next step is to create a new retention policy and add retention tags to it, including the Recoverable Items RPT that you created in Step 1. This new policy will be applied to mailboxes on hold in the next step.

Before you create the new retention policy, determine the additional retention tags that you want to add. For a list of the retention tags that are added to the Default MRM Policy and for information about creating new retention tags, see the following:

- Default Retention Policy in Exchange Online
- Default folders that support Retention Policy Tags
- The "Create a retention tag" section in the Create a Retention Policy topic.

You can use the EAC or Exchange Online PowerShell to create a retention policy.

Use the EAC to create a retention policy

1. In the EAC, go to **Compliance management** > **Retention policies**, and then click **Add +**.
2. On the **New retention policy** page, under **Name**, type a name that describes the purpose of the retention policy; for example, **MRM Policy for Mailboxes on Hold**.
3. Under **Retention tags**, click **Add +**.
4. In the list of retention tags, select the Recoverable Items RPT that you created in Step 1, and then click **Add**.

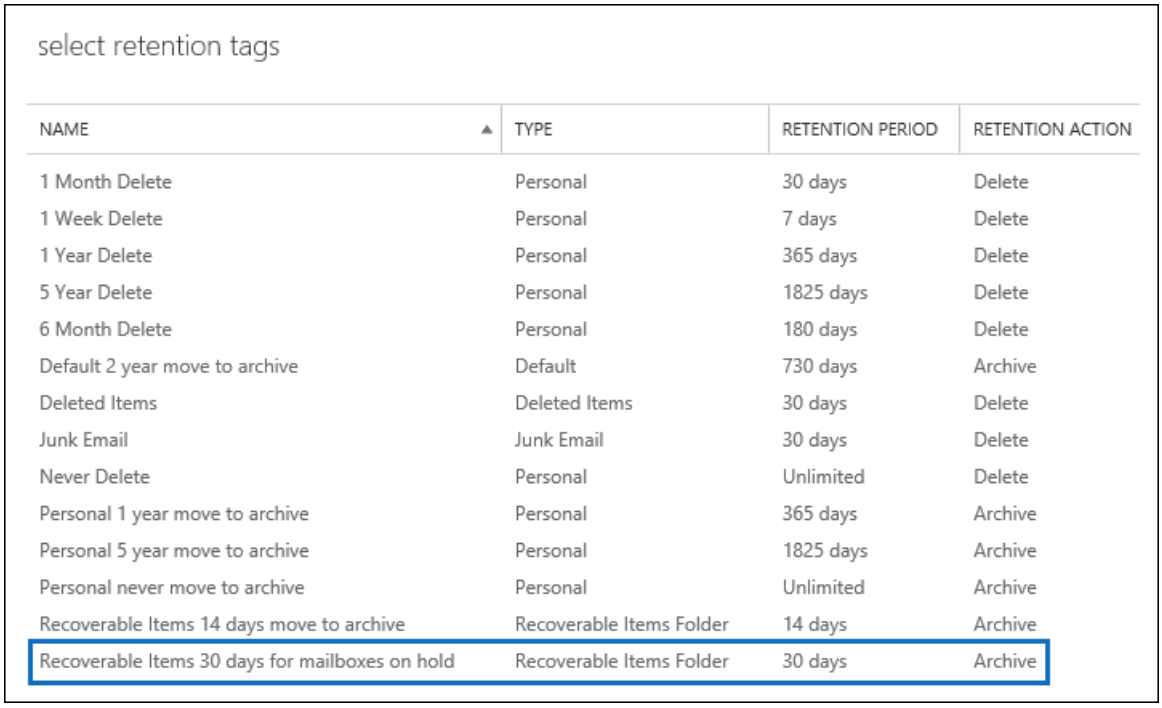

5. Select additional retention tags to add to the retention policy. For example, you might want to add the same tags that are included in the Default MRM Policy.
6. When you're finished adding retention tags, click **OK**.
7. Click **Save** to create the new retention policy.

   Notice that the retention tags linked to the retention policy are displayed in the details pane.

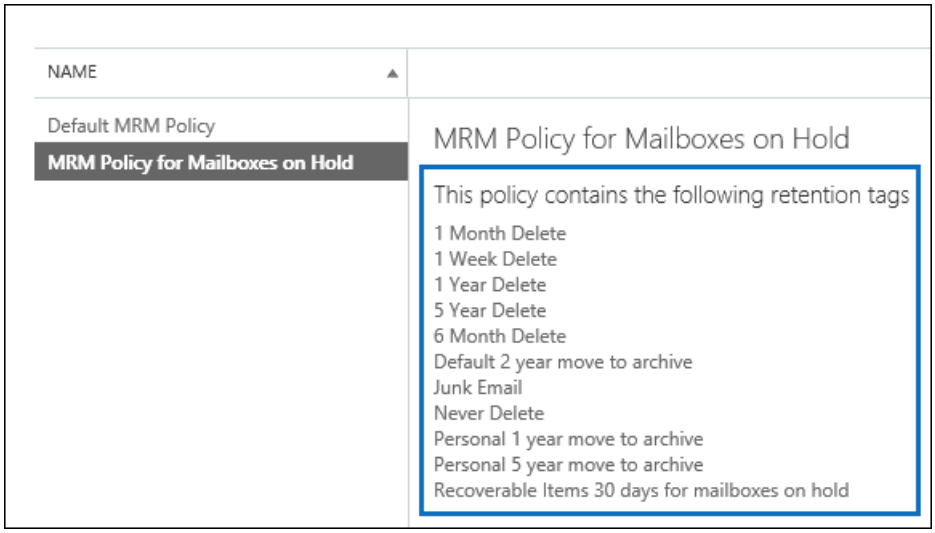

Use Exchange Online PowerShell to create a retention policy

Run the following command to create new retention policy for mailboxes on hold.

```
New-RetentionPolicy <Name of retention policy>  -RetentionPolicyTagLinks <list of retention tags>
```

For example, the following command creates the retention policy and linked retention tags that are displayed in the previous illustration.

```
New-RetentionPolicy "MRM Policy for Mailboxes on Hold"  -RetentionPolicyTagLinks "Recoverable Items 30 days for mailboxes on hold","1 Month Delete","1 Week Delete","1 Year Delete","5 Year Delete","6 Month Delete","Default 2 year move to archive","Junk Email","Never Delete","Personal 1 year move to archive","Personal 5 year move to archive"
```

Step 3: Apply the new Exchange retention policy to mailboxes on hold

The last step is to apply the new retention policy that you created in Step 2 to mailboxes on hold in your organization. You can use the EAC or Exchange Online PowerShell to apply the retention policy to a single mailbox or to multiple mailboxes.

Use the EAC to apply the new retention policy

1. Go to **Recipients** > **Mailboxes**.
2. In the list view, select the mailbox you want to apply the retention policy to, and then click **Edit** ✎.
3. On the **User Mailbox** page, click **Mailbox features**.
4. Under **Retention policy**, select the retention policy that you created in Step 2, and then click **Save**.

You can also use the EAC to apply the retention policy to multiple mailboxes.

1. Go to **Recipients** > **Mailboxes**.
2. In the list view, use the Shift or Ctrl keys to select multiple mailboxes.
3. In the details pane, click **More options**.
4. Under **Retention Policy**, click **Update**.
5. On the **Bulk assign retention policy** page, select the retention policy that you created in Step 2, and then click **Save**.

Use Exchange Online PowerShell to apply the new retention policy

You can use Exchange Online PowerShell to apply a new retention policy to a single mailbox. But the real power of PowerShell is that you can use it to quickly identify all the mailboxes in your organization that are on either Litigation Hold or In-Place Hold, and then apply the new retention policy to all mailboxes on hold in a single command. Here are some examples of using Exchange PowerShell to apply a retention policy to one or more mailboxes. All of the examples apply the retention policy that was created in Step 2.

This example applies the new retention policy to Pilar Pinilla's mailbox.

```
Set-Mailbox "Pilar Pinilla" -RetentionPolicy "MRM Policy for Mailboxes on Hold"
```

This example applies the new retention policy to all mailboxes in the organization that are on Litigation Hold.

```
$LitigationHolds = Get-Mailbox -ResultSize unlimited | Where-Object {$_.LitigationHoldEnabled -eq 'True'}
```

```
$LitigationHolds.DistinguishedName | Set-Mailbox -RetentionPolicy "MRM Policy for Mailboxes on Hold"
```

This example applies the new retention policy to all mailboxes in the organization that are on In-Place Hold.

```
$InPlaceHolds = Get-Mailbox -ResultSize unlimited | Where-Object {$_.InPlaceHolds -ne $null}
```

```
$InPlaceHolds.DistinguishedName | Set-Mailbox -RetentionPolicy "MRM Policy for Mailboxes on Hold"
```

You can use the **Get-Mailbox** cmdlet to verify that the new retention policy was applied.

Here are some examples to verify that the commands in the previous examples applied the "MRM Policy for Mailboxes on Hold" retention policy to mailboxes on Litigation Hold and mailboxes on In-Place Hold.

```
Get-Mailbox "Pilar Pinilla" | Select RetentionPolicy
```

```
Get-Mailbox -ResultSize unlimited | Where-Object {$_.LitigationHoldEnabled -eq 'True'} | FT
DisplayName,RetentionPolicy -Auto
```

```
Get-Mailbox -ResultSize unlimited | Where-Object {$_.InPlaceHolds -ne $null} | FT
DisplayName,RetentionPolicy -Auto
```

(Optional) Step 4: Run the Managed Folder Assistant to apply the new retention settings

After you apply the new Exchange retention policy to mailboxes on hold, it can take up to 7 days in Exchange Online for the Managed Folder Assistant to process these mailboxes using the settings in the new retention policy. Instead of waiting for the Managed Folder Assistant to run, you can use the **Start-ManagedFolderAssistant** cmdlet to manually trigger the assistant to process the mailboxes that you applied the new retention policy to.

Run the following command to start the Managed Folder Assistant for Pilar Pinilla's mailbox.

```
Start-ManagedFolderAssistant "Pilar Pinilla"
```

Run the following commands to start the Managed Folder Assistant for all mailboxes on hold.

```
$MailboxesOnHold = Get-Mailbox -ResultSize unlimited | Where-Object {($_.InPlaceHolds -ne $null) -or
($_.LitigationHoldEnabled -eq "True")}
```

```
$MailboxesOnHold.DistinguishedName | Start-ManagedFolderAssistant
```

More information

- After you enable a user's archive mailbox, consider telling the user that other items in their mailbox (not just items in the Recoverable Items folder) might be moved to the archive mailbox. This is because the Default MRM Policy that's assigned to Exchange Online mailboxes contains a retention tag (named Default 2 years move to archive) that moves items to the archive mailbox two years after the date the item was delivered to the mailbox or created by the user. For more information, see Default Retention Policy in Exchange Online
- After you enable a user's archive mailbox, you might also tell the user that they can recover deleted items in the Recoverable Items folder in their archive mailbox. They can do this in Outlook by selecting the **Deleted Items** folder in the archive mailbox, and then clicking **Recover Deleted Items from Server** on the **Home** tab. For more information about recovering deleted items, see Recover deleted items in Outlook for Windows.

Preserve Bcc and expanded distribution group recipients for eDiscovery

2/18/2021 • 5 minutes to read • Edit Online

In-Place Hold, Litigation Hold, and Microsoft 365 retention policies (created in the Security & Compliance Center) allow you to preserve mailbox content to meet regulatory compliance and eDiscovery requirements. Information about recipients directly addressed in the To and Cc fields of a message is included in all messages by default. But your organization may require the ability to search for and reproduce details about all recipients of a message. This includes:

- **Recipients addressed using the Bcc field of a message:** Bcc recipients are stored in the message in the sender's mailbox, but not included in headers of the message delivered to recipients.
- **Expanded distribution group recipients:** Recipients who receive the message because they're members of a distribution group to which the message was addressed, either in the To, Cc or Bcc fields.

Exchange Online and Exchange Server 2013 (Cumulative Update 7 and later versions) retain information about Bcc and expanded distribution group recipients. You can search for this information by using an In-Place eDiscovery search in the Exchange admin center (EAC) or a Content Search in the Security & Compliance Center.

How Bcc recipients and expanded distribution group recipients are preserved

As stated earlier, information about Bcc'ed recipients is stored with the message in the sender's mailbox. This information is indexed and available to eDiscovery searches and holds.

Information about expanded distribution group recipients is stored with the message after you place a mailbox on In-Place Hold or Litigation Hold. In Office 365, this information is also stored when a Microsoft 365 retention policy is applied to a mailbox. Distribution group membership is determined at the time the message is sent. The expanded recipients list stored with the message is not impacted by changes to membership of the group after the message is sent.

INFORMATION ABOUT...	IS STORED IN...	IS STORED BY DEFAULT?	IS ACCESSIBLE TO...
To and Cc recipients	Message properties in the sender and recipients' mailboxes.	Yes	Sender, recipients, and compliance officers
Bcc recipients	Message property in the sender's mailbox.	Yes	Sender and compliance officers
Expanded distribution group recipients	Message properties in the sender's mailbox.	No. Expanded distribution group recipient information is stored after a mailbox is placed on In-Place Hold or Litigation Hold, or assigned to a Microsoft 365 retention policy.	Compliance officers

Searching for messages sent to Bcc and expanded distribution group

recipients

When searching for messages sent to a recipient, eDiscovery search results now include messages sent to a distribution group that the recipient is a member of. The following table shows the scenarios where messages sent to Bcc and expanded distribution group recipients are returned in eDiscovery searches.

Scenario 1: John is a member of the US-Sales distribution group. This table shows eDiscovery search results when Bob sends a message to John directly or indirectly via a distribution group.

WHEN YOU SEARCH BOB'S MAILBOX FOR MESSAGES SENT...	AND THE MESSAGE IS SENT WITH...	RESULTS INCLUDE MESSAGE?
To:John	John on TO	Yes
To:John	US-Sales on TO	Yes
To:US-Sales	US-Sales on TO	Yes
Cc:John	John on CC	Yes
Cc:John	US-Sales on CC	Yes
Cc:US-Sales	US-Sales on CC	Yes

Scenario 2: Bob sends an email to John (To/Cc) and Jack (Bcc directly, or indirectly via a distribution group). The table below shows eDiscovery search results.

WHEN YOU SEARCH...	FOR MESSAGES SENT...	RESULTS INCLUDE MESSAGE?	NOTES
Bob's mailbox	To/Cc:John	Yes	Presents an indication that Jack was Bcc'ed.
Bob's mailbox	Bcc:Jack	Yes	Presents an indication that Jack was Bcc'ed.
Bob's mailbox	Bcc:Jack (via distribution group)	Yes	List of members of the Bcc'ed distribution group, expanded when the message was sent, is visible in eDiscovery search preview, export, and logs.
John's mailbox	To/Cc:John	Yes	No indication of Bcc recipients.
John's mailbox	Bcc:Jack (directly or via distribution group)	No	Bcc information is not stored in the message delivered to recipients. You must search the sender's mailbox.
Jack's mailbox	To/Cc:John (directly or via distribution group)	Yes	To/Cc information is included in message delivered to all recipients.

WHEN YOU SEARCH...	FOR MESSAGES SENT...	RESULTS INCLUDE MESSAGE?	NOTES
Jack's mailbox	Bcc:Jack (directly or via distribution group)	No	Bcc information is not stored in the message delivered to recipients. You must search the sender's mailbox.

Frequently asked questions

Q. When and where is Bcc recipient information stored?

A. Bcc recipient information is preserved by default in the original message in sender's mailbox. If the Bcc recipient is a distribution group, distribution group membership is only expanded if the sender's mailbox is on hold or assigned to a Microsoft 365 retention policy.

Q. When and where is the list of expanded distribution group recipients stored?

A. Group membership is expanded at the time the message is sent. The list of expanded distribution group members is stored in the original message in the sender's mailbox. The sender's mailbox must be on In-Place Hold, Litigation Hold, or assigned to a Microsoft 365 retention policy.

Q. Can the To/Cc recipients see which recipients were Bcc'ed?

A. No. This information is not included in message headers, and isn't visible to To/Cc recipients. The sender can see the Bcc field stored in the original message stored in their mailbox. Compliance officers can see this information when searching the sender's mailbox.

Q. How can I ensure that expanded distribution group recipients are always preserved?

A. To ensure that expanded distribution group members are always preserved with a message, Place all mailboxes on hold or create an organization-wide Microsoft 365 retention policy.

Q. Which types of groups are supported?

A. Distribution groups, mail-enabled security groups, and dynamic distribution groups are supported.

Q. Is there a limit on the number of distribution group recipients that are expanded and stored in the message?

A. Up to 10,000 members of a distribution group is preserved.

Q. Are nested distribution groups supported?

A. Yes, 25 levels of nested distribution groups are expanded.

Q. Where is the Bcc and expanded distribution group recipient information visible?

A. Bcc and expanded distribution group recipients information is visible to Compliance officers when performing an eDiscovery search. Bcc and expanded distribution group recipients are included in search results copied to a Discovery mailbox or exported to a PST file and in the eDiscovery log included in search results. Bcc recipient information is also available in search preview.

Q. What happens if a member of a distribution group is hidden from the organization's global address list (GAL)?

A. There's no impact. If recipients are hidden from the GAL, they are still included in the list of recipients for the expanded distribution group.

Search the audit log in the compliance center

2/18/2021 • 73 minutes to read • Edit Online

Need to find if a user viewed a specific document or purged an item from their mailbox? If so, you can use the Microsoft 365 compliance center to search the unified audit log to view user and administrator activity in your organization. Why a unified audit log? Because you can search for the following types of user and admin activity in Microsoft 365:

- User activity in SharePoint Online and OneDrive for Business
- User activity in Exchange Online (Exchange mailbox audit logging)
- Admin activity in SharePoint Online
- Admin activity in Azure Active Directory (the directory service for Microsoft 365)
- Admin activity in Exchange Online (Exchange admin audit logging)
- eDiscovery activities in the security and compliance center
- User and admin activity in Power BI
- User and admin activity in Microsoft Teams
- User and admin activity in Dynamics 365
- User and admin activity in Yammer
- User and admin activity in Microsoft Power Automate
- User and admin activity in Microsoft Stream
- Analyst and admin activity in Microsoft Workplace Analytics
- User and admin activity in Microsoft Power Apps
- User and admin activity in Microsoft Forms
- User and admin activity for sensitivity labels for sites that use SharePoint Online or Microsoft Teams

Requirements to search the audit log

Be sure to read the following items before you start searching the audit log.

- Audit log search is turned on by default for Microsoft 365 and Office 365 enterprise organizations. This includes organizations with E3/G3 or E5/G5 subscriptions. To verify that audit log search is turned on, you can run the following command in Exchange Online PowerShell:

  ```
  Get-AdminAuditLogConfig | FL UnifiedAuditLogIngestionEnabled
  ```

  The value of `True` for the *UnifiedAuditLogIngestionEnabled* property indicates that audit log search is turned on. For more information, see Turn audit log search on or off.

- You have to be assigned the View-Only Audit Logs or Audit Logs role in Exchange Online to search the audit log. By default, these roles are assigned to the Compliance Management and Organization Management role groups on the **Permissions** page in the Exchange admin center. Note global

administrators in Office 365 and Microsoft 365 are automatically added as members of the Organization Management role group in Exchange Online. To give a user the ability to search the audit log with the minimum level of privileges, you can create a custom role group in Exchange Online, add the View-Only Audit Logs or Audit Logs role, and then add the user as a member of the new role group. For more information, see Manage role groups in Exchange Online.

> **IMPORTANT**
> If you assign a user the View-Only Audit Logs or Audit Logs role on the **Permissions** page in the Security & Compliance Center, they won't be able to search the audit log. You have to assign the permissions in Exchange Online. This is because the underlying cmdlet used to search the audit log is an Exchange Online cmdlet.

- When an audited activity is performed by a user or admin, an audit record is generated and stored in the audit log for your organization. The length of time that an audit record is retained (and searchable in the audit log) depends on your Office 365 or Microsoft 365 Enterprise subscription, and specifically the type of the license that is assigned to specific users.
  - For users assigned an Office 365 E5 or Microsoft 365 E5 license (or users with a Microsoft 365 E5 Compliance or Microsoft 365 E5 eDiscovery and Audit add-on license), audit records for Azure Active Directory, Exchange, and SharePoint activity are retained for one year by default. Organizations can also create audit log retention policies to retain audit records for activities in other services for up to one year. For more information, see Manage audit log retention policies.

    > **NOTE**
    > If your organization participated in the private preview program for the one-year retention of audit records, the retention duration for audit records that were generated before the general availability rollout date will not be reset.

  - For users assigned any other (non-E5) Office 365 or Microsoft 365 license, audit records are retained for 90 days. For a list of Office 365 and Microsoft 365 subscriptions that support unified audit logging, see the security and compliance center service description.

    > **NOTE**
    > Even when mailbox auditing on by default is turned on, you might notice that mailbox audit events for some users aren't found in audit log searches in the Security & Compliance Center or via the Office 365 Management Activity API. For more information, see More information about mailbox audit logging.

- If you want to turn off audit log search for your organization, you can run the following command in remote PowerShell connected to your Exchange Online organization:

  ```
  Set-AdminAuditLogConfig -UnifiedAuditLogIngestionEnabled $false
  ```

  To turn on audit search again, you can run the following command in Exchange Online PowerShell:

  ```
  Set-AdminAuditLogConfig -UnifiedAuditLogIngestionEnabled $true
  ```

  For more information, see Turn off audit log search.

- As previously stated, the underlying cmdlet used to search the audit log is an Exchange Online cmdlet, which is **Search-UnifiedAuditLog**. That means you can use this cmdlet to search the audit log instead

of using the **Audit log search** page in the Security & Compliance Center. You have to run this cmdlet in remote PowerShell connected to your Exchange Online organization. For more information, see Search-UnifiedAuditLog.

For information about exporting the search results returned by the **Search-UnifiedAuditLog** cmdlet to a CSV file, see the "Tips for exporting and viewing the audit log" section in Export, configure, and view audit log records.

- If you want to programmatically download data from the audit log, we recommend that you use the Office 365 Management Activity API instead of using a PowerShell script. The Office 365 Management Activity API is a REST web service that you can use to develop operations, security, and compliance monitoring solutions for your organization. For more information, see Office 365 Management Activity API reference.
- It can take up to 30 minutes or up to 24 hours after an event occurs for the corresponding audit log record to be returned in the results of an audit log search. The following table shows the time it takes for the different services in Office 365.

MICROSOFT 365 SERVICE OR FEATURE	30 MINUTES	24 HOURS	
Defender for Office 365 and Threat Intelligence	✓		
Azure Active Directory (user login events)		✓	
Azure Active Directory (admin events)		✓	
Data Loss Prevention	✓		
Dynamics 365 CRM		✓	
eDiscovery	✓		
Exchange Online	✓		
Microsoft Power Automate		✓	
Microsoft Project	✓		
Microsoft Stream	✓		
Microsoft Teams	✓		
Power Apps		✓	
Power BI	✓		
Security & Compliance Center	✓		

MICROSOFT 365 SERVICE OR FEATURE	30 MINUTES	24 HOURS	
Sensitivity labels		✓	
SharePoint Online and OneDrive for Business	✓		
Workplace Analytics	✓		
Yammer		✓	
Microsoft Forms	✓		

- Azure Active Directory (Azure AD) is the directory service for Office 365. The unified audit log contains user, group, application, domain, and directory activities performed in the Microsoft 365 admin center or in the Azure management portal. For a complete list of Azure AD events, see Azure Active Directory Audit Report Events.
- Audit logging for Power BI isn't enabled by default. To search for Power BI activities in the audit log, you have to enable auditing in the Power BI admin portal. For instructions, see the "Audit logs" section in Power BI admin portal.

Search the audit log

> **NOTE**
> There was an issue with Azure AD activities being unavailable in the audit log search tool from October 22, 2020 to November 6, 2020. These activites include Azure AD user administration activities, group administration activities, application administration activities, role administration activities, and directory administration activities. The missing events for the period of impact will be available over the next few days, and is expected to take no later than November 20, 2020 to complete. In some cases, customers might notice duplicate event data for events generated between October 26, 2020 and November 05, 2020.

Here's the process for searching the audit log in Office 365.

Step 1: Run an audit log search

Step 2: View the search results

Step 3: Filter the search results

Step 4: Export the search results to a file

Step 1: Run an audit log search

1. Go to https://protection.office.com.

> **TIP**
> Use a private browsing session (not a regular session) to access the Security & Compliance Center because this will prevent the credential that you are currently logged on with from being used. To open an InPrivate Browsing session in Internet Explorer or Microsoft Edge, just press CTRL+SHIFT+P. To open a private browsing session in Google Chrome (called an incognito window), press CTRL+SHIFT+N.

2. Sign in using your work or school account.
3. In the left pane of the Security & Compliance Center, click **Search**, and then click **Audit log search**.

   The **Audit log search** page is displayed.

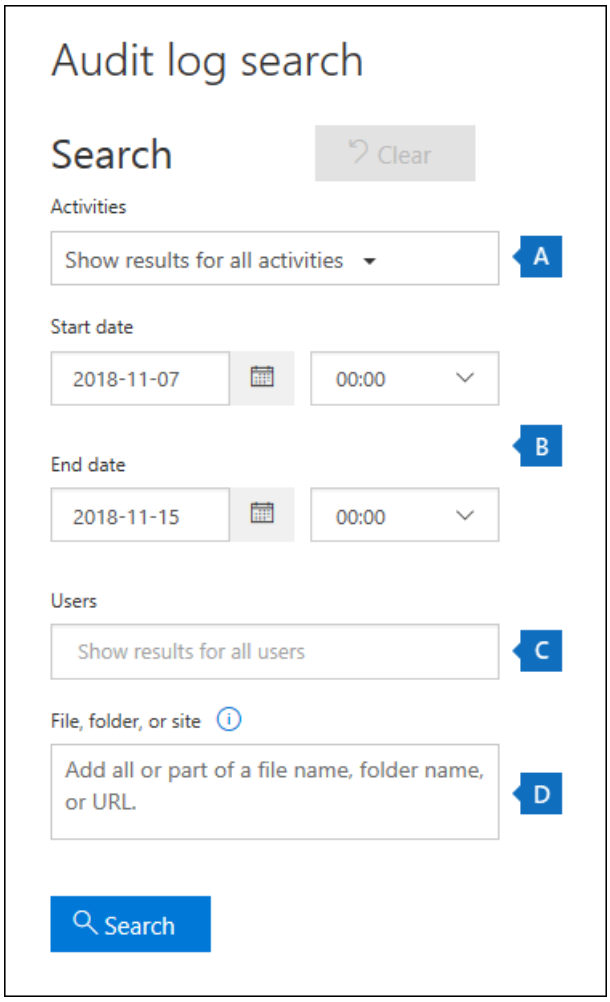

> **NOTE**
> You have to first turn on audit logging before you can run an audit log search. If the **Start recording user and admin activity** link is displayed, click it to turn on auditing. If you don't see this link, auditing has already been turned on for your organization.

4. Configure the following search criteria:

   a. **Activities**: Click the drop-down list to display the activities that you can search for. User and admin activities are organized into groups of related activities. You can select specific activities or you can click the activity group name to select all activities in the group. You can also click a selected activity to clear the selection. After you run the search, only the audit log entries for the selected activities are displayed. Selecting **Show results for all activities** displays results for all activities performed by the selected user or group of users.

      Over 100 user and admin activities are logged in the audit log. Click the **Audited activities** tab at the topic of this article to see the descriptions of every activity in each of the different services.

   b. **Start date** and **End date**: The last seven days are selected by default. Select a date and time range to display the events that occurred within that period. The date and time are presented in Coordinated Universal Time (UTC) format. The maximum date range that you can specify is 90 days. An error is displayed if the selected date range is greater than 90 days.

> **TIP**
>
> If you're using the maximum date range of 90 days, select the current time for the **Start date**. Otherwise, you'll receive an error saying that the start date is earlier than the end date. If you've turned on auditing within the last 90 days, the maximum date range can't start before the date that auditing was turned on.

c. **Users**: Click in this box and then select one or more users to display search results for. The audit log entries for the selected activity performed by the users you select in this box are displayed in the list of results. Leave this box blank to return entries for all users (and service accounts) in your organization.

d. **File, folder, or site**: Type some or all of a file or folder name to search for activity related to the file of folder that contains the specified keyword. You can also specify a URL of a file or folder. If you use a URL, be sure the type the full URL path or if you type a portion of the URL, don't include any special characters or spaces.

   Leave this box blank to return entries for all files and folders in your organization.

> **TIP**
>
> - If you're looking for all activities related to a **site**, add the wildcard symbol (*) after the URL to return all entries for that site; for example, `"https://contoso-my.sharepoint.com/personal*"`.
> - If you're looking for all activities related to a **file**, add the wildcard symbol (*) before the file name to return all entries for that file; for example, `"*Customer_Profitability_Sample.csv"`.

5. Click **Search** to run the search using your search criteria.

   The search results are loaded, and after a few moments they are displayed under **Results**. When the search is finished, the number of results found is displayed. A maximum of 5,000 events will be displayed in the **Results** pane in increments of 150 events. If more than 5,000 events meet the search criteria, the most recent 5,000 events are displayed.

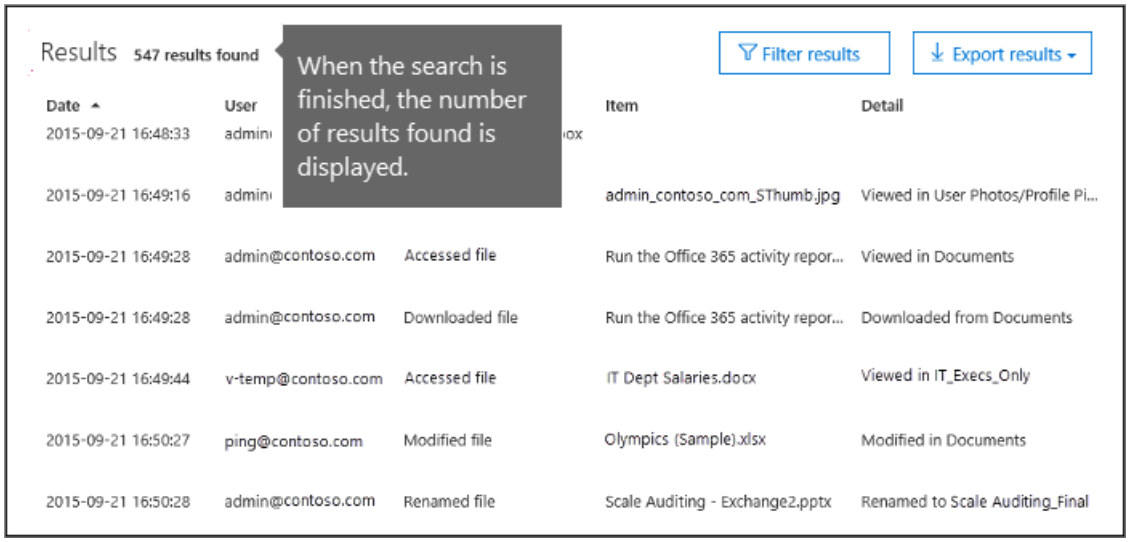

Tips for searching the audit log

- You can select specific activities to search for by clicking the activity name. Or you can search for all activities in a group (such as **File and folder activities**) by clicking the group name. If an activity is selected, you can click it to cancel the selection. You can also use the search box to display the activities that contain the keyword that you type.

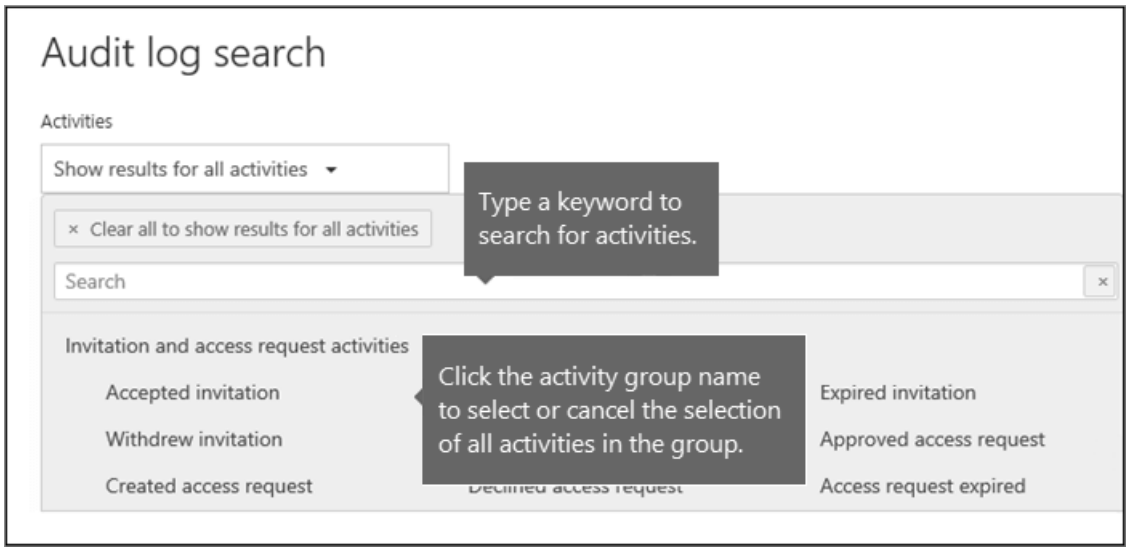

- You have to select **Show results for all activities** in the **Activities** list to display events from the Exchange admin audit log. Events from this audit log display a cmdlet name (for example, **Set-Mailbox**) in the **Activity** column in the results. For more information, click the **Audited activities** tab in this topic and then click **Exchange admin activities**.

  Similarly, there are some auditing activities that don't have a corresponding item in the **Activities** list. If you know the name of the operation for these activities, you can search for all activities, then filter the results by typing the name of the operation in the box for the **Activity** column. See Step 3: Filter the search results for more information about filtering the results.

- Click **Clear** to clear the current search criteria. The date range returns to the default of the last seven days. You can also click **Clear all to show results for all activities** to cancel all selected activities.

- If 5,000 results are found, you can probably assume that there are more than 5,000 events that met the search criteria. You can either refine the search criteria and rerun the search to return fewer results, or you can export all of the search results by selecting **Export results** > **Download all results**.

Step 2: View the search results

The results of an audit log search are displayed under **Results** on the **Audit log search** page. As previously stated a maximum of 5,000 (newest) events are displayed in increments of 150 events. To display more events you can use the scroll bar in the **Results** pane or you can press **Shift + End** to display the next 150 events.

The results contain the following information about each event returned by the search:

- **Date**: The date and time (in UTC format) when the event occurred.

- **IP address**: The IP address of the device that was used when the activity was logged. The IP address is displayed in either an IPv4 or IPv6 address format.

  > **NOTE**
  > For some services, the value displayed in this field might be the IP address for a trusted application (for example, Office on the web apps) calling into the service on behalf of a user and not the IP address of the device used by person who performed the activity. Also, for admin activity (or activity performed by a system account) for Azure Active Directory-related events, the IP address isn't logged and the value displayed in this field is `null`.

- **User**: The user (or service account) who performed the action that triggered the event.

- **Activity**: The activity performed by the user. This value corresponds to the activities that you selected in the **Activities** drop down list. For an event from the Exchange admin audit log, the value in this column is an Exchange cmdlet.

- **Item**: The object that was created or modified as a result of the corresponding activity. For example, the file that was viewed or modified or the user account that was updated. Not all activities have a value in this column.
- **Detail**: Additional information about an activity. Again, not all activities have a value.

> **TIP**
> Click a column header under **Results** to sort the results. You can sort the results from A to Z or Z to A. Click the **Date** header to sort the results from oldest to newest or newest to oldest.

View the details for a specific event

You can view more details about an event by clicking the event record in the list of search results. A **Details** page is displayed that contains the detailed properties from the event record. The properties that are displayed depend on the service in which the event occurs. To display these details, click **More information**. For descriptions, see Detailed properties in the audit log.

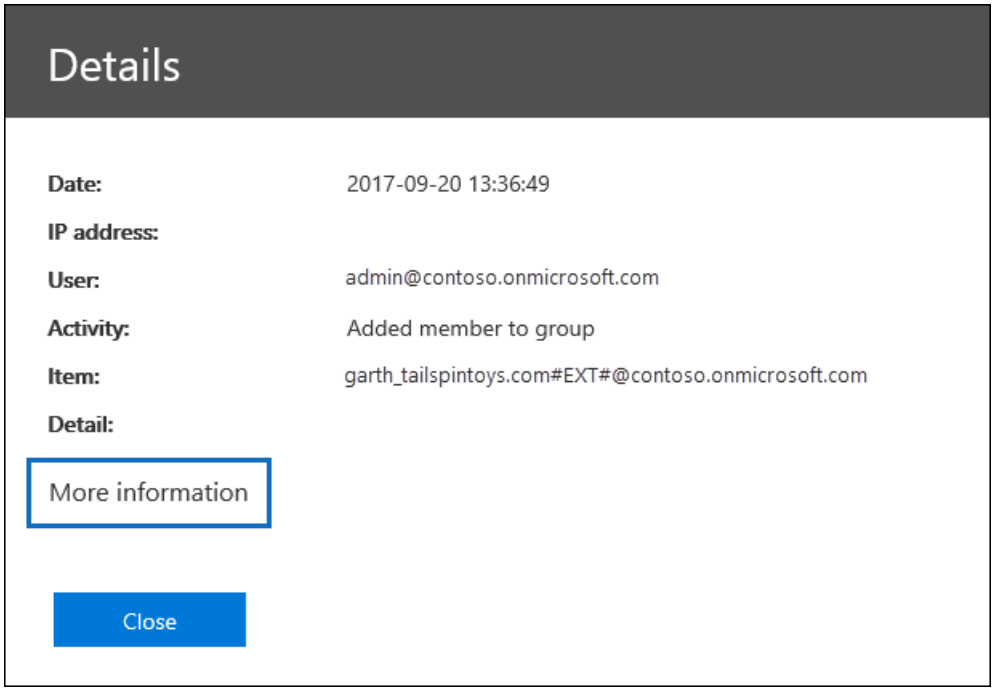

Step 3: Filter the search results

In addition to sorting, you can also filter the results of an audit log search. This is a great feature that can help you quickly filter the results for a specific user or activity. You can initially create a wide search and then quickly filter the results to see specific events. Then you can narrow the search criteria and rerun the search to return a smaller, more concise set of results.

To filter the results:

1. Run an audit log search.
2. When the results are displayed, click **Filter results**.

   Keyword boxes are displayed under each column header.
3. Click one of the boxes under a column header and type a word or phrase, depending on the column you're filtering on. The results will dynamically readjust to display the events that match your filter.

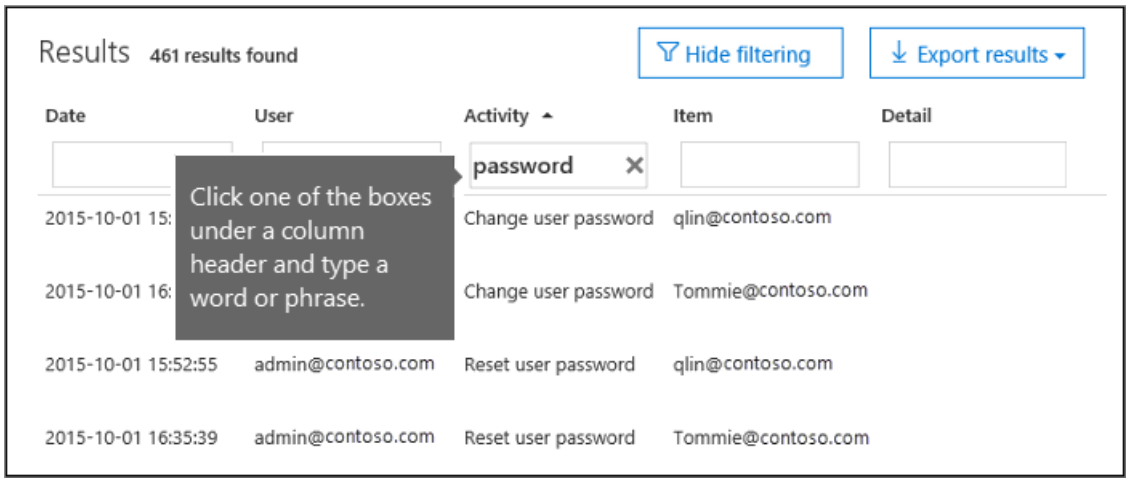

4. To clear a filter, click the **X** in the filter box or click **Hide filtering**.

> **TIP**
> To display events from the Exchange admin audit log, type a - (dash) in the **Activity** filter box. This will display cmdlet names, which are displayed in the **Activity** column for Exchange admin events. Then you can sort the cmdlet names in alphabetical order.

Step 4: Export the search results to a file

You can export the results of an audit log search to a comma-separated value (CSV) file on your local computer. You can open this file in Microsoft Excel and use features such as search, sorting, filtering, and splitting a single column (that contains multiple properties) into multiple columns.

1. Run an audit log search, and then revise the search criteria until you have the desired results.

2. Click **Export results** and select one of the following options:

   - **Save loaded results**: Choose this option to export only the entries that are displayed under **Results** on the **Audit log search** page. The CSV file that is downloaded contains the same columns (and data) displayed on the page (Date, User, Activity, Item, and Details). An extra column (named **More**) is included in the CSV file that contains more information from the audit log entry. Because you're exporting the same results that are loaded (and viewable) on the **Audit log search** page, a maximum of 5,000 entries are exported.

   - **Download all results**: Choose this option to export all entries from the audit log that meet the search criteria. For a large set of search results, choose this option to download all entries from the audit log in addition to the 5,000 audit records that can be displayed on the **Audit log search** page. This option downloads the raw data from the audit log to a CSV file, and contains additional information from the audit log entry in a column named **AuditData**. It may take longer to download the file if you choose this export option because the file may be much larger than the one that's downloaded if you choose the other option.

     > **IMPORTANT**
     > You can download a maximum of 50,000 entries to a CSV file from a single audit log search. If 50,000 entries are downloaded to the CSV file, you can probably assume there are more than 50,000 events that met the search criteria. To export more than this limit, try using a date range to reduce the number of audit log entries. You might have to run multiple searches with smaller date ranges to export more than 50,000 entries.

3. After you select an export option, a message is displayed at the bottom of the window that prompts you

to open the CSV file, save it to the Downloads folder, or save it to a specific folder.

More information about exporting and viewing audit log search results

- If you download all search results, the CSV file contains a column named **AuditData**, which contains additional information about each event. The data in this column consists of a JSON object that contains multiple properties from the audit log record. Each *property:value* pair in the JSON object is separated by a comma. You can use the JSON transform tool in the Power Query Editor in Excel to split **AuditData** column into multiple columns so that each property in the JSON object has its own column. This lets you sort and filter on one or more of these properties. For step-by-step instructions using the Power Query Editor to transform the JSON object, see Export, configure, and view audit log records.

  After you split the **AuditData** column, you can filter on the **Operations** column to display the detailed properties for a specific type of activity.

- The **Download all results** option downloads the raw data from the audit log to a CSV file. This file contains different column names (CreationDate, UserIds, Operation, AuditData) than the file that's downloaded if you select the **Save loaded results** option. The values in the two different CSV files for the same activity may also be different. For example, the activity in the **Action** column in the CSV file and may have a different value than the "user-friendly" name that's displayed in the **Activity** column on the **Audit log search** page. For example, MailboxLogin vs. User signed in to mailbox.

- When you download all results from a search query that contains events from different services, the **AuditData** column in the CSV file contains different properties depending on which service the action was performed in. For example, entries from Exchange and Azure AD audit logs include a property named **ResultStatus** that indicates if the action was successful or not. This property isn't included for events in SharePoint. Similarly, SharePoint events have a property that identifies the site URL for file and folder-related activities. To mitigate this behavior, consider using different searches to export the results for activities from a single service.

  For a description of many of the properties that are listed in the **AuditData** column in the CSV file when you download all results, and the service each one applies to, see Detailed properties in the audit log.

Audited activities

The tables in this section describe the activities that are audited in Office 365. You can search for these events by searching the audit log in the security and compliance center.

These tables group related activities or the activities from a specific service. The tables include the friendly name that's displayed in the **Activities** drop-down list and the name of the corresponding operation that appears in the detailed information of an audit record and in the CSV file when you export the search results. For descriptions of the detailed information, see Detailed properties in the audit log.

Click one of the following links to go to a specific table.

User administration activities

Azure AD group administration activities

Application administration activities

Role administration activities

Directory administration activities

eDiscovery activities

Advanced eDiscovery activities

Power BI activities

Microsoft Workplace Analytics

Microsoft Teams activities

Microsoft Teams Healthcare activities

Microsoft Teams Shifts activities

Yammer activities

Microsoft Power Automate activities

Microsoft Power Apps activities

Microsoft Stream activities

Content explorer activities

Quarantine activities

Microsoft Forms activities

Sensitivity label activities

Retention policy and retention label activities

Exchange admin activities

File and page activities

The following table describes the file and page activities in SharePoint Online and OneDrive for Business.

FRIENDLY NAME	OPERATION	DESCRIPTION
Accessed file	FileAccessed	User or system account accesses a file.

FRIENDLY NAME	OPERATION	DESCRIPTION
(none)	FileAccessedExtended	This is related to the "Accessed file" (FileAccessed) activity. A FileAccessedExtended event is logged when the same person continually accesses a file for an extended period (up to 3 hours). The purpose of logging FileAccessedExtended events is to reduce the number of FileAccessed events that are logged when a file is continually accessed. This helps reduce the noise of multiple FileAccessed records for what is essentially the same user activity, and lets you focus on the initial (and more important) FileAccessed event.
Changed retention label for a file	ComplianceSettingChanged	A retention label was applied to or removed from a document. This event is triggered when a retention label is manually or automatically applied to a message.
Changed record status to locked	LockRecord	The record status of a retention label that classifies a document as a record was locked. This means the document can't be modified or deleted. Only users assigned at least the contributor permission for a site can change the record status of a document.
Changed record status to unlocked	UnlockRecord	The record status of a retention label that classifies a document as a record was unlocked. This means that the document can be modified or deleted. Only users assigned at least the contributor permission for a site can change the record status of a document.
Checked in file	FileCheckedIn	User checks in a document that they checked out from a document library.
Checked out file	FileCheckedOut	User checks out a document located in a document library. Users can check out and make changes to documents that have been shared with them.
Copied file	FileCopied	User copies a document from a site. The copied file can be saved to another folder on the site.
Deleted file	FileDeleted	User deletes a document from a site.
Deleted file from recycle bin	FileDeletedFirstStageRecycleBin	User deletes a file from the recycle bin of a site.

FRIENDLY NAME	OPERATION	DESCRIPTION
Deleted file from second-stage recycle bin	FileDeletedSecondStageRecycleBin	User deletes a file from the second-stage recycle bin of a site.
Deleted file marked as a record	RecordDelete	A document or email that was marked as a record was deleted. An item is considered a record when a retention label that marks items as a record is applied to content.
Detected document sensitivity mismatch	DocumentSensitivityMismatchDetected	User uploads a document to a site that's protected with a sensitivity label and the document has a higher priority sensitivity label than the sensitivity label applied to the site. For example, a document labeled Confidential is uploaded to a site labeled General. This event isn't triggered if the document has a lower priority sensitivity label than the sensitivity label applied to the site. For example, a document labeled General is uploaded to a site labeled Confidential. For more information about sensitivity label priority, see Label priority (order matters).
Detected malware in file	FileMalwareDetected	SharePoint anti-virus engine detects malware in a file.
Discarded file checkout	FileCheckOutDiscarded	User discards (or undoes) a checked out file. That means any changes they made to the file when it was checked out are discarded, and not saved to the version of the document in the document library.
Downloaded file	FileDownloaded	User downloads a document from a site.
Modified file	FileModified	User or system account modifies the content or the properties of a document on a site.

FRIENDLY NAME	OPERATION	DESCRIPTION
(none)	FileModifiedExtended	This is related to the "Modified file" (FileModified) activity. A FileModifiedExtended event is logged when the same person continually modifies a file for an extended period (up to 3 hours). The purpose of logging FileModifiedExtended events is to reduce the number of FileModified events that are logged when a file is continually modified. This helps reduce the noise of multiple FileModified records for what is essentially the same user activity, and lets you focus on the initial (and more important) FileModified event.
Moved file	FileMoved	User moves a document from its current location on a site to a new location.
(none)	FilePreviewed	User previews files on a SharePoint or OneDrive for Business site. These events typically occur in high volumes based on a single activity, such as viewing an image gallery.
Performed search query	SearchQueryPerformed	User or system account performs a search in SharePoint or OneDrive for Business. Some common scenarios where a service account performs a search query include applying an eDiscovery holds and retention policy to sites and OneDrive accounts, and auto-applying retention or sensitivity labels to site content.
Recycled all minor versions of file	FileVersionsAllMinorsRecycled	User deletes all minor versions from the version history of a file. The deleted versions are moved to the site's recycle bin.
Recycled all versions of file	FileVersionsAllRecycled	User deletes all versions from the version history of a file. The deleted versions are moved to the site's recycle bin.
Recycled version of file	FileVersionRecycled	User deletes a version from the version history of a file. The deleted version is moved to the site's recycle bin.
Renamed file	FileRenamed	User renames a document on a site.
Restored file	FileRestored	User restores a document from the recycle bin of a site.

FRIENDLY NAME	OPERATION	DESCRIPTION
Uploaded file	FileUploaded	User uploads a document to a folder on a site.
Viewed page	PageViewed	User views a page on a site. This doesn't include using a Web browser to view files located in a document library.
(none)	PageViewedExtended	This is related to the "Viewed page" (PageViewed) activity. A PageViewedExtended event is logged when the same person continually views a web page for an extended period (up to 3 hours). The purpose of logging PageViewedExtended events is to reduce the number of PageViewed events that are logged when a page is continually viewed. This helps reduce the noise of multiple PageViewed records for what is essentially the same user activity, and lets you focus on the initial (and more important) PageViewed event.
View signaled by client	ClientViewSignaled	A user's client (such as website or mobile app) has signaled that the indicated page has been viewed by the user. This activity is often logged following a PagePrefetched event for a page. **NOTE**: Because ClientViewSignaled events are signaled by the client, rather than the server, it's possible the event may not be logged by the server and therefore may not appear in the audit log. It's also possible that information in the audit record may not be trustworthy. However, because the user's identity is validated by the token used to create the signal, the user's identity listed in the corresponding audit record is accurate.

FRIENDLY NAME	OPERATION	DESCRIPTION
(none)	PagePrefetched	A user's client (such as website or mobile app) has requested the indicated page to help improve performance if the user browses to it. This event is logged to indicate that the page content has been served to the user's client. This event isn't a definitive indication that the user navigated to the page. When the page content is rendered by the client (as per the user's request) a ClientViewSignaled event should be generated. Not all clients support indicating a pre-fetch, and therefore some pre-fetched activities might instead be logged as PageViewed events.

Frequently asked questions about FileAccessed and FilePreviewed events

Could any non-user activities trigger FilePreviewed audit records that contain a user agent like "OneDriveMpc-Transform_Thumbnail"?

We aren't aware of scenarios where non-user actions generate events like these. User actions like opening a user profile card (by clicking their name or email address in a message in Outlook on the web) would generate similar events.

Are calls to the OneDriveMpc-Transform_Thumbnail always intentionally being triggered by the user?

No. But similar events can be logged as a result of browser pre-fetch.

If we see a FilePreviewed event coming from a Microsoft-registered IP address, does that mean that the preview was displayed on the screen of the user's device?

No. The event might have been logged as a result of browser pre-fetch.

Are there scenarios where a user previewing a document generates FileAccessed events?

Both the FilePreviewed and FileAccessed events indicate that a user's call led to a read of the file (or a read of a thumbnail rendering of the file). While these events are intended to align with preview vs. access intention, the event distinction isn't a guarantee of the user's intent.

The app@sharepoint user in audit records

In audit records for some file activities (and other SharePoint-related activities), you may notice the user who performed the activity (identified in the User and UserId fields) is app@sharepoint. This indicates that the "user" who performed the activity was an application. In this case, the application was granted permissions in SharePoint to perform organization-wide actions (such as search a SharePoint site or OneDrive account) on behalf of a user, admin, or service. This process of giving permissions to an application is called *SharePoint App-Only* access. This indicates that the authentication presented to SharePoint to perform an action was made by an application, instead of a user. This is why the app@sharepoint user is identified in certain audit records. For more information, see Grant access using SharePoint App-Only.

For example, app@sharepoint is often identified as the user for "Performed search query" and "Accessed file" events. That's because an application with SharePoint App-Only access in your organization performs search queries and accesses files when applying retention policies to sites and OneDrive accounts.

Here are a few other scenarios where app@sharepoint may be identified in an audit record as the user who performed an activity:

- Microsoft 365 Groups. When a user or admin creates a new group, audit records are generated for creating a site collection, updating lists, and adding members to a SharePoint group. These tasks are performed an application on behalf of the user who created the group.
- Microsoft Teams. Similar to Microsoft 365 Groups, audit records are generated for creating a site collection, updating lists, and adding members to a SharePoint group when a team is created.
- Compliance features. When an admin implements compliance features, such as retention policies, eDiscovery holds, and auto-applying sensitivity labels.

In these and other scenarios, you'll also notice that multiple audit records with app@sharepoint as the specified user were created within a short time frame, often within a few seconds of each other. This also indicates they were probably triggered by the same user-initiated task. Also, the ApplicationDisplayName and EventData fields in the audit record may help you identify the scenario or application that triggered the event.

Folder activities

The following table describes the folder activities in SharePoint Online and OneDrive for Business. As previously explained, audit records for some SharePoint activities will indicate the app@sharepoint user performed the activity of behalf of the user or admin who initiated the action. For more information, see The app@sharepoint user in audit records.

FRIENDLY NAME	OPERATION	DESCRIPTION
Copied folder	FolderCopied	User copies a folder from a site to another location in SharePoint or OneDrive for Business.
Created folder	FolderCreated	User creates a folder on a site.
Deleted folder	FolderDeleted	User deletes a folder from a site.
Deleted folder from recycle bin	FolderDeletedFirstStageRecycleBin	User deletes a folder from the recycle bin on a site.
Deleted folder from second-stage recycle bin	FolderDeletedSecondStageRecycleBin	User deletes a folder from the second-stage recycle bin on a site.
Modified folder	FolderModified	User modifies a folder on a site. This includes changing the folder metadata, such as changing tags and properties.
Moved folder	FolderMoved	User moves a folder to a different location on a site.
Renamed folder	FolderRenamed	User renames a folder on a site.
Restored folder	FolderRestored	User restores a deleted folder from the recycle bin on a site.

SharePoint list activities

The following table describes activities related to when users interact with lists and list items in SharePoint Online. As previously explained, audit records for some SharePoint activities will indicate the app@sharepoint

user performed the activity of behalf of the user or admin who initiated the action. For more information, see The app@sharepoint user in audit records.

FRIENDLY NAME	OPERATION	DESCRIPTION
Created list	ListCreated	A user created a SharePoint list.
Created list column	ListColumnCreated	A user created a SharePoint list column. A list column is a column that's attached to one or more SharePoint lists.
Created list content type	ListContentTypeCreated	A user created a list content type. A list content type is a content type that's attached to one or more SharePoint lists.
Created list item	ListItemCreated	A user created an item in an existing SharePoint list.
Created site column	SiteColumnCreated	A user created a SharePoint site column. A site column is a column that isn't attached to a list. A site column is also a metadata structure that can be used by any list in a given web.
Created site content type	Site ContentType Created	A user created a site content type. A site content type is a content type that's attached to the parent site.
Deleted list	ListDeleted	A user deleted a SharePoint list.
Deleted list column	List Column Deleted	A user deleted a SharePoint list column.
Deleted list content type	ListContentTypeDeleted	A user deleted a list content type.
Deleted list item	List Item Deleted	A user deleted a SharePoint list item.
Deleted site column	SiteColumnDeleted	A user deleted a SharePoint site column.
Deleted site content type	SiteContentTypeDeleted	A user deleted a site content type.
Recycled list item	ListItemRecycled	A user moved a SharePoint list item to the Recycle Bin.
Restored list	ListRestored	A user restored a SharePoint list from the Recycle Bin.
Restored list item	ListItemRestored	A user restored a SharePoint list item from the Recycle Bin.
Updated list	ListUpdated	A user updated a SharePoint list by modifying one or more properties.

FRIENDLY NAME	OPERATION	DESCRIPTION
Updated list column	ListColumnUpdated	A user updated a SharePoint list column by modifying one or more properties.
Updated list content type	ListContentTypeUpdated	A user updated a list content type by modifying one or more properties.
Updated list item	ListItemUpdated	A user updated a SharePoint list item by modifying one or more properties.
Updated site column	SiteColumnUpdated	A user updated a SharePoint site column by modifying one or more properties.
Updated site content type	SiteContentTypeUpdated	A user updated a site content type by modifying one or more properties.

Sharing and access request activities

The following table describes the user sharing and access request activities in SharePoint Online and OneDrive for Business. For sharing events, the **Detail** column under **Results** identifies the name of the user or group the item was shared with and whether that user or group is a member or guest in your organization. For more information, see Use sharing auditing in the audit log.

> **NOTE**
> Users can be either *members* or *guests* based on the UserType property of the user object. A member is usually an employee, and a guest is usually a collaborator outside of your organization. When a user accepts a sharing invitation (and isn't already part of your organization), a guest account is created for them in your organization's directory. Once the guest user has an account in your directory, resources may be shared directly with them (without requiring an invitation).

FRIENDLY NAME	OPERATION	DESCRIPTION
Added permission level to site collection	PermissionLevelAdded	A permission level was added to a site collection.
Accepted access request	AccessRequestAccepted	An access request to a site, folder, or document was accepted and the requesting user has been granted access.
Accepted sharing invitation	SharingInvitationAccepted	User (member or guest) accepted a sharing invitation and was granted access to a resource. This event includes information about the user who was invited and the email address that was used to accept the invitation (they could be different). This activity is often accompanied by a second event that describes how the user was granted access to the resource, for example, adding the user to a group that has access to the resource.

FRIENDLY NAME	OPERATION	DESCRIPTION
Blocked sharing invitation	SharingInvitationBlocked	A sharing invitation sent by a user in your organization is blocked because of an external sharing policy that either allows or denies external sharing based on the domain of the target user. In this case, the sharing invitation was blocked because: The target user's domain isn't included in the list of allowed domains. Or The target user's domain is included in the list of blocked domains. For more information about allowing or blocking external sharing based on domains, see Restricted domains sharing in SharePoint Online and OneDrive for Business.
Created access request	AccessRequestCreated	User requests access to a site, folder, or document they don't have permissions to access.
Created a company shareable link	CompanyLinkCreated	User created a company-wide link to a resource. company-wide links can only be used by members in your organization. They can't be used by guests.
Created an anonymous link	AnonymousLinkCreated	User created an anonymous link to a resource. Anyone with this link can access the resource without having to be authenticated.
Created secure link	SecureLinkCreated	A secure sharing link was created to this item.
Created sharing invitation	SharingInvitationCreated	User shared a resource in SharePoint Online or OneDrive for Business with a user who isn't in your organization's directory.
Deleted secure link	SecureLinkDeleted	A secure sharing link was deleted.
Denied access request	AccessRequestDenied	An access request to a site, folder, or document was denied.
Removed a company shareable link	CompanyLinkRemoved	User removed a company-wide link to a resource. The link can no longer be used to access the resource.
Removed an anonymous link	AnonymousLinkRemoved	User removed an anonymous link to a resource. The link can no longer be used to access the resource.

FRIENDLY NAME	OPERATION	DESCRIPTION
Shared file, folder, or site	SharingSet	User (member or guest) shared a file, folder, or site in SharePoint or OneDrive for Business with a user in your organization's directory. The value in the **Detail** column for this activity identifies the name of the user the resource was shared with and whether this user is a member or a guest. This activity is often accompanied by a second event that describes how the user was granted access to the resource. For example, adding the user to a group that has access to the resource.
Updated access request	AccessRequestUpdated	An access request to an item was updated.
Updated an anonymous link	AnonymousLinkUpdated	User updated an anonymous link to a resource. The updated field is included in the EventData property when you export the search results.
Updated sharing invitation	SharingInvitationUpdated	An external sharing invitation was updated.
Used an anonymous link	AnonymousLinkUsed	An anonymous user accessed a resource by using an anonymous link. The user's identity might be unknown, but you can get other details such as the user's IP address.
Unshared file, folder, or site	SharingRevoked	User (member or guest) unshared a file, folder, or site that was previously shared with another user.
Used a company shareable link	CompanyLinkUsed	User accessed a resource by using a company-wide link.
Used secure link	SecureLinkUsed	A user used a secure link.
User added to secure link	AddedToSecureLink	A user was added to the list of entities who can use a secure sharing link.
User removed from secure link	RemovedFromSecureLink	A user was removed from the list of entities who can use a secure sharing link.
Withdrew sharing invitation	SharingInvitationRevoked	User withdrew a sharing invitation to a resource.

Synchronization activities

The following table lists file synchronization activities in SharePoint Online and OneDrive for Business.

FRIENDLY NAME	OPERATION	DESCRIPTION
Allowed computer to sync files	ManagedSyncClientAllowed	User successfully establishes a sync relationship with a site. The sync relationship is successful because the user's computer is a member of a domain that's been added to the list of domains (called the *safe recipients list*) that can access document libraries in your organization. For more information about this feature, see Use Windows PowerShell cmdlets to enable OneDrive sync for domains that are on the safe recipients list.
Blocked computer from syncing files	UnmanagedSyncClientBlocked	User tries to establish a sync relationship with a site from a computer that isn't a member of your organization's domain or is a member of a domain that hasn't been added to the list of domains (called the *safe recipients list)* that can access document libraries in your organization. The sync relationship is not allowed, and the user's computer is blocked from syncing, downloading, or uploading files on a document library. For information about this feature, see Use Windows PowerShell cmdlets to enable OneDrive sync for domains that are on the safe recipients list.
Downloaded files to computer	FileSyncDownloadedFull	User establishes a sync relationship and successfully downloads files for the first time to their computer from a document library.
Downloaded file changes to computer	FileSyncDownloadedPartial	User successfully downloads any changes to files from a document library. This activity indicates that any changes that were made to files in the document library were downloaded to the user's computer. Only changes were downloaded because the document library was previously downloaded by the user (as indicated by the **Downloaded files to computer** activity).
Uploaded files to document library	FileSyncUploadedFull	User establishes a sync relationship and successfully uploads files for the first time from their computer to a document library.

FRIENDLY NAME	OPERATION	DESCRIPTION
Uploaded file changes to document library	FileSyncUploadedPartial	User successfully uploads changes to files on a document library. This event indicates that any changes made to the local version of a file from a document library are successfully uploaded to the document library. Only changes are uploaded because those files were previously uploaded by the user (as indicated by the **Uploaded files to document library** activity).

Site permissions activities

The following table lists events related to assigning permissions in SharePoint and using groups to give (and revoke) access to sites. As previously explained, audit records for some SharePoint activities will indicate the app@sharepoint user performed the activity of behalf of the user or admin who initiated the action. For more information, see The app@sharepoint user in audit records.

FRIENDLY NAME	OPERATION	DESCRIPTION
Added site collection admin	SiteCollectionAdminAdded	Site collection administrator or owner adds a person as a site collection administrator for a site. Site collection administrators have full control permissions for the site collection and all subsites. This activity is also logged when an admin gives themselves access to a user's OneDrive account (by editing the user profile in the SharePoint admin center or by using the Microsoft 365 admin center).
Added user or group to SharePoint group	AddedToGroup	User added a member or guest to a SharePoint group. This might have been an intentional action or the result of another activity, such as a sharing event.
Broke permission level inheritance	PermissionLevelsInheritanceBroken	An item was changed so that it no longer inherits permission levels from its parent.
Broke sharing inheritance	SharingInheritanceBroken	An item was changed so that it no longer inherits sharing permissions from its parent.
Created group	GroupAdded	Site administrator or owner creates a group for a site, or performs a task that results in a group being created. For example, the first time a user creates a link to share a file, a system group is added to the user's OneDrive for Business site. This event can also be a result of a user creating a link with edit permissions to a shared file.

FRIENDLY NAME	OPERATION	DESCRIPTION
Deleted group	GroupRemoved	User deletes a group from a site.
Modified access request setting	WebRequestAccessModified	The access request settings were modified on a site.
Modified 'Members Can Share' setting	WebMembersCanShareModified	The **Members Can Share** setting was modified on a site.
Modified permission level on a site collection	PermissionLevelModified	A permission level was changed on a site collection.
Modified site permissions	SitePermissionsModified	Site administrator or owner (or system account) changes the permission level that is assigned to a group on a site. This activity is also logged if all permissions are removed from a group. **NOTE**: This operation has been deprecated in SharePoint Online. To find related events, you can search for other permission-related activities such as **Added site collection admin**, **Added user or group to SharePoint group**, **Allowed user to create groups**, **Created group**, and **Deleted group.**
Removed permission level from site collection	PermissionLevelRemoved	A permission level was removed from a site collection.
Removed site collection admin	SiteCollectionAdminRemoved	Site collection administrator or owner removes a person as a site collection administrator for a site. This activity is also logged when an admin removes themselves from the list of site collection administrators for a user's OneDrive account (by editing the user profile in the SharePoint admin center). To return this activity in the audit log search results, you have to search for all activities.
Removed user or group from SharePoint group	RemovedFromGroup	User removed a member or guest from a SharePoint group. This might have been an intentional action or the result of another activity, such as an unsharing event.
Requested site admin permissions	SiteAdminChangeRequest	User requests to be added as a site collection administrator for a site collection. Site collection administrators have full control permissions for the site collection and all subsites.

FRIENDLY NAME	OPERATION	DESCRIPTION
Restored sharing inheritance	SharingInheritanceReset	A change was made so that an item inherits sharing permissions from its parent.
Updated group	GroupUpdated	Site administrator or owner changes the settings of a group for a site. This can include changing the group's name, who can view or edit the group membership, and how membership requests are handled.

Site administration activities

The following table lists events that result from site administration tasks in SharePoint Online. As previously explained, audit records for some SharePoint activities will indicate the app@sharepoint user performed the activity of behalf of the user or admin who initiated the action. For more information, see The app@sharepoint user in audit records.

FRIENDLY NAME	OPERATION	DESCRIPTION
Added allowed data location	AllowedDataLocationAdded	A SharePoint or global administrator added an allowed data location in a multi-geo environment.
Added exempt user agent	ExemptUserAgentSet	A SharePoint or global administrator added a user agent to the list of exempt user agents in the SharePoint admin center.
Added geo location admin	GeoAdminAdded	A SharePoint or global administrator added a user as a geo admin of a location.
Allowed user to create groups	AllowGroupCreationSet	Site administrator or owner adds a permission level to a site that allows a user assigned that permission to create a group for that site.
Canceled site geo move	SiteGeoMoveCancelled	A SharePoint or global administrator successfully cancels a SharePoint or OneDrive site geo move. The Multi-Geo capability lets an organization span multiple Microsoft datacenter geographies, which are called geos. For more information, see Multi-Geo Capabilities in OneDrive and SharePoint Online.

FRIENDLY NAME	OPERATION	DESCRIPTION
Changed a sharing policy	SharingPolicyChanged	A SharePoint or global administrator changed a SharePoint sharing policy by using the Microsoft 365 admin portal, SharePoint admin portal, or SharePoint Online Management Shell. Any change to the settings in the sharing policy in your organization will be logged. The policy that was changed is identified in the **ModifiedProperties** field in the detailed properties of the event record.
Changed device access policy	DeviceAccessPolicyChanged	A SharePoint or global administrator changed the unmanaged devices policy for your organization. This policy controls access to SharePoint, OneDrive, and Microsoft 365 from devices that aren't joined to your organization. Configuring this policy requires an Enterprise Mobility + Security subscription. For more information, see Control access from unmanaged devices.
Changed exempt user agents	CustomizeExemptUsers	A SharePoint or global administrator customized the list of exempt user agents in the SharePoint admin center. You can specify which user agents to exempt from receiving an entire web page to index. This means when a user agent you've specified as exempt encounters an InfoPath form, the form will be returned as an XML file, instead of an entire web page. This makes indexing InfoPath forms faster.
Changed network access policy	NetworkAccessPolicyChanged	A SharePoint or global administrator changed the location-based access policy (also called a trusted network boundary) in the SharePoint admin center or by using SharePoint Online PowerShell. This type of policy controls who can access SharePoint and OneDrive resources in your organization based on authorized IP address ranges that you specify. For more information, see Control access to SharePoint Online and OneDrive data based on network location.
Completed site geo move	SiteGeoMoveCompleted	A site geo move that was scheduled by a global administrator in your organization was successfully completed. The Multi-Geo capability lets an organization span multiple Microsoft datacenter geographies, which are called geos. For more information, see Multi-Geo Capabilities in OneDrive and SharePoint Online in Office 365.

FRIENDLY NAME	OPERATION	DESCRIPTION
Created Sent To connection	SendToConnectionAdded	A SharePoint or global administrator creates a new Send To connection on the Records management page in the SharePoint admin center. A Send To connection specifies settings for a document repository or a records center. When you create a Send To connection, a Content Organizer can submit documents to the specified location.
Created site collection	SiteCollectionCreated	A SharePoint or global administrator creates a site collection in your SharePoint Online organization or a user provisions their OneDrive for Business site.
Deleted orphaned hub site	HubSiteOrphanHubDeleted	A SharePoint or global administrator deleted an orphan hub site, which is a hub site that doesn't have any sites associated with it. An orphaned hub is likely caused by the deletion of the original hub site.
Deleted Sent To connection	SendToConnectionRemoved	A SharePoint or global administrator deletes a Send To connection on the Records management page in the SharePoint admin center.
Deleted site	SiteDeleted	Site administrator deletes a site.
Enabled document preview	PreviewModeEnabledSet	Site administrator enables document preview for a site.
Enabled legacy workflow	LegacyWorkflowEnabledSet	Site administrator or owner adds the SharePoint 2013 Workflow Task content type to the site. Global administrators can also enable work flows for the entire organization in the SharePoint admin center.
Enabled Office on Demand	OfficeOnDemandSet	Site administrator enables Office on Demand, which lets users access the latest version of Office desktop applications. Office on Demand is enabled in the SharePoint admin center and requires a Microsoft 365 subscription that includes full, installed Office applications.
Enabled result source for People Searches	PeopleResultsScopeSet	Site administrator creates the result source for People Searches for a site.

FRIENDLY NAME	OPERATION	DESCRIPTION
Enabled RSS feeds	NewsFeedEnabledSet	Site administrator or owner enables RSS feeds for a site. Global administrators can enable RSS feeds for the entire organization in the SharePoint admin center.
Joined site to hub site	HubSiteJoined	A site owner associates their site with a hub site.
Registered hub site	HubSiteRegistered	A SharePoint or global administrator creates a hub site. The results are that the site is registered to be a hub site.
Removed allowed data location	AllowedDataLocationDeleted	A SharePoint or global administrator removed an allowed data location in a multi-geo environment.
Removed geo location admin	GeoAdminDeleted	A SharePoint or global administrator removed a user as a geo admin of a location.
Renamed site	SiteRenamed	Site administrator or owner renames a site
Scheduled site geo move	SiteGeoMoveScheduled	A SharePoint or global administrator successfully schedules a SharePoint or OneDrive site geo move. The Multi-Geo capability lets an organization span multiple Microsoft datacenter geographies, which are called geos. For more information, see Multi-Geo Capabilities in OneDrive and SharePoint Online in Office 365.
Set host site	HostSiteSet	A SharePoint or global administrator changes the designated site to host personal or OneDrive for Business sites.
Set storage quota for geo location	GeoQuotaAllocated	A SharePoint or global administrator configured the storage quota for a geo location in a multi-geo environment.
Unjoined site from hub site	HubSiteUnjoined	A site owner disassociates their site from a hub site.
Unregistered hub site	HubSiteUnregistered	A SharePoint or global administrator unregisters a site as a hub site. When a hub site is unregistered, it no longer functions as a hub site.

Exchange mailbox activities

The following table lists the activities that can be logged by mailbox audit logging. Mailbox activities performed by the mailbox owner, a delegated user, or an administrator are automatically logged in the audit log for up to

90 days. It's possible for an admin to turn off mailbox audit logging for all users in your organization. In this case, no mailbox actions for any user are logged. For more information, see Manage mailbox auditing.

You can also search for mailbox activities by using the Search-MailboxAuditLog cmdlet in Exchange Online PowerShell.

FRIENDLY NAME	OPERATION	DESCRIPTION
Accessed mailbox items	MailItemsAccessed	Messages were read or accessed in mailbox. Audit records for this activity are triggered in one of two ways: when a mail client (such as Outlook) performs a bind operation on messages or when mail protocols (such as Exchange ActiveSync or IMAP) sync items in a mail folder. This activity is only logged for users with an Office 365 or Microsoft 365 E5 license. Analyzing audit records for this activity is useful when investigating compromised email account. For more information, see the "Access to crucial events for investigations" section in Advanced Audit.
Added delegate mailbox permissions	AddMailboxPermissions	An administrator assigned the FullAccess mailbox permission to a user (known as a delegate) to another person's mailbox. The FullAccess permission allows the delegate to open the other person's mailbox, and read and manage the contents of the mailbox.
Added or removed user with delegate access to calendar folder	UpdateCalendarDelegation	A user was added or removed as a delegate to the calendar of another user's mailbox. Calendar delegation gives someone else in the same organization permissions to manage the mailbox owner's calendar.
Added permissions to folder	AddFolderPermissions	A folder permission was added. Folder permissions control which users in your organization can access folders in a mailbox and the messages located in those folders.
Copied messages to another folder	Copy	A message was copied to another folder.
Created mailbox item	Create	An item is created in the Calendar, Contacts, Notes, or Tasks folder in the mailbox. For example, a new meeting request is created. Creating, sending, or receiving a message isn't audited. Also, creating a mailbox folder is not audited.

FRIENDLY NAME	OPERATION	DESCRIPTION
Created new inbox rule in Outlook web app	New-InboxRule	A mailbox owner or other user with access to the mailbox created an inbox rule in the Outlook web app.
Deleted messages from Deleted Items folder	SoftDelete	A message was permanently deleted or deleted from the Deleted Items folder. These items are moved to the Recoverable Items folder. Messages are also moved to the Recoverable Items folder when a user selects it and presses **Shift+Delete**.
Labeled message as a record	ApplyRecordLabel	A message was classified as a record. This occurs when a retention label that classifies content as a record is manually or automatically applied to a message.
Moved messages to another folder	Move	A message was moved to another folder.
Moved messages to Deleted Items folder	MoveToDeletedItems	A message was deleted and moved to the Deleted Items folder.
Modified folder permission	UpdateFolderPermissions	A folder permission was changed. Folder permissions control which users in your organization can access mailbox folders and the messages in the folder.
Modified inbox rule from Outlook web app	Set-InboxRule	A mailbox owner or other user with access to the mailbox modified an inbox rule using the Outlook web app.
Purged messages from the mailbox	HardDelete	A message was purged from the Recoverable Items folder (permanently deleted from the mailbox).
Removed delegate mailbox permissions	Remove-MailboxPermission	An administrator removed the FullAccess permission (that was assigned to a delegate) from a person's mailbox. After the FullAccess permission is removed, the delegate can't open the other person's mailbox or access any content in it.
Removed permissions from folder	RemoveFolderPermissions	A folder permission was removed. Folder permissions control which users in your organization can access folders in a mailbox and the messages located in those folders.

FRIENDLY NAME	OPERATION	DESCRIPTION
Sent message	Send	A message was sent, replied to or forwarded. This activity is only logged for users with an Office 365 or Microsoft 365 E5 license. For more information, see the "Access to crucial events for investigations" section in Advanced Audit.
Sent message using Send As permissions	SendAs	A message was sent using the SendAs permission. This means that another user sent the message as though it came from the mailbox owner.
Sent message using Send On Behalf permissions	SendOnBehalf	A message was sent using the SendOnBehalf permission. This means that another user sent the message on behalf of the mailbox owner. The message indicates to the recipient whom the message was sent on behalf of and who actually sent the message.
Updated inbox rules from Outlook client	UpdateInboxRules	A mailbox owner or other user with access to the mailbox modified an inbox rule in the Outlook client.
Updated message	Update	A message or its properties was changed.
User signed in to mailbox	MailboxLogin	The user signed in to their mailbox.
Label message as a record		A user applied a retention label to an email message and that label is configured to mark the item as a record.

User administration activities

The following table lists user administration activities that are logged when an admin adds or changes a user account by using the Microsoft 365 admin center or the Azure management portal.

ACTIVITY	OPERATION	DESCRIPTION
Added user	Add user	A user account was created.
Changed user license	Change user license	The license assigned to a user what changed. To see what licenses were changes, see the corresponding **Updated user** activity.

ACTIVITY	OPERATION	DESCRIPTION
Changed user password	Change user password	A user changes their password. Self-service password reset has to be enabled (for all or selected users) in your organization to allow users to reset their password. You can also track self-service password reset activity in Azure Active Directory. For more information, see Reporting options for Azure AD password management.
Deleted user	Delete user	A user account was deleted.
Reset user password	Reset user password	Administrator resets the password for a user.
Set property that forces user to change password	Set force change user password	Administrator set the property that forces a user to change their password the next time the user signs in to Office 365.
Set license properties	Set license properties	Administrator modifies the properties of a licensed assigned to a user.
Updated user	Update user	Administrator changes one or more properties of a user account. For a list of the user properties that can be updated, see the "Update user attributes" section in Azure Active Directory Audit Report Events.

Azure AD group administration activities

The following table lists group administration activities that are logged when an admin or a user creates or changes a Microsoft 365 group or when an admin creates a security group by using the Microsoft 365 admin center or the Azure management portal. For more information about groups in Office 365, see View, create, and delete Groups in the Microsoft 365 admin center.

FRIENDLY NAME	OPERATION	DESCRIPTION
Added group	Add group	A group was created.
Added member to group	Add member to group	A member was added to a group.
Deleted group	Delete group	A group was deleted.
Removed member from group	Remove member from group	A member was removed from a group.
Updated group	Update group	A property of a group was changed.

Application administration activities

The following table lists application admin activities that are logged when an admin adds or changes an

application that's registered in Azure AD. Any application that relies on Azure AD for authentication must be registered in the directory.

FRIENDLY NAME	OPERATION	DESCRIPTION
Added delegation entry	Add delegation entry	An authentication permission was created/granted to an application in Azure AD.
Added service principal	Add service principal	An application was registered in Azure AD. An application is represented by a service principal in the directory.
Added credentials to a service principal	Add service principal credentials	Credentials were added to a service principal in Azure AD. A service principle represents an application in the directory.
Removed delegation entry	Remove delegation entry	An authentication permission was removed from an application in Azure AD.
Removed a service principal from the directory	Remove service principal	An application was deleted/unregistered from Azure AD. An application is represented by a service principal in the directory.
Removed credentials from a service principal	Remove service principal credentials	Credentials were removed from a service principal in Azure AD. A service principle represents an application in the directory.
Set delegation entry	Set delegation entry	An authentication permission was updated for an application in Azure AD.

Role administration activities

The following table lists Azure AD role administration activities that are logged when an admin manages admin roles in the Microsoft 365 admin center or in the Azure management portal.

FRIENDLY NAME	OPERATION	DESCRIPTION
Add member to Role	Add role member to role	Added a user to an admin role in Microsoft 365.
Removed a user from a directory role	Remove role member from role	Removed a user to from an admin role in Microsoft 365.
Set company contact information	Set company contact information	Updated the company-level contact preferences for your organization. This includes email addresses for subscription-related email sent by Microsoft 365, and technical notifications about services.

Directory administration activities

The following table lists Azure AD directory and domain-related activities that are logged when an administrator manages their organization in the Microsoft 365 admin center or in the Azure management portal.

FRIENDLY NAME	OPERATION	DESCRIPTION
Added domain to company	Add domain to company	Added a domain to your organization.
Added a partner to the directory	Add partner to company	Added a partner (delegated administrator) to your organization.
Removed domain from company	Remove domain from company	Removed a domain from your organization.
Removed a partner from the directory	Remove partner from company	Removed a partner (delegated administrator) from your organization.
Set company information	Set company information	Updated the company information for your organization. This includes email addresses for subscription-related email sent by Microsoft 365, and technical notifications about Microsoft 365 services.
Set domain authentication	Set domain authentication	Changed the domain authentication setting for your organization.
Updated the federation settings for a domain	Set federation settings on domain	Changed the federation (external sharing) settings for your organization.
Set password policy	Set password policy	Changed the length and character constraints for user passwords in your organization.
Turned on Azure AD sync	Set DirSyncEnabled flag on company	Set the property that enables a directory for Azure AD Sync.
Updated domain	Update domain	Updated the settings of a domain in your organization.
Verified domain	Verify domain	Verified that your organization is the owner of a domain.
Verified email verified domain	Verify email verified domain	Used email verification to verify that your organization is the owner of a domain.

eDiscovery activities

Content Search and eDiscovery-related activities that are performed in the security and compliance center or by running the corresponding PowerShell cmdlets are logged in the audit log. This includes the following activities:

- Creating and managing eDiscovery cases
- Creating, starting, and editing Content Searches
- Performing Content Search actions, such as previewing, exporting, and deleting search results

- Configuring permissions filtering for Content Search
- Managing the eDiscovery Administrator role

For a list and detailed description of the eDiscovery activities that are logged, see Search for eDiscovery activities in the audit log.

> **NOTE**
>
> It takes up to 30 minutes for events that result from the activities listed under **eDiscovery activities** and **Advanced eDiscovery activities** in the **Activities** drop-down list to be displayed in the search results. Conversely, it takes up to 24 hours for the corresponding events from eDiscovery cmdlet activities to appear in the search results.

Advanced eDiscovery activities

You can also search the audit log for activities in Advanced eDiscovery. For a description of these activities, see the "Advanced eDiscovery activities" section in Search for eDiscovery activities in the audit log.

Power BI activities

You can search the audit log for activities in Power BI. For information about Power BI activities, see the "Activities audited by Power BI" section in Using auditing within your organization.

Audit logging for Power BI isn't enabled by default. To search for Power BI activities in the audit log, you have to enable auditing in the Power BI admin portal. For instructions, see the "Audit logs" section in Power BI admin portal.

Microsoft Workplace Analytics activities

Workplace Analytics provides insight into how groups collaborate across your organization. The following table lists activities performed by users that are assigned the Administrator role or the Analyst roles in Workplace Analytics. Users assigned the Analyst role have full access to all service features and use the product to do analysis. Users assigned the Administrator role can configure privacy settings and system defaults, and can prepare, upload, and verify organizational data in Workplace Analytics. For more information, see Workplace Analytics.

FRIENDLY NAME	OPERATION	DESCRIPTION
Accessed OData link	AccessedOdataLink	Analyst accessed the OData link for a query.
Canceled query	CanceledQuery	Analyst canceled a running query.
Created meeting exclusion	MeetingExclusionCreated	Analyst created a meeting exclusion rule.
Deleted result	DeletedResult	Analyst deleted a query result.
Downloaded report	DownloadedReport	Analyst downloaded a query result file.
Executed query	ExecutedQuery	Analyst ran a query.
Updated data access setting	UpdatedDataAccessSetting	Admin updated data access settings.
Updated privacy setting	UpdatedPrivacySetting	Admin updated privacy settings; for example, minimum group size.

FRIENDLY NAME	OPERATION	DESCRIPTION
Uploaded organization data	UploadedOrgData	Admin uploaded organizational data file.
Viewed Explore	ViewedExplore	Analyst viewed visualizations in one or more Explore page tabs.

Microsoft Teams activities

You can search the audit log for user and admin activities in Microsoft Teams. Teams is a chat-centered workspace in Office 365. It brings a team's conversations, meetings, files, and notes together into a single place. For descriptions of the Teams activities that are audited, see Search the audit log for events in Microsoft Teams.

Microsoft Teams Healthcare activities

If your organization is using the Patients application in Microsoft Teams, you can search the audit log for activities related to the using the Patients app. If your environment is configured to support Patients app, an additional activity group for these activities is available in the **Activities** picker list.

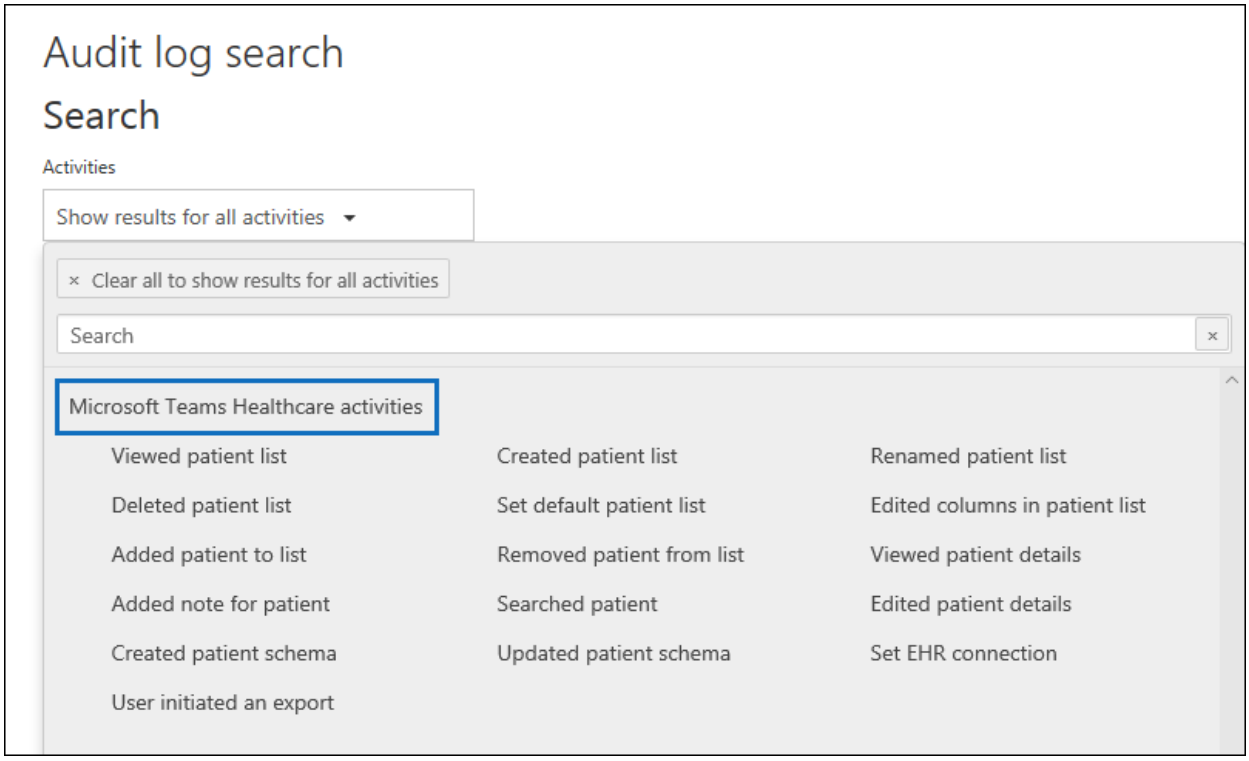

For a description of the Patients app activities, see Audit logs for Patients app.

Microsoft Teams Shifts activities

If your organization is using the Shifts app in Microsoft Teams, you can search the audit log for activities related to the using the Shifts app. If your environment is configured to support Shifts apps, an additional activity group for these activities is available in the **Activities** picker list.

For a description of Shifts app activities, see Search the audit log for events in Microsoft Teams.

Yammer activities

The following table lists the user and admin activities in Yammer that are logged in the audit log. To return Yammer-related activities from the audit log, you have to select **Show results for all activities** in the **Activities** list. Use the date range boxes and the **Users** list to narrow the search results.

FRIENDLY NAME	OPERATION	DESCRIPTION
Changed data retention policy	SoftDeleteSettingsUpdated	Verified admin updates the setting for the network data retention policy to either Hard Delete or Soft Delete. Only verified admins can perform this operation.
Changed network configuration	NetworkConfigurationUpdated	Network or verified admin changes the Yammer network's configuration. This includes setting the interval for exporting data and enabling chat.
Changed network profile settings	ProcessProfileFields	Network or verified admin changes the information that appears on member profiles for network users network.
Changed private content mode	SupervisorAdminToggled	Verified admin turns *Private Content Mode* on or off. This mode lets an admin view the posts in private groups and view private messages between individual users (or groups of users). Only verified admins only can perform this operation.
Changed security configuration	NetworkSecurityConfigurationUpdated	Verified admin updates the Yammer network's security configuration. This includes setting password expiration policies and restrictions on IP addresses. Only verified admins can perform this operation.
Created file	FileCreated	User uploads a file.
Created group	GroupCreation	User creates a group.
Deleted group	GroupDeletion	A group is deleted from Yammer.
Deleted message	MessageDeleted	User deletes a message.
Downloaded file	FileDownloaded	User downloads a file.
Exported data	DataExport	Verified admin exports Yammer network data. Only verified admins can perform this operation.
Shared file	FileShared	User shares a file with another user.
Suspended network user	NetworkUserSuspended	Network or verified admin suspends (deactivates) a user from Yammer.
Suspended user	UserSuspension	User account is suspended (deactivated).
Updated file description	FileUpdateDescription	User changes the description of a file.
Updated file name	FileUpdateName	User changes the name of a file.

FRIENDLY NAME	OPERATION	DESCRIPTION
Viewed file	FileVisited	User views a file.

Microsoft Power Automate activities

You can search the audit log for activities in Power Automate (formerly called Microsoft Flow). These activities include creating, editing, and deleting flows, and changing flow permissions. For information about auditing for Power Automate activities, see the blog Microsoft Flow audit events now available in Security & Compliance Center.

Microsoft Power Apps activities

You can search the audit log for app-related activities in Power Apps. These activities include creating, launching, and publishing an app. Assigning permissions to apps is also audited. For a description of all Power Apps activities, see Activity logging for Power Apps.

Microsoft Stream activities

You can search the audit log for activities in Microsoft Stream. These activities include video activities performed by users, group channel activities, and admin activities such as managing users, managing organization settings, and exporting reports. For a description of these activities, see the "Actions logged in Stream" section in Audit Logs in Microsoft Stream.

Content explorer activities

The following table lists the activities in content explorer that are logged in the audit log. Content explorer, which is accessed on the Data classifications tool in the Microsoft 365 compliance center. For more information, see Using data classification content explorer.

FRIENDLY NAME	OPERATION	DESCRIPTION
Accessed item	LabelContentExplorerAccessedItem	An admin (or a user who's a member of the Content Explorer Content Viewer role group) uses content explorer to view an email message or SharePoint/OneDrive document.

Quarantine activities

The following table lists the quarantine activities that you can search for in the audit log. For more information about quarantine, see Quarantine email messages in Office 365.

FRIENDLY NAME	OPERATION	DESCRIPTION
Deleted quarantine message	QuarantineDelete	A user deleted an email message that was deemed to be harmful.
Exported quarantine message	QuarantineExport	A user exported an email message that was deemed to be harmful.
Previewed quarantine message	QuarantinePreview	A user previewed an email message that was deemed to be harmful.

FRIENDLY NAME	OPERATION	DESCRIPTION
Released quarantine message	QuarantineRelease	A user released an email message from quarantine that was deemed to be harmful.
Viewed quarantine message's header	QuarantineViewHeader	A user viewed the header an email message that was deemed to be harmful.

Microsoft Forms activities

The following table lists the user and admin activities in Microsoft Forms that are logged in the audit log. Microsoft Forms is a forms/quiz/survey tool used to collect data for analysis.

Where noted below in the descriptions, some operations contain additional activity parameters.

> **NOTE**
> If a Forms activity is performed by a co-author or an anonymous responder, it will be logged slightly differently. For more information, see the Forms activities performed by co-authors and anonymous responders section.

FRIENDLY NAME	OPERATION	DESCRIPTION
Created comment	CreateComment	Form owner adds comment or score to a quiz.
Created form	CreateForm	Form owner creates a new form.
Edited form	EditForm	Form owner edits a form such, as creating, removing, or editing a question. The property *EditOperation:string* indicates the edit operation name. The possible operations are: - CreateQuestion - CreateQuestionChoice - DeleteQuestion - DeleteQuestionChoice - DeleteFormImage - DeleteQuestionImage - UpdateQuestion - UpdateQuestionChoice - UploadFormImage/Bing/Onedrive - UploadQuestionImage - ChangeTheme FormImage includes any place within Forms that user can upload an image, such as in a query or as a background theme.

FRIENDLY NAME	OPERATION	DESCRIPTION
Moved form	MoveForm	Form owner moves a form. Property DestinationUserId:string indicates the user ID of the person who moved the form. Property NewFormId:string is the new ID for the newly copied form.
Deleted form	DeleteForm	Form owner deletes a form. This includes SoftDelete (delete option used and form moved to recycle bin) and HardDelete (Recycle bin is emptied).
Viewed form (design time)	ViewForm	Form owner opens an existing form for editing.
Previewed form	PreviewForm	Form owner previews a form using the Preview function.
Exported form	ExportForm	Form owner exports results to Excel. Property ExportFormat:string indicates if the Excel file is Download or Online.
Allowed share form for copy	AllowShareFormForCopy	Form owner creates a template link to share the form with other users. This event is logged when the form owner clicks to generate template URL.
Disallowed share form for copy	DisallowShareFormForCopy	Form owner deletes template link.
Added form coauthor	AddFormCoauthor	A user uses a collaboration link to help design for/view responses. This event is logged when a user uses a collab URL (not when collab URL is first generated).
Removed form coauthor	RemoveFormCoauthor	Form owner deletes a collaboration link.
Viewed response page	ViewRuntimeForm	User has opened a response page to view. This event is logged regardless of whether the user submits a response or not.
Created response	CreateResponse	Similar to receiving a new response. A user has submitted a response to a form. Property ResponseId:string and Property ResponderId:string indicates which result is being viewed. For an anonymous responder, the ResponderId property will be null.

FRIENDLY NAME	OPERATION	DESCRIPTION
Updated response	UpdateResponse	Form owner has updated a comment or score on a quiz. Property ResponseId:string and Property ResponderId:string indicates which result is being viewed. For an anonymous responder, the ResponderId property will be null.
Deleted all responses	DeleteAllResponses	Form owner deletes all response data.
Deleted Response	DeleteResponse	Form owner deletes one response. Property ResponseId:string indicates the response being deleted.
Viewed responses	ViewResponses	Form owner views the aggregated list of responses. Property ViewType:string indicates whether form owner is viewing Detail or Aggregate
Viewed response	ViewResponse	Form owner views a particular response. Property ResponseId:string and Property ResponderId:string indicates which result is being viewed. For an anonymous responder, the ResponderId property will be null.
Created summary link	GetSummaryLink	Form owner creates summary results link to share results.
Deleted summary link	DeleteSummaryLink	Form owner deletes summary results link.
Updated form phishing status	UpdatePhishingStatus	This event is logged whenever the detailed value for the internal security status was changed, regardless of whether this changed the final security state (for example, form is now Closed or Opened). This means you may see duplicate events without a final security state change. The possible status values for this event are: - Take Down - Take Down by Admin - Admin Unblocked - Auto Blocked - Auto Unblocked - Customer Reported - Reset Customer Reported

FRIENDLY NAME	OPERATION	DESCRIPTION
Updated user phishing status	UpdateUserPhishingStatus	This event is logged whenever the value for the user security status was changed. The value of the user status in the audit record is **Confirmed as Phisher** when the user created a phishing form that was taken down by the Microsoft Online safety team. If an admin unblocks the user, the value of the user's status is set to **Reset as Normal User**.
Sent Forms Pro invitation	ProInvitation	User clicks to activate a Pro trial.
Updated form setting	UpdateFormSetting	Form owner updates a form setting. Property FormSettingName:string indicates the setting's name and new value.
Updated user setting	UpdateUserSetting	Form owner updates a user setting. Property UserSettingName:string indicates the setting's name and new value
Listed forms	ListForms	Form owner is viewing a list of forms. Property ViewType:string indicates which view the form owner is looking at: All Forms, Shared with Me, or Group Forms
Submitted response	SubmitResponse	A user submits a response to a form. Property IsInternalForm:boolean indicates if the responder is within the same organization as the form owner.

Forms activities performed by coauthors and anonymous responders

Forms supports collaboration when forms are designed and when analyzing responses. A form collaborator is known as a *coauthor*. Coauthors can do everything a form owner can do, except delete or move a form. Forms also allows you to create a form that can be responded to anonymously. This means the responder doesn't have to be signed into your organization to respond to a form.

The following table describes the auditing activities and information in the audit record for activities performed by coauthors and anonymous responders.

ACTIVITY TYPE	INTERNAL OR EXTERNAL USER	USER ID THAT'S LOGGED	ORGANIZATION LOGGED IN TO	FORMS USER TYPE
Coauthoring activities	Internal	UPN	Form owner's org	Coauthor
Coauthoring activities	External	UPN	Coauthor's org	Coauthor

ACTIVITY TYPE	INTERNAL OR EXTERNAL USER	USER ID THAT'S LOGGED	ORGANIZATION LOGGED IN TO	FORMS USER TYPE
Coauthoring activities	External	`urn:forms:coauthor#a0bfc2d3@mms.office.com` (The second part of the ID is a hash, which will differ for different users)	Form owner's org	Coauthor
Response activities	External	UPN	Responder's org	Responder
Response activities	External	`urn:forms:external#a0bfc2d3@mms.office.com` (The second part of the User ID is a hash, which will differ for different users)	Form owner's org	Responder
Response activities	Anonymous	`urn:forms:anonymous#a0bfc2d3@mms.office.com` (The second part of the User ID is a hash, which will differ for different users)	Form owner's org	Responder

Sensitivity label activities

The following table lists events that result from labeling activities for SharePoint Online and Teams sites.

FRIENDLY NAME	OPERATION	DESCRIPTION
Applied sensitivity label to site	SensitivityLabelApplied	A sensitivity label was applied to a SharePoint or Teams site.
Removed sensitivity label from site	SensitivityLabelRemoved	A sensitivity label was removed from a SharePoint or Teams site.
Applied sensitivity label to file	FileSensitivityLabelApplied	A sensitivity label was applied to a document by using Office on the web or an auto-labeling policy.
Changed sensitivity label applied to file	FileSensitivityLabelChanged	A different sensitivity label was applied to a document by using Office on the web or an auto-labeling policy.
Removed sensitivity label from file	FileSensitivityLabelRemoved	A sensitivity label was removed from a document by using Office on the web, an auto-labeling policy, or by using the Unlock-SPOSensitivityLabelEncryptedFile cmdlet.

Retention policy and retention label activities

FRIENDLY NAME	OPERATION	DESCRIPTION
Configured settings for a retention policy	NewRetentionComplianceRule	Administrator configured the retention settings for a new retention policy. Retention settings include how long items are retained, and what happens to items when the retention period expires (such as deleting items, retaining items, or retaining and then deleting them). This activity also corresponds to running the New-RetentionComplianceRule cmdlet.
Created retention label	NewComplianceTag	Administrator created a new retention label.
Created retention policy	NewRetentionCompliancePolicy	Administrator created a new retention policy.
Deleted settings from a retention policy	RemoveRetentionComplianceRule	Administrator deleted the configuration settings of a retention policy. Most likely, this activity is logged when an administrator deletes a retention policy or runs the Remove-RetentionComplianceRule cmdlet.
Deleted retention label	RemoveComplianceTag	Administrator deleted a retention label.
Deleted retention policy	RemoveRetentionCompliancePolicy	Administrator deleted a retention policy.
Enabled regulatory record option for retention labels	SetRestrictiveRetentionUI	Administrator ran the Set-RegulatoryComplianceUI cmdlet so that an administrator can then select the UI configuration option for a retention label to mark content as a regulatory record.
Updated settings for a retention policy	SetRetentionComplianceRule	Administrator changed the retention settings for an existing retention policy. Retention settings include how long items are retained, and what happens to items when the retention period expires (such as deleting items, retaining items, or retaining and then deleting them). This activity also corresponds to running the Set-RetentionComplianceRule cmdlet.
Updated retention label	SetComplianceTag	Administrator updated an existing retention label.
Updated retention policy	SetRetentionCompliancePolicy	Administrator updated an existing a retention policy. Updates that trigger this event include adding or excluding content locations that the retention policy is applied to.

Exchange admin audit log

Exchange administrator audit logging (which is enabled by default in Office 365) logs an event in the audit log when an administrator (or a user who has been assigned administrative permissions) makes a change in your Exchange Online organization. Changes made by using the Exchange admin center or by running a cmdlet in Exchange Online PowerShell are logged in the Exchange admin audit log. Cmdlets that begin with the verbs **Get-**, **Search-**, or **Test-** are not logged in the audit log. For more detailed information about admin audit logging in Exchange, see Administrator audit logging.

> **IMPORTANT**
>
> Some Exchange Online cmdlets that aren't logged in the Exchange admin audit log (or in the audit log). Many of these cmdlets are related to maintaining the Exchange Online service and are run by Microsoft datacenter personnel or service accounts. These cmdlets aren't logged because they would result in a large number of "noisy" auditing events. If there's an Exchange Online cmdlet that isn't being audited, please submit a suggestion to the Security & Compliance User Voice forum and request that it is enabled for auditing. You can also submit a design change request (DCR) to Microsoft Support.

Here are some tips for searching for Exchange admin activities when searching the audit log:

- To return entries from the Exchange admin audit log, you have to select **Show results for all activities** in the **Activities** list. Use the date range boxes and the **Users** list to narrow the search results for cmdlets run by a specific Exchange administrator within a specific date range.
- To display events from the Exchange admin audit log, filter the search results and type a - (dash) in the **Activity** filter box. This displays cmdlet names, which are displayed in the **Activity** column for Exchange admin events. Then you can sort the cmdlet names in alphabetical order.

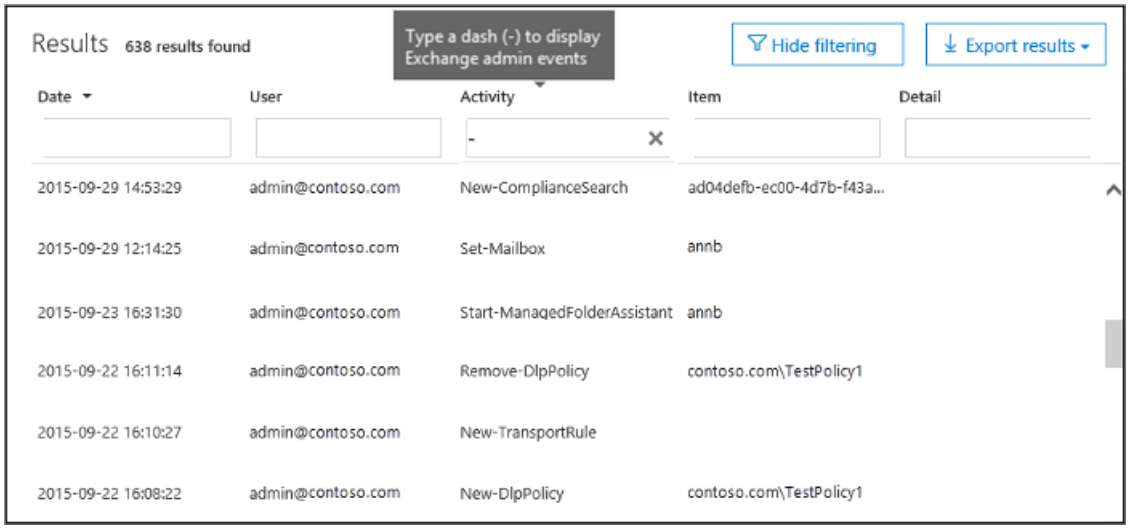

- To get information about what cmdlet was run, which parameters and parameter values were used, and what objects were affected, you can export the search results by selecting the **Download all results** option. For more information, see Export, configure, and view audit log records.
- You can also use the `Search-UnifiedAuditLog -RecordType ExchangeAdmin` command in Exchange Online PowerShell to return only audit records from the Exchange admin audit log. It may take up to 30 minutes after an Exchange cmdlet is run for the corresponding audit log entry to be returned in the search results. For more information, see Search-UnifiedAuditLog. For information about exporting the search results returned by the **Search-UnifiedAuditLog** cmdlet to a CSV file, see the "Tips for exporting and viewing the audit log" section in Export, configure, and view audit log records.
- You can also view events in the Exchange admin audit log by using the Exchange admin center or running the **Search-AdminAuditLog** in Exchange Online PowerShell. This is a good way to specifically search for activity performed by Exchange Online administrators. For instructions, see:

- View the administrator audit log
- Search-AdminAuditLog

Keep in mind that the same Exchange admin activities are logged in both the Exchange admin audit log and audit log.

Frequently asked questions

What are different Microsoft 365 services that are currently audited?

The most used services like Exchange Online, SharePoint Online, OneDrive for Business, Azure Active Directory, Microsoft Teams, Dynamics 365, Defender for Office 365, and Power BI are audited. See the beginning of this article for a list of services that are audited.

What activities are audited by auditing service in Office 365?

See the Audited activities section in this article for a list and description of the activities that are audited.

How long does it take for an auditing record to be available after an event has occurred?

Most auditing data is available within 30 minutes but it may take up to 24 hours after an event occurs for the corresponding audit log entry to be displayed in the search results. See the table in the Requirements to search the audit log section of this article that shows the time it takes for events in the different services to be available.

How long are the audit records retained for?

As previously explained, audit records for activities performed by users assigned an Office 365 E5 or Microsoft E5 license (or users with a Microsoft 365 E5 add-on license) are retained for one year. For all other subscriptions that support unified audit logging, audit records are retained for 90 days.

Can I access the auditing data programmatically?

Yes. The Office 365 Management Activity API is used to fetch the audit logs programmatically. To get started, see Get started with Office 365 Management APIs.

Are there other ways to get auditing logs other than using the security and compliance center or the Office 365 Management Activity API?

No. These are the only two ways to get data from the auditing service.

Do I need to individually enable auditing in each service that I want to capture audit logs for?

In most services, auditing is enabled by default after you initially turn on auditing for your organization (as described in the Requirements to search the audit log section in this article).

Does the auditing service support de-duplication of records?

No. The auditing service pipeline is near real time, and therefore can't support de-duplication.

Does auditing data flow across geographies?

No. We currently have auditing pipeline deployments in the NA (North America), EMEA (Europe, Middle East, and Africa) and APAC (Asia Pacific) regions. However, we may flow the data across these regions for load-balancing and only during live-site issues. When we do perform these activities, the data in transit is encrypted.

Is auditing data encrypted?

Auditing data is stored in Exchange mailboxes (data at rest) in the same region where the unified auditing pipeline is deployed. Mailbox data at rest is not encrypted by Exchange. However, service-level encryption encrypts all mailbox data because Exchange servers in Microsoft datacenters are encrypted via BitLocker. For

more information, see Office 365 Encryption for Skype for Business, OneDrive for Business, SharePoint Online, and Exchange Online.

Mail data in transit is always encrypted.

Use a PowerShell script to search the audit log

2/18/2021 • 8 minutes to read • Edit Online

Security, compliance, and auditing have become a top priority for IT administrators in today's world. Microsoft 365 has several built-in capabilities to help organizations manage security, compliance, and auditing. In particular, unified audit logging can help you investigate security incidents and compliance issues. You can retrieve audit logs by using the following methods:

- The Office 365 Management Activity API
- The audit log search tool in the Microsoft 365 compliance center
- The Search-UnifiedAuditLog cmdlet in Exchange Online PowerShell

If you need to retrieve audit logs on a regular basis, you should consider a solution that uses the Office 365 Management Activity API because it that can provide large organizations with the scalability and performance to retrieve millions of audit records on an ongoing basis. Using the audit log search tool in Microsoft 365 compliance center is a good way to quickly find audit records for specific operations that occur in shorter time range. Using longer time ranges in the audit log search tool, especially for large organizations, might return too many records to easily manage or export.

When there are situations where you need to manually retrieve auditing data for a specific investigation or incident, particularly for longer date ranges in larger organizations, using the **Search-UnifiedAuditLog** cmdlet may be the best option. This article includes a PowerShell script that uses the cmdlet to retrieve up to 50,000 audit records and then export them to a CSV file that you can format using Power Query in Excel to help with your review. Using the script in this article also minimizes the chance that large audit log searches will time out in the service.

Before you run the script

- Audit logging has to be enabled for your organization to successfully use the script to return audit records. Audit logging is turned on by default for Microsoft 365 and Office 365 enterprise organizations. To verify that audit log search is turned on for your organization, you can run the following command in Exchange Online PowerShell:

  ```
  Get-AdminAuditLogConfig | FL UnifiedAuditLogIngestionEnabled
  ```

  The value of `True` for the **UnifiedAuditLogIngestionEnabled** property indicates that audit log search is turned on.

- You have to be assigned the View-Only Audit Logs or Audit Logs role in Exchange Online to run successfully the script. By default, these roles are assigned to the Compliance Management and Organization Management role groups on the Permissions page in the Exchange admin center. For more information, see the "Requirements to search the audit log" section in Search the audit log in the compliance center.

- It may take a long time for the script to complete. How long it takes to run depends on the date range and the size of the interval that you configure the script to retrieve audit records for. Larger date ranges and smaller intervals will result in a long running time. See the table in Step 2 for more information about the date range and intervals.

- The sample script provided in this article isn't supported under any Microsoft standard support program

or service. The sample script is provided AS IS without warranty of any kind. Microsoft further disclaims all implied warranties including, without limitation, any implied warranties of merchantability or of fitness for a particular purpose. The entire risk arising out of the use or performance of the sample script and documentation remains with you. In no event shall Microsoft, its authors, or anyone else involved in the creation, production, or delivery of the script be liable for any damages whatsoever (including, without limitation, damages for loss of business profits, business interruption, loss of business information, or other pecuniary loss) arising out of the use of or inability to use the sample script or documentation, even if Microsoft has been advised of the possibility of such damages.

Step 1: Connect to Exchange Online PowerShell

The first step is to connect to Exchange Online PowerShell. You can connect using modern authentication or with multi-factor authentication (MFA). For step-by-step instructions, see Connect to Exchange Online PowerShell.

Step 2: Modify and run the script to retrieve audit records

After you've connected to Exchange Online PowerShell, the next step is to create, modify, and run the script to retrieve the auditing data. The first seven lines in the audit log search script contain the following variables that you can modify to configure your search. See the table in step 2 for a description of these variables.

1. Save the following text to a Windows PowerShell script by using a filename suffix of .ps1. For example, SearchAuditLog.ps1.

```
#Modify the values for the following variables to configure the audit log search.
$logFile = "d:\AuditLogSearch\AuditLogSearchLog.txt"
$outputFile = "d:\AuditLogSearch\AuditLogRecords.csv"
[DateTime]$start = [DateTime]::UtcNow.AddDays(-1)
[DateTime]$end = [DateTime]::UtcNow
$record = "AzureActiveDirectory"
$resultSize = 5000
$intervalMinutes = 60

#Start script
[DateTime]$currentStart = $start
[DateTime]$currentEnd = $start

Function Write-LogFile ([String]$Message)
{
    $final = [DateTime]::Now.ToUniversalTime().ToString("s") + ":" + $Message
    $final | Out-File $logFile -Append
}

Write-LogFile "BEGIN: Retrieving audit records between $($start) and $($end), RecordType=$record,
PageSize=$resultSize."
Write-Host "Retrieving audit records for the date range between $($start) and $($end), RecordType=$record,
ResultsSize=$resultSize"

$totalCount = 0
while ($true)
{
    $currentEnd = $currentStart.AddMinutes($intervalMinutes)
    if ($currentEnd -gt $end)
    {
        $currentEnd = $end
    }

    if ($currentStart -eq $currentEnd)
    {
        break
    }

    $sessionID = [Guid]::NewGuid().ToString() + "_" +  "ExtractLogs" + (Get-
```

```
Date).ToString("yyyyMMddHHmmssfff")
    Write-LogFile "INFO: Retrieving audit records for activities performed between $($currentStart) and
$($currentEnd)"
    Write-Host "Retrieving audit records for activities performed between $($currentStart) and
$($currentEnd)"
    $currentCount = 0

    $sw = [Diagnostics.StopWatch]::StartNew()
    do
    {
        $results = Search-UnifiedAuditLog -StartDate $currentStart -EndDate $currentEnd -RecordType $record
-SessionId $sessionID -SessionCommand ReturnLargeSet -ResultSize $resultSize

        if (($results | Measure-Object).Count -ne 0)
        {
            $results | export-csv -Path $outputFile -Append -NoTypeInformation

            $currentTotal = $results[0].ResultCount
            $totalCount += $results.Count
            $currentCount += $results.Count
            Write-LogFile "INFO: Retrieved $($currentCount) audit records out of the total $($currentTotal)"

            if ($currentTotal -eq $results[$results.Count - 1].ResultIndex)
            {
                $message = "INFO: Successfully retrieved $($currentTotal) audit records for the current time
range. Moving on!"
                Write-LogFile $message
                Write-Host "Successfully retrieved $($currentTotal) audit records for the current time
range. Moving on to the next interval." -foregroundColor Yellow
                ""
                break
            }
        }
    }
    while (($results | Measure-Object).Count -ne 0)

    $currentStart = $currentEnd
}

Write-LogFile "END: Retrieving audit records between $($start) and $($end), RecordType=$record,
PageSize=$resultSize, total count: $totalCount."
Write-Host "Script complete! Finished retrieving audit records for the date range between $($start) and
$($end). Total count: $totalCount" -foregroundColor Green
```

2. Modify the variables listed in the following table to configure the search criteria. The script includes sample values for these variables, but you should change them (unless stated otherwise) to meet your specific requirements.

VARIABLE	SAMPLE VALUE	DESCRIPTION
`$logFile`	"d:\temp\AuditSearchLog.txt"	Specifies the name and location for the log file that contains information about the progress of the audit log search performed by the script. The script writes UTC timestamps to the log file.
`$outputFile`	"d:\temp\AuditRecords.csv"	Specifies the name and location of the CSV file that contains the audit records returned by the script.

VARIABLE	SAMPLE VALUE	DESCRIPTION
`[DateTime]$start` and `[DateTime]$end`	[DateTime]::UtcNow.AddDays(-1) [DateTime]::UtcNow	Specifies the date range for the audit log search. The script will return records for audit activities that occurred within the specified date range. For example, to return activities performed in January 2021, you can use a start date of `"2021-01-01"` and an end date of `"2021-01-31"` (be sure to surround the values in double-quotation marks) The sample value in the script returns records for activities performed in the previous 24 hours. If you don't include a timestamp in the value, the default timestamp is 12:00 AM (midnight) on the specified date.
`$record`	"AzureActiveDirectory"	Specifies the record type of the audit activities (also called *operations*) to search for. This property indicates the service or feature that an activity was triggered in. For a list of record types that you can use for this variable, see Audit log record type. You can use the record type name or ENUM value. **Tip:** To return audit records for all record types, use the value `$null` (without double-quotations marks).
`$resultSize`	5000	Specifies the number of results returned each time the **Search-UnifiedAuditLog** cmdlet is called by the script (called a *result set*). The value of 5,000 is the maximum value supported by the cmdlet. Leave this value as-is.
`$intervalMinutes`	60	To help overcome the limit of 5000 records returned, this variable takes the data range you specified and slices it up into smaller time intervals. Now each interval, not the entire date range, is subject to the 5000 record output limit of the command. The default value of 5000 records per 60-minute interval within the date range should be sufficient for most organizations. But, if the script returns an error that says, `maximum results limitation reached`, decrease the time interval (for example, to 30 minutes or even 15 minutes) and rerun the script.

VARIABLE	SAMPLE VALUE	DESCRIPTION

Most of the variables listed in the previous table correspond to parameters for the **Search-UnifiedAuditLog** cmdlet. For more information about these parameters, see Search-UnifiedAuditLog.

3. On your local computer, open Windows PowerShell and go to the folder where you saved the modified script.
4. Run the script in Exchange Online PowerShell; for example:

```
.\SearchAuditLog.ps1
```

The script displays progress messages while it's running. After the script is finished running, it creates the log file and the CSV file that contains the audit records and saves them to the folders defined by the `$logFile` and `$outputFile` variables.

> **IMPORTANT**
> There is a 50,000 limit for the maximum number of audit records returned each time you run this script. If you run this script and it returns 50,000 results, then it's likely that audit records for activities that occurred within the date range weren't included. If this happens, we recommend that you divide the date range into smaller durations and then rerun the script for each date range. For example, if a date range of 90 days returns 50,000 results then you can rerun the script twice, once for the first 45 days in the date range and then again for the next 45 days.

Step 3: Format and view the audit records

After you've run the script and exported the audit records to a CSV file, you may want to format the CSV to make easier to review and analyze the audit records. One way to do this is to the Power Query JSON transform feature in Excel to split each property in the JSON object in the **AuditData** column into its own column. For step-by-step instructions, see "Step 2: Format the exported audit log using the Power Query Editor" in Export, configure, and view audit log records.

Turn audit log search on or off

2/18/2021 • 3 minutes to read • Edit Online

Audit logging is turned on by default for Microsoft 365 and Office 365 enterprise organizations. This includes organizations with E3/G3 or E5/G5 subscriptions. When audit log search in the compliance center is turned on, user and admin activity from your organization is recorded in the audit log and retained for 90 days, and up to one year depending on the license assigned to users. However, your organization may have reasons for not wanting to record and retain audit log data. In those cases, a global admin may decide to turn off auditing in Microsoft 365.

> **IMPORTANT**
> If you turn off audit log search in Microsoft 365, you can't use the Office 365 Management Activity API or Azure Sentinel to access auditing data for your organization. Turning off audit log search by following the steps in this article means that no results will be returned when you search the audit log using the Security & Compliance Center or when you run the **Search-UnifiedAuditLog** cmdlet in Exchange Online PowerShell. This also means that audit logs won't be available through the Office 365 Management Activity API or Azure Sentinel.

Before you turn audit log search on or off

- You have to be assigned the Audit Logs role in Exchange Online to turn audit log search on or off in your Microsoft 365 organization. By default, this role is assigned to the Compliance Management and Organization Management role groups on the **Permissions** page in the Exchange admin center. Global admins in Microsoft 365 are members of the Organization Management role group in Exchange Online.

  > **NOTE**
  > Users have to be assigned permissions in Exchange Online to turn audit log search on or off. If you assign users the Audit Logs role on the **Permissions** page in the Security & Compliance Center, they won't be able to turn audit log search on or off. This is because the underlying cmdlet is an Exchange Online PowerShell cmdlet.

- For step-by-step instructions on searching the audit log, see Search the audit log in the Security & Compliance Center. For more information about the Microsoft 365 Management Activity API, see Get started with Microsoft 365 Management APIs.

- To verify that audit log search is turned on, you can run the following command in Exchange Online PowerShell:

  ```
  Get-AdminAuditLogConfig | FL UnifiedAuditLogIngestionEnabled
  ```

  The value of `True` for the *UnifiedAuditLogIngestionEnabled* property indicates that audit log search is turned on.

Turn on audit log search

If audit log search is not turned on for your organization, you can turn it on in the compliance center or by using Exchange Online PowerShell. It may take several hours after you turn on audit log search before you can return results when you search the audit log.

Use the compliance center to turn on audit log search

1. Go to the compliance center and sign in.
2. In the compliance center, go to **Search** > **Audit log search**.

   If audit log search is not turned on for your organization, a banner is displayed saying that auditing has to be turned on to record user and admin activity.
3. Click **Turn on auditing**.

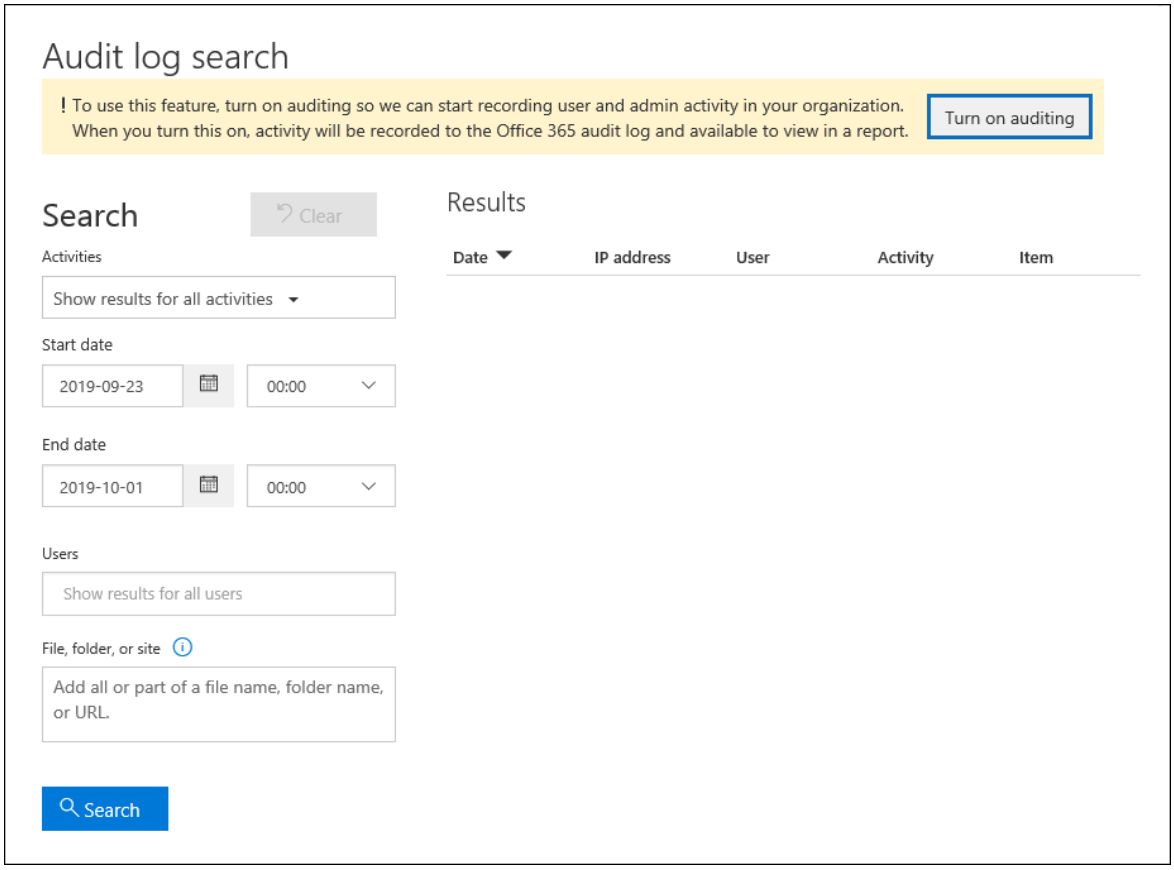

   The banner is updated to say the audit log is being prepared and that you can search for user and admin activity in a few hours.

Use PowerShell to turn on audit log search

1. Connect to Exchange Online PowerShell
2. Run the following PowerShell command to turn on audit log search in Office 365.

```
Set-AdminAuditLogConfig -UnifiedAuditLogIngestionEnabled $true
```

   A message is displayed saying that it may take up to 60 minutes for the change to take effect.

Turn off audit log search

You have to use Exchange Online PowerShell to turn off audit log search.

1. Connect to Exchange Online PowerShell
2. Run the following PowerShell command to turn off audit log search.

```
Set-AdminAuditLogConfig -UnifiedAuditLogIngestionEnabled $false
```

3. After a while, verify that audit log search is turned off (disabled). There are two ways to do this:

   - In Exchange Online PowerShell, run the following command:

     ```
     Get-AdminAuditLogConfig | FL UnifiedAuditLogIngestionEnabled
     ```

     The value of `False` for the *UnifiedAuditLogIngestionEnabled* property indicates that audit log search is turned off.

   - In the compliance center, go to **Search** > **Audit log search**.

     A banner is displayed saying that auditing has to be turned on in order to record user and admin activity.

Detailed properties in the audit log

2/18/2021 • 10 minutes to read • Edit Online

When you export the results of an audit log search from the Security & Compliance Center, you have the option to download all the results that meet your search criteria. You do this by selecting **Export results** > **Download all results** on the **Audit log search** page. For more information, see Search the audit log.

When your export all results for an audit log search, the raw data from the unified audit log is copied to a comma-separated value (CSV) file that is downloaded to your local computer. This file contains additional information from each audit record in a column named **AuditData**. This column contains a multi-value property for multiple properties from the audit log record. Each of the **property: value** pairs in this multi-value property are separated by a comma.

The following table describes the properties that are included (depending on the service in which an event occurs) in the multi-property **AuditData** column. The **Office 365 service that has this property** column indicates the service and type of activity (user or admin) that includes the property. For more detailed information about these properties or about properties that may not be listed in this topic, see Management Activity API Schema.

TIP

You can use the JSON transform feature in Power Query in Excel to split the **AuditData** column into multiple columns so that each property has its own column. This lets you sort and filter on one or more of these properties. To learn how to do this, see Export, configure, and view audit log records.

PROPERTY	DESCRIPTION	MICROSOFT 365 SERVICE THAT HAS THIS PROPERTY
Actor	The user or service account that performed the action.	Azure Active Directory
AddOnName	The name of an add-on that was added, removed, or updated in a team. The type of add-ons in Microsoft Teams is a bot, a connector, or a tab.	Microsoft Teams
AddOnType	The type of an add-on that was added, removed, or updated in a team. The following values indicate the type of add-on. **1** - Indicates a bot. **2** - Indicates a connector. **3** - Indicates a tab.	Microsoft Teams
AzureActiveDirectoryEventType	The type of Azure Active Directory event. The following values indicate the type of event. **0** - Indicates an account login event. **1** - Indicates an Azure application security event.	Azure Active Directory

PROPERTY	DESCRIPTION	MICROSOFT 365 SERVICE THAT HAS THIS PROPERTY
ChannelGuid	The ID of a Microsoft Teams channel. The team that the channel is located in is identified by the **TeamName** and **TeamGuid** properties.	Microsoft Teams
ChannelName	The name of a Microsoft Teams channel. The team that the channel is located in is identified by the **TeamName** and **TeamGuid** properties.	Microsoft Teams
Client	The client device, the device OS, and the device browser used for the login event (for example, Nokia Lumia 920; Windows Phone 8; IE Mobile 11).	Azure Active Directory
ClientInfoString	Information about the email client that was used to perform the operation, such as a browser version, Outlook version, and mobile device information	Exchange (mailbox activity)
ClientIP	The IP address of the device that was used when the activity was logged. The IP address is displayed in either an IPv4 or IPv6 address format. For some services, the value displayed in this property might be the IP address for a trusted application (for example, Office on the web apps) calling into the service on behalf of a user and not the IP address of the device used by person who performed the activity. Also, for admin activity (or activity performed by a system account) for Azure Active Directory-related events, the IP address isn't logged and the value for the ClientIP property is `null`.	Azure Active Directory, Exchange, SharePoint
CreationTime	The date and time in Coordinated Universal Time (UTC) when the user performed the activity.	All
DestinationFileExtension	The file extension of a file that is copied or moved. This property is displayed only for the FileCopied and FileMoved user activities.	SharePoint
DestinationFileName	The name of the file is copied or moved. This property is displayed only for the FileCopied and FileMoved actions.	SharePoint

PROPERTY	DESCRIPTION	MICROSOFT 365 SERVICE THAT HAS THIS PROPERTY
DestinationRelativeUrl	The URL of the destination folder where a file is copied or moved. The combination of the values for the **SiteURL**, the **DestinationRelativeURL**, and the **DestinationFileName** property is the same as the value for the **ObjectID** property, which is the full path name for the file that was copied. This property is displayed only for the FileCopied and FileMoved user activities.	SharePoint
EventSource	Identifies that an event occurred in SharePoint. Possible values are **SharePoint** and **ObjectModel**.	SharePoint
ExternalAccess	For Exchange admin activity, specifies whether the cmdlet was run by a user in your organization, by Microsoft datacenter personnel or a datacenter service account, or by a delegated administrator. The value **False** indicates that the cmdlet was run by someone in your organization. The value **True** indicates that the cmdlet was run by datacenter personnel, a datacenter service account, or a delegated administrator. For Exchange mailbox activity, specifies whether a mailbox was accessed by a user outside your organization.	Exchange
ExtendedProperties	The extended properties for an Azure Active Directory event.	Azure Active Directory
ID	The ID of the report entry. The ID uniquely identifies the report entry.	All
InternalLogonType	Reserved for internal use.	Exchange (mailbox activity)
ItemType	The type of object that was accessed or modified. Possible values include **File**, **Folder**, **Web**, **Site**, **Tenant**, and **DocumentLibrary**.	SharePoint
LoginStatus	Identifies login failures that might have occurred.	Azure Active Directory

PROPERTY	DESCRIPTION	MICROSOFT 365 SERVICE THAT HAS THIS PROPERTY
LogonType	The type of mailbox access. The following values indicate the type of user who accessed the mailbox. 0 - Indicates a mailbox owner. 1 - Indicates an administrator. 2 - Indicates a delegate. 3 - Indicates the transport service in the Microsoft datacenter. 4 - Indicates a service account in the Microsoft datacenter. 6 - Indicates a delegated administrator.	Exchange (mailbox activity)
MailboxGuid	The Exchange GUID of the mailbox that was accessed.	Exchange (mailbox activity)
MailboxOwnerUPN	The email address of the person who owns the mailbox that was accessed.	Exchange (mailbox activity)
Members	Lists the users that have been added or removed from a team. The following values indicate the Role type assigned to the user. 1 - Indicates the Owner role. 2 - Indicates the Member role. 3 - Indicates the Guest role. The Members property also includes the name of your organization, and the member's email address.	Microsoft Teams
ModifiedProperties (Name, NewValue, OldValue)	The property is included for admin events, such as adding a user as a member of a site or a site collection admin group. The property includes the name of the property that was modified (for example, the Site Admin group) the new value of the modified property (such the user who was added as a site admin, and the previous value of the modified object.	All (admin activity)
ObjectId	For Exchange admin audit logging, the name of the object that was modified by the cmdlet. For SharePoint activity, the full URL path name of the file or folder accessed by a user. For Azure AD activity, the name of the user account that was modified.	All

PROPERTY	DESCRIPTION	MICROSOFT 365 SERVICE THAT HAS THIS PROPERTY
Operation	The name of the user or admin activity. The value of this property corresponds to the value that was selected in the **Activities** drop down list. If **Show results for all activities** was selected, the report will included entries for all user and admin activities for all services. For a description of the operations/activities that are logged in the audit log, see the **Audited activities** tab in Search the audit log in the Office 365. For Exchange admin activity, this property identifies the name of the cmdlet that was run.	All
OrganizationId	The GUID for your organization.	All
Path	The name of the mailbox folder where the message that was accessed is located. This property also identifies the folder a where a message is created in or copied/moved to.	Exchange (mailbox activity)
Parameters	For Exchange admin activity, the name and value for all parameters that were used with the cmdlet that is identified in the Operation property.	Exchange (admin activity)
RecordType	The type of operation indicated by the record. This property indicates the service or feature that the operation was triggered in. For a list of record types and their corresponding ENUM value (which is the value displayed in the **RecordType** property in an audit record), see Audit log record type.	
ResultStatus	Indicates whether the action (specified in the **Operation** property) was successful or not. For Exchange admin activity, the value is either **True** (successful) or **False** (failed).	All
SecurityComplianceCenterEventType	Indicates that the activity was a Security & Compliance Center event. All Security & Compliance Center activities will have a value of **0** for this property.	Security & Compliance Center
SharingType	The type of sharing permissions that was assigned to the user that the resource was shared with. This user is identified in the **UserSharedWith** property.	SharePoint

PROPERTY	DESCRIPTION	MICROSOFT 365 SERVICE THAT HAS THIS PROPERTY
Site	The GUID of the site where the file or folder accessed by the user is located.	SharePoint
SiteUrl	The URL of the site where the file or folder accessed by the user is located.	SharePoint
SourceFileExtension	The file extension of the file that was accessed by the user. This property is blank if the object that was accessed is a folder.	SharePoint
SourceFileName	The name of the file or folder accessed by the user.	SharePoint
SourceRelativeUrl	The URL of the folder that contains the file accessed by the user. The combination of the values for the **SiteURL**, the **SourceRelativeURL**, and the **SourceFileName** property is the same as the value for the **ObjectID** property, which is the full path name for the file accessed by the user.	SharePoint
Subject	The subject line of the message that was accessed.	Exchange (mailbox activity)
TabType	The type of tab added, removed, or updated in a team. The possible values for this property are: **Excel pin** - An Excel tab. **Extension** - All first-party and third-party apps; such as Class Schedule, VSTS, and Forms. **Notes** - OneNote tab. **Pdfpin** - A PDF tab. **Powerbi** - A Power BI tab. **Powerpointpin** - A PowerPoint tab. **Sharepointfiles** - A SharePoint tab. **Webpage** - A pinned website tab. **Wiki-tab** - A wiki tab. **Wordpin** - A Word tab.	Microsoft Teams
Target	The user that the action (identified in the **Operation** property) was performed on. For example, if a guest user is added to SharePoint or a Microsoft Team, that user would be listed in this property.	Azure Active Directory
TeamGuid	The ID of a team in Microsoft Teams.	Microsoft Teams
TeamName	The name of a team in Microsoft Teams.	Microsoft Teams

PROPERTY	DESCRIPTION	MICROSOFT 365 SERVICE THAT HAS THIS PROPERTY
UserAgent	Information about the user's browser. This information is provided by the browser.	SharePoint
UserDomain	Identity information about the tenant organization of the user (actor) who performed the action.	Azure Active Directory
UserId	The user who performed the action (specified in the **Operation** property) that resulted in the record being logged. Audit records for activity performed by system accounts (such as SHAREPOINT\system or NT AUTHORITY\SYSTEM) are also included in the audit log. Another common value for the UserId property is app@sharepoint. This indicates that the "user" who performed the activity was an application that has the necessary permissions in SharePoint to perform organization-wide actions (such as search a SharePoint site or OneDrive account) on behalf of a user, admin, or service. For more information, see The app@sharepoint user in audit records.	All
UserKey	An alternative ID for the user identified in the **UserID** property. For example, this property is populated with the passport unique ID (PUID) for events performed by users in SharePoint. This property also might specify the same value as the **UserID** property for events occurring in other services and events performed by system accounts.	All
UserSharedWith	The user that a resource was shared with. This property is included if the value for the **Operation** property is **SharingSet**. This user is also listed in the **Shared with** column in the report.	SharePoint

PROPERTY	DESCRIPTION	MICROSOFT 365 SERVICE THAT HAS THIS PROPERTY
UserType	The type of user that performed the operation. The following values indicate the user type. **0** - A regular user. **2** - An administrator in your Microsoft 365 organization.[1] **3** - A Microsoft datacenter administrator or datacenter system account. **4** - A system account. **5** - An application. **6** - A service principal. **7** - A custom policy. **8** - A system policy.	All
Version	Indicates the version number of the activity (identified by the **Operation** property) that's logged.	All
Workload	The Microsoft 365 service where the activity occurred.	All

NOTE

[1] For Azure Active Directory-related events, the value for an administrator isn't used in an audit record. Audit records for activities performed by administrators will indicate that a regular user (for example, **UserType**: **0**) performed the activity. The **UserID** property will identify the person (regular user or administrator) who performed the activity.

The properties described above are also displayed when you click **More information** when viewing the details of a specific event.

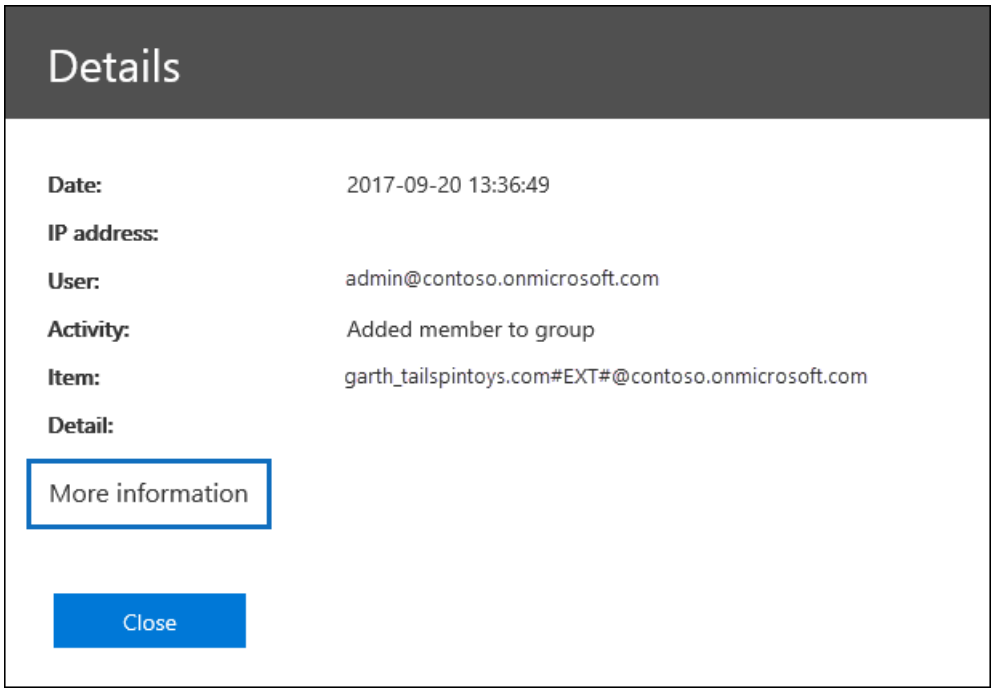

Export, configure, and view audit log records

11/2/2020 • 6 minutes to read • Edit Online

After you search the audit log and download the search results to a CSV file, the file contains a column named **AuditData**, which contains additional information about each event. The data in this column is formatted as a JSON object, which contains multiple properties that are configured as *property:value* pairs separated by commas. You can use the JSON transform feature in the Power Query Editor in Excel to split each property in the JSON object in the **AuditData** column into multiple columns so that each property has its own column. This lets you sort and filter on one or more of these properties, which can help you quickly locate the specific auditing data you're looking for.

Step 1: Export audit log search results

The first step is to search the audit log and then export the results in a comma-separated value (CSV) file to your local computer.

1. Run an audit log search and revise the search criteria if necessary until you have the desired results.
2. Click **Export results** and select **Download all results**.

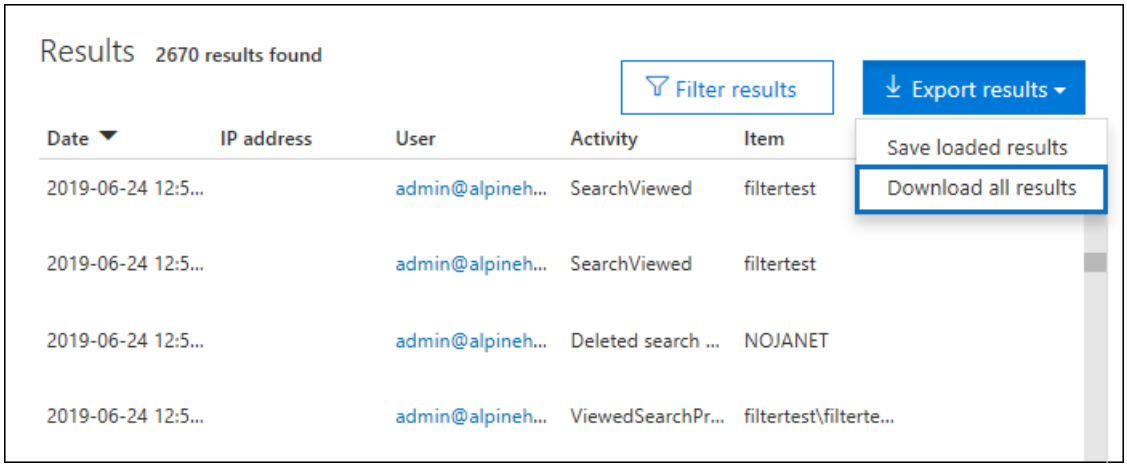

   This option to exports all the audit records from the audit log search you ran in step 1, and downloads the raw data from the audit log to a CSV file.

   A message is displayed at the bottom of the window that prompts you to open or save the CSV file.

3. Click **Save** > **Save as** and save the CSV file to your local computer. It takes a while to download many search results. This is typically the case when searching for all activities or a broad date range. A message at the bottom of the windows is displayed when the CSV file is finished downloading.

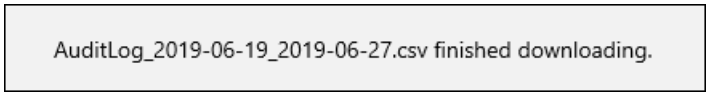

NOTE

You can download a maximum of 50,000 entries to a CSV file from a single audit log search. If 50,000 entries are downloaded to the CSV file, you can probably assume there are more than 50,000 events that met the search criteria. To export more than this limit, try using a date range to reduce the number of audit log records. You might have to run multiple searches with smaller date ranges to export more than 50,000 entries.

Step 2: Format the exported audit log using the Power Query Editor

The next step is to use the JSON transform feature in the Power Query Editor in Excel to split each property in the JSON object in the **AuditData** column into its own column. Then you filter columns to view records based on the values of specific properties. This can help you quickly locate the specific auditing data you're looking for.

1. Open a blank workbook in Excel for Office 365, Excel 2019, or Excel 2016.
2. On the **Data** tab, in the **Get & Transform Data** ribbon group, click **From Text/CSV**.

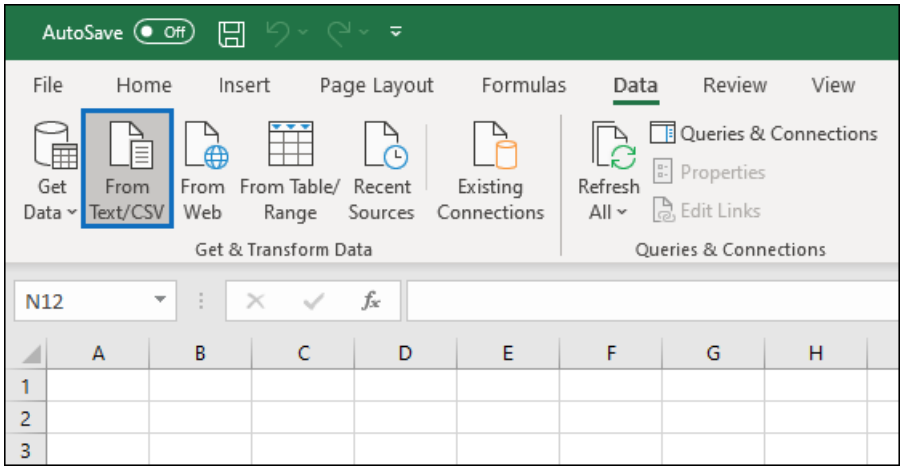

3. Open the CSV file that you downloaded in Step 1.
4. In the window that's displayed, click **Transform Data**.

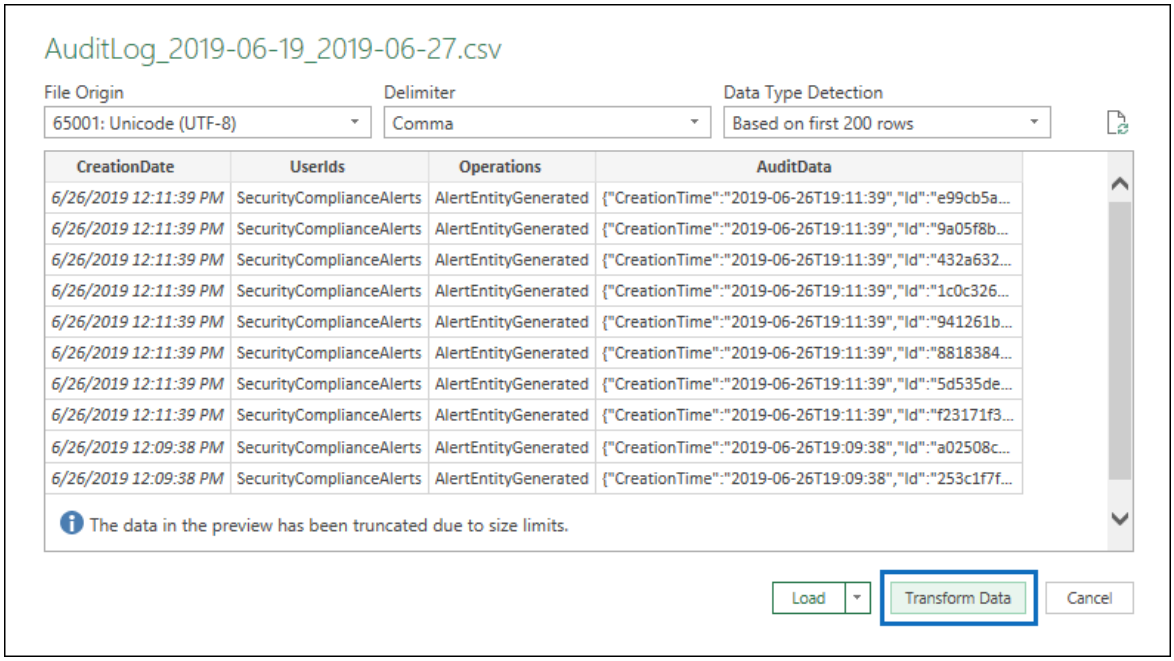

   The CSV file is opened in the **Query Editor**. There are four columns: **CreationDate**, **UserIds**, **Operations**, and **AuditData**. The **AuditData** column is a JSON object that contains multiple properties. The next step is to create a column for each property in the JSON object.

5. Right-click the title in the **AuditData** column, click **Transform**, and then click **JSON**.

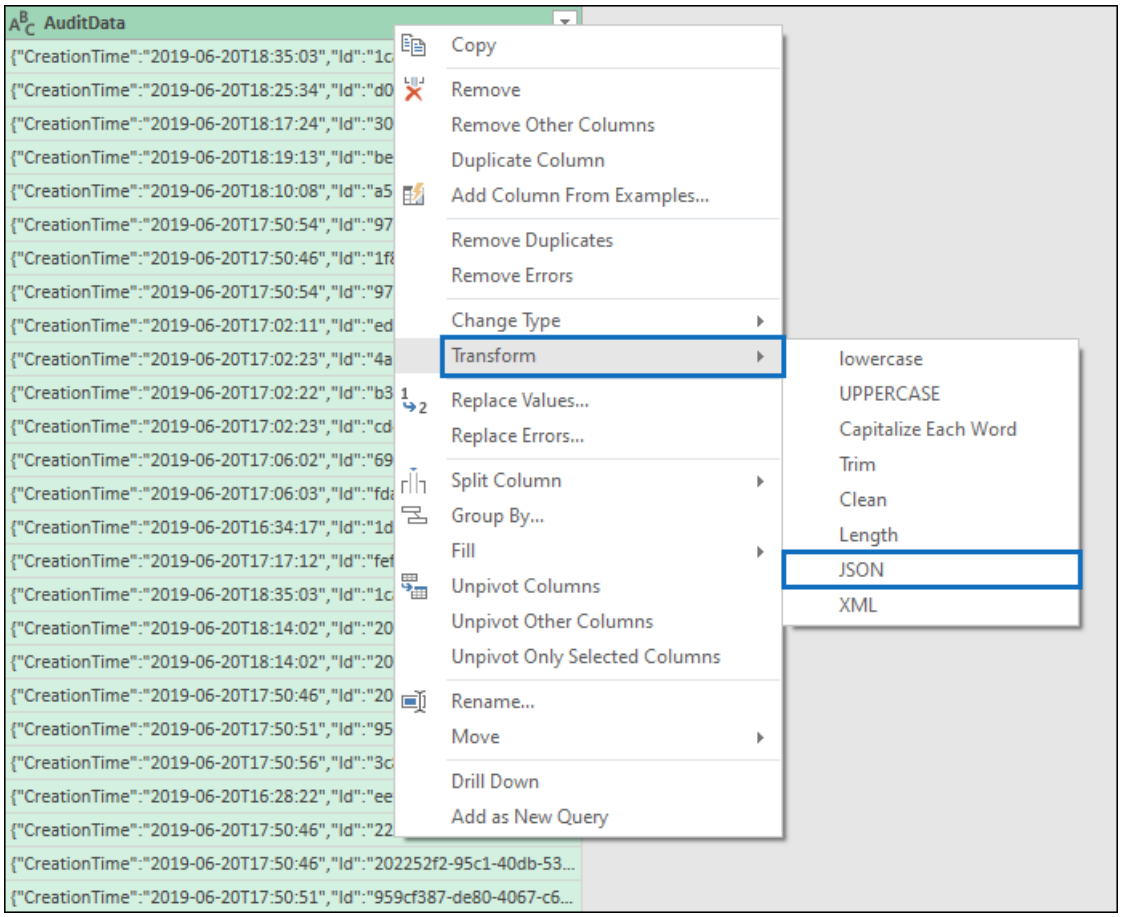

6. In the upper-right corner of the **AuditData** column, click the expand icon.

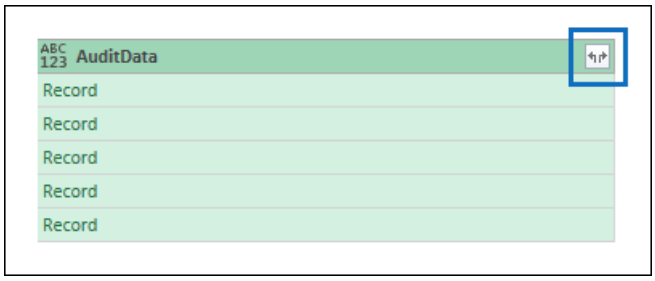

A partial list of the properties in the JSON objects in the **AuditData** column is displayed.

7. Click **Load more** to display all properties in the JSON objects in the **AuditData** column.

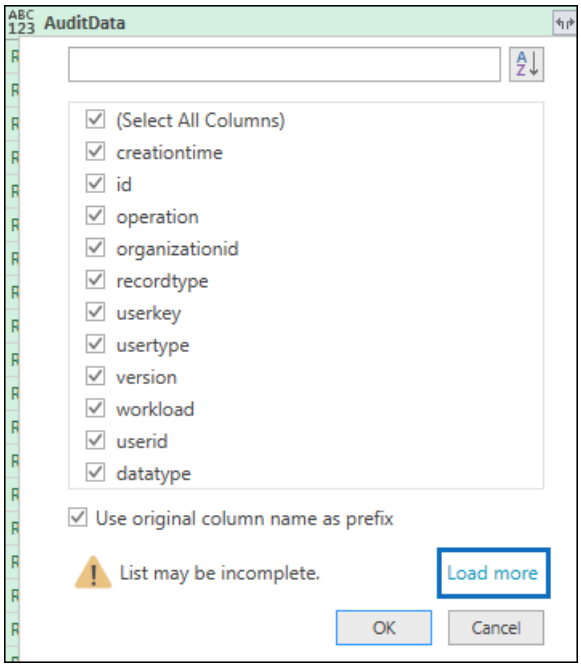

You can unselect the checkbox next to any property that you don't want to include. Eliminating columns that aren't useful for your investigation is a good way to reduce the amount of data displayed in the audit log.

> **NOTE**
>
> The JSON properties displayed in the previous screenshot (after you click **Load more**) are based on the properties found in the **AuditData** column from the first 1,000 rows in the CSV file. If there are different JSON properties in records after the first 1,000 rows, these properties (and a corresponding column) won't be included when the **AuditData** column is split into multiple columns. To help prevent this, consider re-running the audit log search and narrow the search criteria so that fewer records are returned. Another workaround is to filter items in the **Operations** column to reduce the number of rows (before you perform step 5 above) before transforming the JSON object in the **AuditData** column.

8. Do one of the following things to format the title of the columns that are added for each JSON property that's selected.

   - Unselect the **Use original column name as prefix** checkbox to use the name of the JSON property as the column names; for example, **RecordType** or **SourceFileName**.
   - Leave the **Use original column name as prefix** checkbox selected to add the AuditData prefix to the column names; for example, **AuditData.RecordType** or **AuditData.SourceFileName**.

9. Click **OK**.

   The **AuditData** column is split into multiple columns. Each new column corresponds to a property in the AuditData JSON object. Each row in the column contains the value for the property. If the property doesn't contain a value, the *null* value is displayed. In Excel, cells with null values are empty.

10. On the **Home** tab, click **Close & Load** to close the Power Query Editor and open the transformed CSV file in an Excel workbook.

Use PowerShell to search and export audit log records

Instead of using the audit log search tool in the Security & Compliance Center, you can use the Search-UnifiedAuditLog cmdlet in Exchange Online PowerShell to export the results of an audit log search to a CSV file. Then you can follow the same procedure described in Step 2 to format the audit log using the Power Query editor. One advantage of using the PowerShell cmdlet is that you can search for events from a specific service by using the *RecordType* parameter. Here are few examples of using PowerShell to export audit records to a CSV file so you can use the Power Query editor to transform the JSON object in the **AuditData** column as described in Step 2.

In this example, run the following commands to return all records related to SharePoint sharing operations.

```
$auditlog = Search-UnifiedAuditLog -StartDate 06/01/2019 -EndDate 06/30/2019 -RecordType
SharePointSharingOperation
```

```
$auditlog | Select-Object -Property CreationDate,UserIds,RecordType,AuditData | Export-Csv -Path
c:\AuditLogs\PowerShellAuditlog.csv -NoTypeInformation
```

The search results are exported to a CSV file named *PowerShellAuditlog* that contains four columns: CreationDate, UserIds, RecordType, AuditData).

You can also use the name or enum value for the record type as the value for the *RecordType* parameter. For a

list of record type names and their corresponding enum values, see the *AuditLogRecordType* table in Office 365 Management Activity API schema.

You can only include a single value for the *RecordType* parameter. To search for audit records for other record types, you have to run the two previous commands again to specify a different record type and append those results to the original CSV file. For example, you would run the following two commands to add SharePoint file activities from the same date range to the PowerShellAuditlog.csv file.

```
$auditlog = Search-UnifiedAuditLog -StartDate 06/01/2019 -EndDate 06/30/2019 -RecordType
SharePointFileOperation
```

```
$auditlog | Select-Object -Property CreationDate,UserIds,RecordType,AuditData | Export-Csv -Append -Path
c:\AuditLogs\PowerShellAuditlog.csv -NoTypeInformation
```

Tips for exporting and viewing the audit log

Here are some tips and examples of exporting and viewing the audit log before and after you use the JSON transform feature to split the **AuditData** column into multiple columns.

- Filter the **RecordType** column to display only the records from a specific service or functional area. For example, to show events related to SharePoint sharing, you would select **14** (the enum value for records triggered by SharePoint sharing activities). For a list of the services that correspond to the enum values displayed in the **RecordType** column, see Detailed properties in the audit log.
- Filter the **Operations** column to display the records for specific activities. For a list of most operations that correspond to a searchable activity in the audit log search tool in the Security & Compliance Center, see the "Audited activities" section in Search the audit log in the Security & Compliance Center.

Search the audit log to investigate common support issues

11/2/2020 • 18 minutes to read • Edit Online

This article describes how to use the audit log search tool to help you investigate common support issues. This includes using the audit log to:

- Find the IP address of the computer used to access a compromised account
- Determine who set up email forwarding for a mailbox
- Determine if a user deleted email items in their mailbox
- Determine if a user created an inbox rule
- Investigate why there was a successful login by a user outside your organization
- Search for mailbox activities performed by users with non-E5 licenses
- Search for mailbox activities performed by delegate users

Using the audit log search tool

Each of the troubleshooting scenarios described in this article is based on using the audit log search tool in the Security & Compliance Center. This section lists the permissions required to search the audit log and describes the steps to access and run audit log searches. Each scenario section explains how to configure an audit log search query and what to look for in the detailed information in the audit records that match the search criteria.

Permissions required to use the audit log search tool

You must be assigned the View-Only Audit Logs or Audit Logs role in Exchange Online to search the audit log. By default, these roles are assigned to the Compliance Management and Organization Management role groups on the **Permissions** page in the Exchange admin center. Global administrators in Office 365 and Microsoft 365 are automatically added as members of the Organization Management role group in Exchange Online. For more information, see Manage role groups in Exchange Online.

Running audit log searches

This section describes the basics for creating and running audit log searches. Use these instructions as a starting point for each troubleshooting scenario in this article. For more detailed step-by-step instructions, see Search the audit log.

1. Go to https://protection.office.com/unifiedauditlog and sign in using your work or school account.

   The **Audit log search** page is displayed.

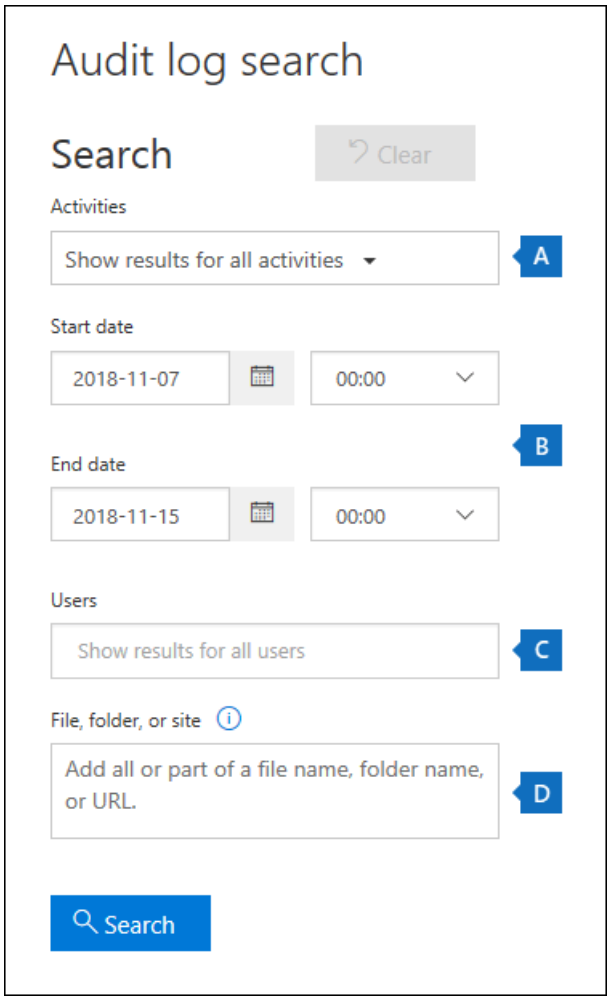

2. You can configure the following search criteria. Each troubleshooting scenario in this article recommends specific guidance for configuring these fields.

   a. **Activities:** Select the drop-down list to display the activities that you can search for. After you run the search, only the audit records for the selected activities are displayed. Selecting **Show results for all activities** displays results for all activities that meet the other search criteria. You'll also have to leave this field blank in some of the troubleshooting scenarios.

   b. **Start date** and **End date:** Select a date and time range to display the events that occurred within that period. The last seven days are selected by default. The date and time are presented in Coordinated Universal Time (UTC) format. The maximum date range that you can specify is 90 days.

   c. **Users:** Click in this box and then select one or more users to display search results for. Audit records for the selected activity performed by the users you select in this box are displayed in the list of results. Leave this box blank to return entries for all users (and service accounts) in your organization.

   d. **File, folder, or site:** Type some or all of a file or folder name to search for activity related to the file of folder that contains the specified keyword. You can also specify a URL of a file or folder. If you use a URL, be sure the type the full URL path or if you only type a portion of the URL, don't include any special characters or spaces. Leave this box blank to return entries for all files and folders in your organization. This field is left blank in all the troubleshooting scenarios in this article.

3. Select **Search** to run the search using your search criteria.

   The search results are loaded, and after a few moments they're displayed under **Results** on the **Audit log search** page. Each of the sections in this article provides guidance about things to look for in the context of the specific troubleshooting scenario.

   For more information about viewing, filtering, or exporting audit log search results, see:

- View search results
- Filter search results
- Export search results

Find the IP address of the computer used to access a compromised account

The IP address corresponding to an activity performed by any user is included in most audit records. Information about the client used is also included in the audit record.

Here's how to configure an audit log search query for this scenario:

Activities: If relevant to your case, select a specific activity to search for. For troubleshooting compromised accounts, consider selecting the **User signed in to mailbox** activity under **Exchange mailbox activities**. This returns auditing records showing the IP address that was use when signing in to the mailbox. Otherwise, leave this field blank to return audit records for all activities.

> **TIP**
> Leaving this field blank will return **UserLoggedIn** activities, which is an Azure Active Directory activity that indicates that someone has signed in to an user account. Use filtering in the search results to display the **UserLoggedIn** audit records.

Start date and **End date:** Select a date range that's applicable to your investigation.

Users: If you're investigating a compromised account, select the user whose account was compromised. This returns audit records for activities performed by that user account.

File, folder, or site: Leave this field blank.

After you run the search, the IP address for each activity is displayed in the **IP address** column in the search results. Select the record in the search results to view more detailed information on the flyout page.

Determine who set up email forwarding for a mailbox

When email forwarding is configured for a mailbox, email messages that are sent to the mailbox are forwarded to another mailbox. Messages can be forwarded to users inside or outside of your organization. When email forwarding is set up on a mailbox, the underlying Exchange Online cmdlet that's used is **Set-Mailbox**.

Here's how to configure an audit log search query for this scenario:

Activities: Leave this field blank so that the search returns audit records for all activities. This is necessary to return any audit records related to the **Set-Mailbox** cmdlet.

Start date and **End date:** Select a date range that's applicable to your investigation.

Users: Unless you're investigating an email forwarding issue for a specific user, leave this field blank. This helps you identify if email forwarding was set up for any user.

File, folder, or site: Leave this field blank.

After you run the search, select **Filter results** on the search results page. In the box under **Activity** column header, type **Set-Mailbox** so that only audit records related to the **Set-Mailbox** cmdlet are displayed.

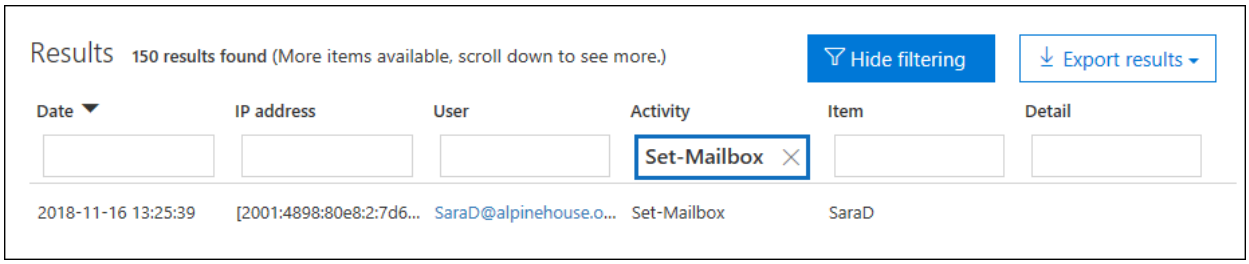

At this point, you have to look at the details of each audit record to determine if the activity is related to email forwarding. Select the audit record to display the **Details** flyout page, and then select **More information**. The following screenshot and descriptions highlight the information that indicates email forwarding was set on the mailbox.

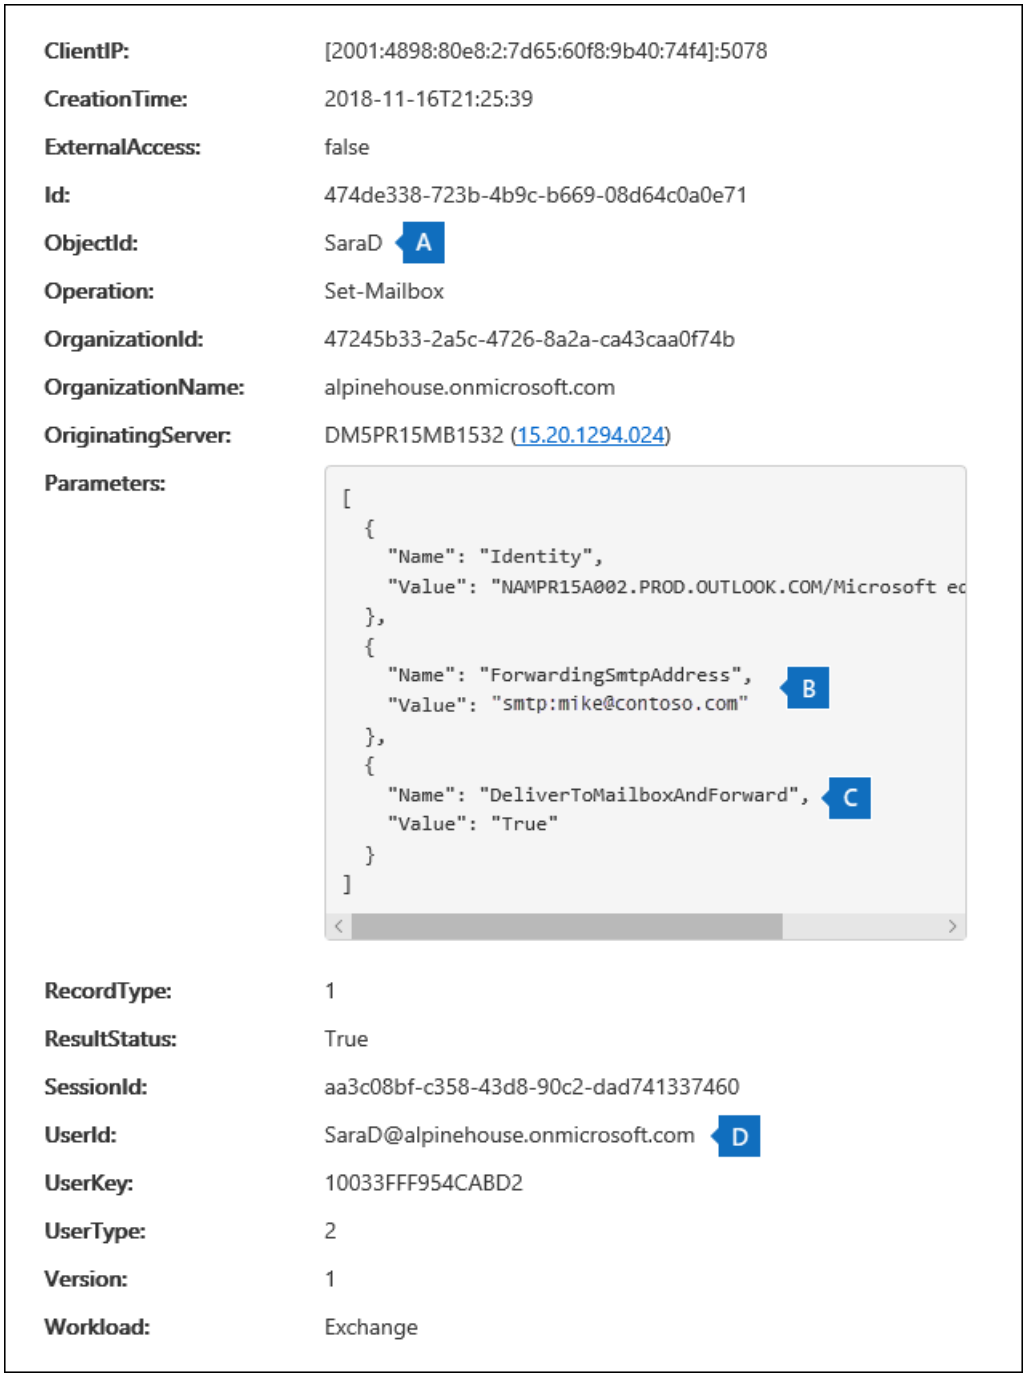

a. In the **ObjectId** field, the alias of the mailbox that email forwarding was set on is displayed. This mailbox is also displayed on the **Item** column in the search results page.

b. In the **Parameters** field, The value *ForwardingSmtpAddress* indicates that email forwarding was set on the mailbox. In this example, mail is being forwarded to the email address mike@contoso.com, which is outside of the alpinehouse.onmicrosoft.com organization.

c. The *True* value for the *DeliverToMailboxAndForward* parameter indicates that a copy of the message is delivered to sarad@alpinehouse.onmicrosoft.com *and* is forwarded to the email address specified by the *ForwardingSmtpAddress* parameter, which in this example is mike@contoso.com. If the value for the *DeliverToMailboxAndForward* parameter is set to *False*, then email is only forwarded to the address specified by the *ForwardingSmtpAddress* parameter. It's not delivered to the mailbox specified in the **ObjectId** field.

d. The **UserId** field indicates the user who set email forwarding on the mailbox specified in the **ObjectId** field. This user is also displayed in the **User** column on the search results page. In this case, it seems that the owner of the mailbox set email forwarding on her mailbox.

If you determine that email forwarding shouldn't be set on the mailbox, you can remove it by running the following command in Exchange Online PowerShell:

```
Set-Mailbox <mailbox alias> -ForwardingSmtpAddress $null
```

For more information about the parameters related to email forwarding, see the Set-Mailbox article.

Determine if a user deleted email items

Starting in January 2019, Microsoft is turning on mailbox audit logging by default for all Office 365 and Microsoft organizations. This means that certain actions performed by mailbox owners are automatically logged, and the corresponding mailbox audit records are available when you search for them in the mailbox audit log. Before mailbox auditing was turned on by default, you had to manually enable it for every user mailbox in your organization.

The mailbox actions logged by default include the SoftDelete and HardDelete mailbox actions performed by mailbox owners. This means you can use the following steps to search the audit log for events related to deleted email items. For more information about mailbox auditing on by default, see Manage mailbox auditing.

Here's how to configure an audit log search query for this scenario:

Activities: Under **Exchange mailbox activities**, select one or both of the following activities:

- **Deleted messages from Deleted Items folder:** This activity corresponds to the **SoftDelete** mailbox auditing action. This activity is also logged when a user permanently deletes an item by selecting it and pressing **Shift+Delete**. After an item is permanently deleted, the user can recover it until the deleted item retention period expires.
- **Purged messages from mailbox:** This activity corresponds to the **HardDelete** mailbox auditing action. This is logged when a user purges an item from the Recoverable Items folder. Admins can use the Content Search tool in the security and compliance center to search for and recover purged items until the deleted item retention period expires or longer if the user's mailbox is on hold.

Start date and **End date:** Select a date range that's applicable to your investigation.

Users: If you select a user in this field, the audit log search tool returns audit records for email items that were deleted (SoftDeleted or HardDeleted) by the user you specify. Sometimes the user who deletes an email might not be the mailbox owner.

File, folder, or site: Leave this field blank.

After you run the search, you can filter the search results to display the audit records for soft-deleted items or for hard-deleted items. Select the audit record to display the **Details** flyout page, and then select **More information**. Additional information about the deleted item, such as the subject line and the location of the item when it was deleted, is displayed in the **AffectedItems** field. The following screenshots show an example of the **AffectedItems** field from a soft-deleted item and a hard-deleted item.

Example of AffectedItems field for soft-deleted item


```
[
  {
    "Id": "RgAAAABupAfm1J7zSbo2hrt6nSJ3BwAjcPtFX4L8RL4xOX9HtjKnAAAA
    "InternetMessageId": "<SN2PR15MB1038A559832DBEF4331B0312F00A0@S
    "ParentFolder": {
      "Id": "LgAAAABupAfm1J7zSbo2hrt6nSJ3AQAjcPtFX4L8RL4xOX9HtjKnAA
      "Path": "\\Deleted Items"
    },
    "Subject": "Search Results-7/22/2016 8:15:48 PM"
  }
]
```


Example of AffectedItems field for hard-deleted item


```
[
  {
    "Id": "RgAAAAAF2Bt1vQxATYXa84f0wpBlBwBNbdf5Q8KQQaZXh+MPAq0fAAA
    "InternetMessageId": "<89423594-a3b8-4886-8d82-83b8896c0c18@SN
    "ParentFolder": {
      "Id": "LgAAAAAF2Bt1vQxATYXa84f0wpBlAQBNbdf5Q8KQQaZXh+MPAq0fA
      "Path": "\\Recoverable Items\\Deletions"
    },
    "Subject": "Items labeled as 'OneDrive retention label' are re
  }
]
```


Recover deleted email items

Users can recover soft-deleted items if the deleted items retention period has not expired. In Exchange Online, the default deleted items retention period is 14 days, but admins can increase this setting to a maximum of 30 days. Point users to the Recover deleted items or email in Outlook on the web article for instructions on recovering deleted items.

As previously explained, admins may be able to recover hard-deleted items if the deleted item retention period hasn't expired or if the mailbox is on hold, in which case items are kept until the hold duration expires. When you run a content search, soft-deleted and hard-deleted items in the Recoverable Items folder are returned in the search results if they match the search query. For more information about running content searches, see Content Search in Office 365.

> **TIP**
> To search for deleted email items, search for all or part of the subject line that's displayed in the **AffectedItems** field in the audit record.

Determine if a user created an inbox rule

When users create an inbox rule for their Exchange Online mailbox, a corresponding audit record is saved to the audit log. For more information about inbox rules, see:

- Use inbox rules in Outlook on the web
- Manage email messages in Outlook by using rules

Here's how to configure an audit log search query for this scenario:

Activities: Under **Exchange mailbox activities**, select **New-InboxRule Create/modify/enable/disable inbox rule**.

Start date and **End date:** Select a date range that's applicable to your investigation.

Users: Unless you're investigating a specific user, leave this field blank. This helps you identify new inbox rules set up by any user.

File, folder, or site: Leave this field blank.

After you run the search, any audit records for this activity are displayed in the search results. Select an audit record to display the **Details** flyout page, and then select **More information**. Information about the inbox rule settings is displayed in the **Parameters** field. The following screenshot and descriptions highlight the information about inbox rules.

ClientIP: [2001:4898:80e8:1:b989:30b9:cd9f:5eea]:13975

CreationTime: 2018-11-27T18:50:21

ExternalAccess: false

Id: d59f8802-3556-40ff-2d0f-08d654992f14

ObjectId: A NAMPR15A002.PROD.OUTLOOK.COM/Microsoft Exchange Hosted Organizations/alpinehouse.onmicrosoft.com/SaraD\Move messages from admin

Operation: New-InboxRule

OrganizationId: 47245b33-2a5c-4726-8a2a-ca43caa0f74b

OrganizationName: alpinehouse.onmicrosoft.com

OriginatingServer: DM5PR15MB1532 (15.20.1361.017)

Parameters:


```
[
  {
    "Name": "AlwaysDeleteOutlookRulesBlob",
    "Value": "False"
  },
  {
    "Name": "Force",
    "Value": "False"
  },
  {
    "Name": "From",
    "Value": "admin@alpinehouse.onmicrosoft.com"
  },
  {
    "Name": "MoveToFolder",
    "Value": "adminsearch"
  },
  {
    "Name": "Name",
    "Value": "Move messages from admin"
  },
  {
    "Name": "StopProcessingRules",
    "Value": "False"
  }
]
```


B

C

RecordType: 1

ResultStatus: True

SessionId: 4bef64fd-7009-4531-a915-4e2eebe1f1ee

UserId: SaraD@alpinehouse.onmicrosoft.com D

UserKey: 10033FFF954CABD2

UserType: 2

Version: 1

Workload: Exchange

a. In the **ObjectId** field, the full name of the inbox rule is displayed. This name includes the alias of the user's mailbox (for example, SaraD) and the name of the inbox rule (for example, "Move messages from admin").

b. In the **Parameters** field, the condition of the inbox rule is displayed. In this example, the condition is specified by the *From* parameter. The value defined for the *From* parameter indicates that the inbox rule acts on email sent by admin@alpinehouse.onmicrosoft.com. For a complete list of the parameters that can be used to define conditions of inbox rules, see the New-InboxRule article.

c. The *MoveToFolder* parameter specifies the action for the inbox rule. In this example, messages received from admin@alpinehouse.onmicrosoft.com are moved to the folder named *AdminSearch*. Also see the New-InboxRule article for a complete list of parameters that can be used to define the action of an inbox rule.

d. The **UserId** field indicates the user who created the inbox rule specified in the **ObjectId** field. This user is also displayed in the **User** column on the search results page.

Investigate why there was a successful login by a user outside your organization

When reviewing audit records in the audit log, you may see records that indicate an external user was authenticated by Azure Active Directory and successfully logged in to your organization. For example, an admin in contoso.onmicrosoft.com may see an audit record showing that a user from a different organization (for example, fabrikam.onmicrosoft.com) successfully logged into contoso.onmicrosoft.com. Similarly, you may see audit records that indicate users with a Microsoft Account (MSA), such as an Outlook.com or Live.com, successfully logged in to your organization. In these situations, the audited activity is **User logged In**.

This behavior is by design. Azure Active Directory (Azure AD), the directory service, allows something called *pass-through authentication* when an external user tries to access a SharePoint site or a OneDrive location in your organization. When the external user tries to do this, they're prompted to enter their credentials. Azure AD uses the credentials to authenticate the user, meaning only Azure AD verifies that the user is who they say they are. The indication of the successful login in the audit record is the result of Azure AD authenticating the user. The successful login doesn't mean that the user was able to access any resources or perform any other actions in your organization. It only indicates that the user was authenticated by Azure AD. In order for a pass-through user to access SharePoint or OneDrive resources, a user in your organization would have to explicitly share a resource with the external user by sending them a sharing invitation or anonymous sharing link.

> **NOTE**
> Azure AD allows pass-through authentication only for *first-party applications*, such as SharePoint Online and OneDrive for Business. It isn't allowed for other third-party applications.

Here's an example and descriptions of relevant properties in an audit record for a **User logged In** event that is a result of pass-through authentication. Select the audit record to display the **Details** flyout page, and then select **More information**.

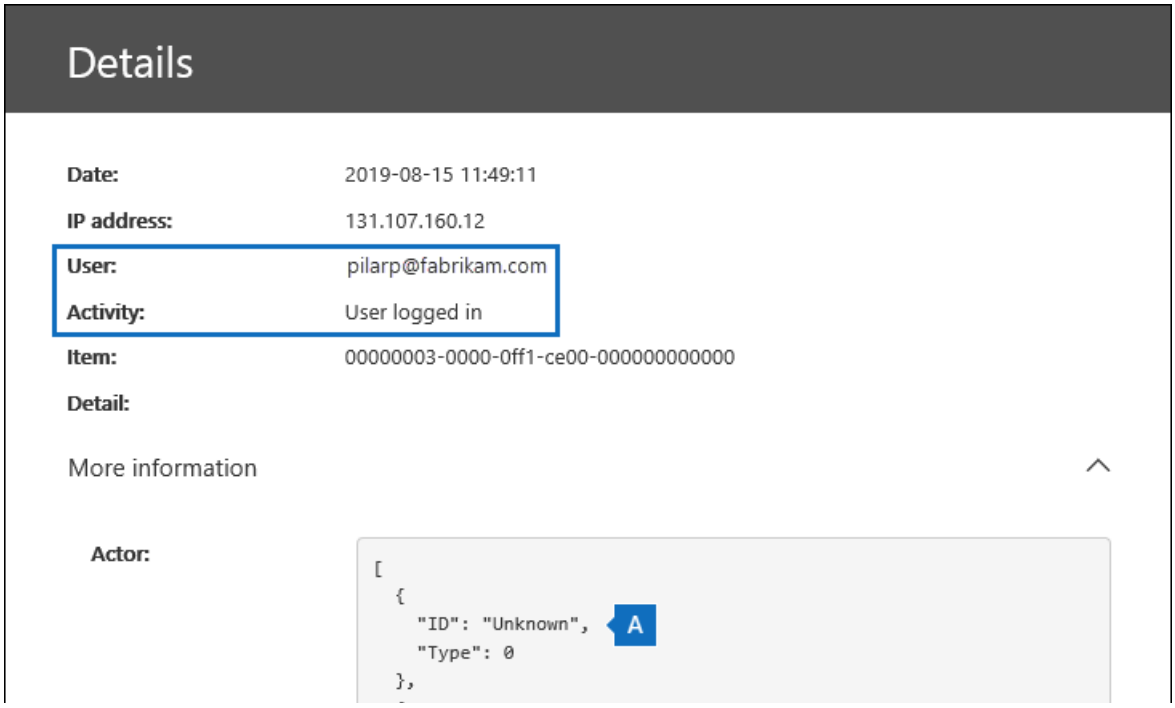

```
  {
    "ID": "pilarp@fabrikam.com",
    "Type": 5
  }
]
```


B

ActorContextId: d4010a10-1844-4528-8b8c-0aa89c666b1f

ActorIpAddress: 131.107.160.12

ApplicationId: 00000003-0000-0ff1-ce00-000000000000

C

AzureActiveDirectoryEver 1

ClientIP: 131.107.160.12

CreationTime: 2019-08-15T18:49:11

ExtendedProperties:


```
[
  {
    "Name": "UserAgent",
    "Value": "Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/
  },
  {
    "Name": "UserAuthenticationMethod",
    "Value": "1"
  },
  {
    "Name": "RequestType",
    "Value": "OAuth2:Authorize"
  },
  {
    "Name": "ResultStatusDetail",
    "Value": "Success"
  },
  {
    "Name": "KeepMeSignedIn",
    "Value": "True"
  }
]
```


D

Id: 8e80bcc0-9c55-40d5-b806-deaa3fd244b4

InterSystemsId: 998cfa9e-a08d-9000-79fd-29c87a9e9e71

IntraSystemId: 7c00fa4a-9aca-4650-83f3-11270b851f00

ModifiedProperties:


```
[]
```


ObjectId: 00000003-0000-0ff1-ce00-000000000000

Operation: UserLoggedIn

OrganizationId: 753f6a8d-8546-48cc-93fb-76d681d7e0a0

RecordType: 15

E

ResultStatus: Succeeded

SupportTicketId:

Target:


```
[
  {
    "ID": "00000003-0000-0ff1-ce00-000000000000",
    "Type": 0
  }
]
```


TargetContextId: 753f6a8d-8546-48cc-93fb-76d681d7e0a0

UserId: pilarp@fabrikam.com

UserKey: Not Available

UserType: 0

Version:	1
Workload:	AzureActiveDirectory

a. This field indicates that the user who attempted to access a resource in your organization wasn't found in your organization's Azure AD.

b. This field displays the UPN of the external user that attempted to access a resource in your organization. This user ID is also identified in the **User** and **UserId** properties in the audit record.

c. The **ApplicationId** property identifies the application that triggered the logon request. The value of 00000003-0000-0ff1-ce00-000000000000 displayed in the ApplicationId property in this audit record indicates SharePoint Online. OneDrive for Business also has this same ApplicationId.

d. This indicates that the pass-through authentication was successful. In other words, the user was successfully authenticated by Azure AD.

e. The **RecordType** value of **15** indicates that the audited activity (UserLoggedIn) is a Secure Token Service (STS) logon event in Azure AD.

For more information about the other properties displayed in a UserLoggedIn audit record, see the Azure AD-related schema information in Office 365 Management Activity API schema.

Here are two examples scenarios that would result in a successful **User logged in** audit activity because of pass-through authentication:

- A user with a Microsoft Account (such as SaraD@outlook.com) has tried to access a document in a OneDrive for Business account in fourthcoffee.onmicrosoft.com and there isn't a corresponding guest user account for SaraD@outlook.com in fourthcoffee.onmicrosoft.com.
- A user with a Work or School account in an organization (such as pilarp@fabrikam.onmicrosoft.com) has tried to access a SharePoint site in contoso.onmicrosoft.com and there isn't a corresponding guest user account for pilarp@fabrikam.com in contoso.onmicrosoft.com.

Tips for investigating successful logins resulting from pass-through authentication

- Search the audit log for activities performed by the external user identified in the **User logged in** audit record. Type the UPN for the external user in the **Users** box and use a date range if relevant to your scenario. For example, you can create a search using the following search criteria:

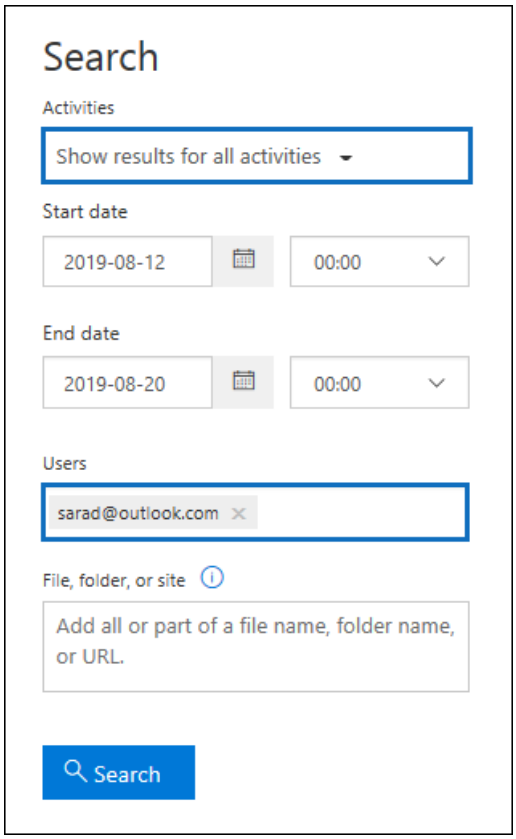

In addition to the **User logged in** activities, other audit records may be returned, such ones that indicate a user in your organization shared resources with the external user and whether the external user accessed, modified, or downloaded a document that was shared with them.

- Search for SharePoint sharing activities that would indicate a file was shared with the external user identified by a **User logged in** audit record. For more information, see Use sharing auditing in the audit log.
- Export the audit log search results that contain records relevant to your investigation so that you can use Excel to search for other activities related to the external user. For more information, see Export, configure, and view audit log records.

Search for mailbox activities performed by users with non-E5 licenses

Even when mailbox auditing on by default is turned on for your organization, you might notice that mailbox audit events for some users aren't found in audit log searches by using the compliance center, the **Search-UnifiedAuditLog** cmdlet, or the Office 365 Management Activity API. The reason for this is that mailbox audit events will be returned only for users with E5 licenses when you one of the previous methods to search the unified audit log.

To retrieve mailbox audit log records for non-E5 users, you can perform one of the following workarounds:

- Manually enable mailbox auditing on individual mailboxes (run the `Set-Mailbox -Identity <MailboxIdentity> -AuditEnabled $true` command in Exchange Online PowerShell). After you do this, search for mailbox audit activities by using the compliance center, the **Search-UnifiedAuditLog** cmdlet, or the Office 365 Management Activity API.

> **NOTE**
>
> If mailbox auditing already appears to be enabled on the mailbox, but your searches return no results, change the value of the *AuditEnabled* parameter to `$false` and then back to `$true`.

- Use the following cmdlets in Exchange Online PowerShell:
  - Search-MailboxAuditLog to search the mailbox audit log for specific users.
  - New-MailboxAuditLogSearch to search the mailbox audit log for specific users and to have the results sent via email to specified recipients.

Search for mailbox activities performed in a specific mailbox (including shared mailboxes)

When you use the **Users** dropdown list in the audit log search tool in the compliance center or the **Search-UnifiedAuditLog -UserIds** command in Exchange Online PowerShell, you can search for activities performed by a specific user. For mailbox audit activities, this type of search will search for activities performed by the specified user. It doesn't guarantee that all activities performed in the same mailbox are returned in the search results. For example, an audit log search won't return audit records for activities performed by a delegate user because searching for mailbox activities performed by a specific user will not return activities performed by a delegate user who's been assigned permissions to access another user's mailbox. (A delegate user is someone who's been assigned the SendAs, SendOnBehalf, or FullAccess mailbox permission to another user's mailbox.)

Also, using the **User** dropdown list in the audit log search tool or the **Search-UnifiedAuditLog -UserIds** will not return results for activities performed in a shared mailbox.

To search for the activities performed in a specific mailbox or to search for activities performed in a shared mailbox, use the following syntax when running the **Search-UnifiedAuditLog** cmdlet:

```
Search-UnifiedAuditLog  -StartDate <date> -EndDate <date> -FreeText (Get-Mailbox <mailbox
identity).ExchangeGuid
```

For example, the following command returns audit records for activities performed in the Contoso Compliance Team shared mailbox between August 2020 and October 2020:

```
Search-UnifiedAuditLog  -StartDate 08/01/2020 -EndDate 10/31/2020 -FreeText (Get-Mailbox
complianceteam@contoso.onmicrosoft.com).ExchangeGuid
```

Alternatively, you can use the **Search-MailboxAuditLog** cmdlet to search for audit records for activity performed in a specific mailbox. This includes searching for activities performed in a shared mailbox.

The following example returns mailbox audit log records for activities performed in the Contoso Compliance Team shared mailbox:

```
Search-MailboxAuditLog -Identity complianceteam@contoso.onmicrosoft.com -StartDate 08/01/2020 -EndDate
10/31/2020 -ShowDetails
```

The following example returns mailbox audit log records for activities performed in the specified mailbox by delegate users:

```
Search-MailboxAuditLog -Identity <mailbox identity> -StartDate <date> -EndDate <date> -LogonTypes Delegate -
ShowDetails
```

You can also use the **New-MailboxAuditLogSearch** cmdlet to search the audit log for a specific mailbox and to have the results sent via email to specified recipients.

Use sharing auditing in the audit log

11/2/2020 • 7 minutes to read • Edit Online

Sharing is a key activity in SharePoint Online and OneDrive for Business, and it's widely used in organizations. Administrators can use sharing auditing in the audit log to determine how sharing is used in their organization.

The SharePoint Sharing schema

Sharing events (not including events related to sharing policy and sharing links) are different from file- and folder-related events in one primary way: one user is performing an action that has an effect on another user. For example, when a resource User A gives User B access to a file. In this example, User A is the *acting user* and User B is the *target user*. In the SharePoint File schema, the acting user's action only affects the file itself. When User A opens a file, the only information needed in the **FileAccessed** event is the acting user. To address this difference, there is a separate schema, called the *SharePoint Sharing schema*, that captures more information about sharing events. This ensures that administrators have visibility into who shared a resource and the user the resource was shared with.

The Sharing schema provides two additional fields in an audit record related to sharing events:

- **TargetUserOrGroupType:** Identifies whether the target user or group is a Member, Guest, SharePointGroup, SecurityGroup, or Partner.
- **TargetUserOrGroupName:** Stores the UPN or name of the target user or group that a resource was shared with (User B in the previous example).

These two fields, in addition to other properties from the audit log schema such as User, Operation, and Date can tell the full story about *which* user shared *what* resource with *whom* and *when*.

There's another schema property that's important to the sharing story. When you export audit log search results, the **AuditData** column in the exported CSV file stores information about sharing events. For example, when a user shares a site with another user, this is accomplished by adding the target user to a SharePoint group. The **AuditData** column captures this information to provide context for administrators. See Step 2 for instructions on how to parse the information in the **AuditData** column.

SharePoint sharing events

Sharing is defined by when a user (the *acting* user) wants to share a resource with another user (the *target* user). Audit records related to sharing a resource with an external user (a user who is outside of your organization and doesn't have a guest account in your organization's Azure Active Directory) are identified by the following events, which are logged in the audit log:

- **SharingInvitationCreated:** A user in your organization tried to share a resource (likely a site) with an external user. This results in an external sharing invitation sent to the target user. No access to the resource is granted at this point.
- **SharingInvitationAccepted:** The external user has accepted the sharing invitation sent by the acting user and now has access to the resource.
- **AnonymousLinkCreated:** An anonymous link (also called an "Anyone" link) is created for a resource. Because an anonymous link can be created and then copied, it's reasonable to assume that any document that has an anonymous link has been shared with a target user.
- **AnonymousLinkUsed:** As the name implies, this event is logged when an anonymous link is used to

access a resource.

- **SecureLinkCreated:** A user has created a "specific people link" to share a resource with a specific person. This target user may be someone who is external to your organization. The person that the resource is shared with is identified in the audit record for the **AddedToSecureLink** event. The time stamps for these two events are nearly identical.
- **AddedToSecureLink:** A user was added to a specific people link. Use the **TargetUserOrGroupName** field in this event to identify the user added to the corresponding specific people link. This target user may be someone who is external to your organization.

Sharing auditing work flow

When a user (the acting user) wants to share a resource with another user (the target user), SharePoint (or OneDrive for Business) first checks if the email address of the target user is already associated with a user account in the organization's directory. If the target user is in the directory (and has a corresponding guest user account), SharePoint does the following things:

- Immediately assigns the target user permissions to access the resource by adding the target user to the appropriate SharePoint group, and logs an **AddedToGroup** event.
- Sends a sharing notification to the email address of the target user.
- Logs a **SharingSet** event. This event has a friendly name of "Shared file, folder, or site" under **Sharing and access request activities** in the activities picker of the audit log search tool. See the screenshot in Step 1.

If a user account for the target user isn't in the directory, SharePoint does the following:

- Logs one of the following events, based on how the resource is shared:
  - **AnonymousLinkCreated**
  - **SecureLinkCreated**
  - **AddedToSecureLink**
  - **SharingInvitationCreated** (this event is logged only when the shared resource is a site)
- When the target user accepts the sharing invitation that's sent to them (by clicking the link in the invitation), SharePoint logs a **SharingInvitationAccepted** event and assigns the target user permissions to access the resource. If the target user is sent an anonymous link, the **AnonymousLinkUsed** event is logged after the target user uses the link to access the resource. For secure links, a **FileAccessed** event is logged when an external user uses the link to access the resource.

Additional information about the target user is also logged, such as the identity of the user the invitation is to and the user who accepts the invitation. In some case, these users (or email addresses) can be different.

How to identify resources shared with external users

A common requirement for administrators is creating a list of all resources that have been shared with users outside of the organization. By using sharing auditing in Office 365, administrators can generate this list. Here's how.

Step 1: Search for sharing events and export the results to a CSV file

The first step is to search the audit log for sharing events. For more information (including the required permissions) about searching the audit log, see Search the audit log in the Security & Compliance Center.

1. Go to https://protection.office.com.

2. Sign in using your work or school account.
3. In the left pane of the Security & Compliance Center, click **Search** > **Audit log search**.

   The **Audit log search** page is displayed.
4. Under **Activities**, click **Sharing and access request activities** to search for sharing-related events.

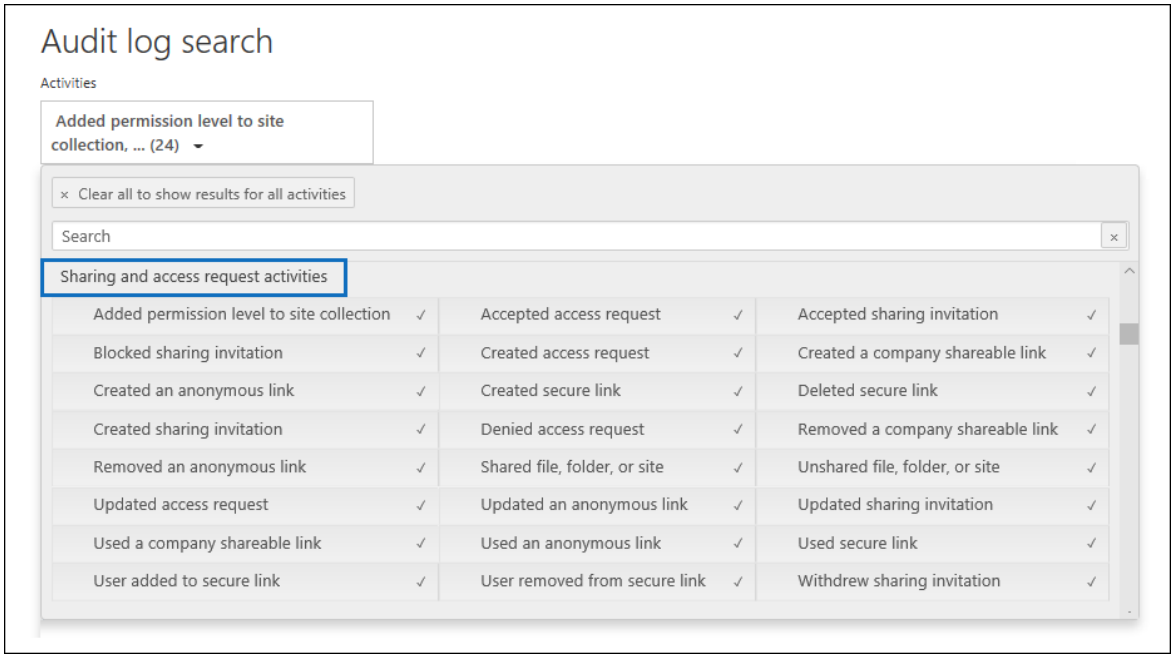

5. Select a date and time range to find the sharing events that occurred within that period.
6. Click **Search** to run the search.
7. When the search is finished running and the results are displayed, click **Export results** > **Download all results**.

   After you select the export option, a message at the bottom of the window prompts you to open or save the CSV file.
8. Click **Save** > **Save as** and save the CSV file to a folder on your local computer.

Step 2: Use the PowerQuery Editor to format the exported audit log

The next step is to use the JSON transform feature in the Power Query Editor in Excel to split each property in the **AuditData** column (which consists of a multi-property JSON object) into its own column. This lets you filter columns to view records related to sharing

For step-by-step instructions, see "Step 2: Format the exported audit log using the Power Query Editor" in Export, configure, and view audit log records.

Step 3: Filter the CSV file for resources shared with external users

The next step is to filter the CSV for the different sharing-related events that were previously described in the SharePoint sharing events section. Alternatively, you can filter the **TargetUserOrGroupType** column to display all records where the value of this property is **Guest**.

After you've followed the instructions in the previous step to prepare the CSV file by using the PowerQuery editor, do the following:

1. Open the Excel file that you created in Step 2.
2. On the **Home** tab, click **Sort & Filter**, and then click **Filter**.
3. In the **Sort & Filter** dropdown list on the **Operations** column, clear all selections, then select one or

more the following sharing-related events and then click **Ok**.

- **SharingInvitationCreated**
- **AnonymousLinkCreated**
- **SecureLinkCreated**
- **AddedToSecureLink**

Excel displays the rows for the events you selected.

4. Go to the column named **TargetUserOrGroupType** and select it.
5. In the **Sort & Filter** dropdown list, clear all selections, then select **TargetUserOrGroupType:Guest**, and click **Ok**.

   Now Excel displays the rows for sharing events AND where the target user is outside of your organization, because external users are identified by the value **TargetUserOrGroupType:Guest**.

> **TIP**
>
> For the audit records that are displayed, the **ObjectId** column identifies the resource that was shared with the target user; for example
>
> ```
> ObjectId:https:\/\/contoso-my.sharepoint.com\/personal\/sarad_contoso_com\/Documents\/Southwater Proposal.docx
> ```
>
> .

Manage mailbox auditing

2/18/2021 • 18 minutes to read • Edit Online

Starting in January 2019, Microsoft is turning on mailbox audit logging by default for all organizations. This means that certain actions performed by mailbox owners, delegates, and admins are automatically logged, and the corresponding mailbox audit records will be available when you search for them in the mailbox audit log. Before mailbox auditing was turned on by default, you had to manually enable it for every user mailbox in your organization.

Here are some benefits of mailbox auditing on by default:

- Auditing is automatically enabled when you create a new mailbox. You don't need to manually enable it for new users.
- You don't need to manage the mailbox actions that are audited. A predefined set of mailbox actions are audited by default for each logon type (Admin, Delegate, and Owner).
- When Microsoft releases a new mailbox action, the action might be automatically added to the list of mailbox actions that are audited by default (subject to the user having the appropriate license). This means you don't need to monitor add new actions on mailboxes.
- You have a consistent mailbox auditing policy across your organization (because you're auditing the same actions for all mailboxes).

NOTE

- The important thing to remember about the release of mailbox auditing on by default is: you don't need to do anything to manage mailbox auditing. However, to learn more, customize mailbox auditing from the default settings, or turn it off altogether, this topic can help you.
- By default, only mailbox audit events for E5 users are available in audit log searches in the Security & Compliance Center or via the Office 365 Management Activity API. For more information, see the More information section in this topic.

Verify mailbox auditing on by default is turned on

To verify that mailbox auditing on by default is turned on for your organization, run the following command in Exchange Online PowerShell:

```
Get-OrganizationConfig | Format-List AuditDisabled
```

The value **False** indicates that mailbox auditing on by default is enabled for the organization. This on by default organizational value overrides the mailbox auditing setting on specific mailboxes. For example, if mailbox auditing is disabled for a mailbox (the *AuditEnabled* property is **False** on the mailbox), the default mailbox actions will still be audited for the mailbox, because mailbox auditing on by default is enabled for the organization.

To keep mailbox auditing disabled for specific mailboxes, you configure mailbox auditing bypass for the mailbox owner and other users who have been delegated access to the mailbox. For more information, see the Bypass mailbox audit logging section in this topic.

> **NOTE**
>
> When mailbox auditing on by default is turned on for the organization, the *AuditEnabled* property for affected mailboxes won't be changed from **False** to **True**. In other words, mailbox auditing on by default ignores the *AuditEnabled* property on mailboxes.

Supported mailbox types

The following table shows the mailbox types that are currently supported by mailbox auditing on by default:

MAILBOX TYPE	SUPPORTED	NOT SUPPORTED
User mailboxes	✓	
Shared mailboxes	✓	
Microsoft 365 Group mailboxes	✓	
Resource mailboxes		✓
Public folder mailboxes		✓

Logon types and mailbox actions

Logon types classify the user that did the audited actions on the mailbox. The following list describes the logon types that are used in mailbox audit logging:

- **Owner**: The mailbox owner (the account that's associated with the mailbox).
- **Delegate**:
  - A user who's been assigned the SendAs, SendOnBehalf, or FullAccess permission to another mailbox.
  - An admin who's been assigned the FullAccess permission to a user's mailbox.
- **Admin**:
  - The mailbox is searched with one of the following Microsoft eDiscovery tools:
    - Content Search in the Compliance center.
    - eDiscovery or Advanced eDiscovery in the Compliance center.
    - In-Place eDiscovery in Exchange Online.
  - The mailbox is accessed by using the Microsoft Exchange Server MAPI Editor.

Mailbox actions for user mailboxes and shared mailboxes

The following table describes the mailbox actions that are available in mailbox audit logging for user mailboxes and shared mailboxes.

- A check mark (✓) indicates the mailbox action can be logged for the logon type (not all actions are available for all logon types).
- An asterisk (*) after the check mark indicates the mailbox action is logged by default for the logon type.

- Remember, an admin with Full Access permission to a mailbox is considered a delegate.

MAILBOX ACTION	DESCRIPTION	ADMIN	DELEGATE	OWNER
AddFolderPermissions	**Note**: Although this value is accepted as a mailbox action, it's already included in the **UpdateFolderPermissions** action and isn't audited separately. In other words, don't use this value.			
ApplyRecord	An item is labeled as a record.	✓ *	✓ *	✓ *
Copy	A message was copied to another folder.	✓		
Create	An item was created in the Calendar, Contacts, Notes, or Tasks folder in the mailbox (for example, a new meeting request is created). Creating, sending, or receiving a message isn't audited. Also, creating a mailbox folder is not audited.	✓ *	✓ *	✓
Default		✓	✓	✓
FolderBind	A mailbox folder was accessed. This action is also logged when the admin or delegate opens the mailbox. **Note**: Audit records for folder bind actions performed by delegates are consolidated. One audit record is generated for individual folder access within a 24-hour period.	✓	✓	
HardDelete	A message was purged from the Recoverable Items folder.	✓ *	✓ *	✓ *

MAILBOX ACTION	DESCRIPTION	ADMIN	DELEGATE	OWNER
MailItemsAccessed	Mail data is accessed by mail protocols and clients. This value is only available for E5 or E5 Compliance add-on subscription users. For more information, see Set up Advanced Audit for users.	✓ *	✓ *	✓ *
MailboxLogin	The user signed into their mailbox.			✓
MessageBind	A message was viewed in the preview pane or opened by an admin. **Note**: Although this value is accepted as a mailbox action, these actions are no longer logged.	✓		
ModifyFolderPermissions	**Note**: Although this value is accepted as a mailbox action, it's already included in the **UpdateFolderPermissions** action and isn't audited separately. In other words, don't use this value.			
Move	A message was moved to another folder.	✓	✓	✓
MoveToDeletedItems	A message was deleted and moved to the Deleted Items folder.	✓ *	✓ *	✓ *
RecordDelete	An item that's labeled as a record was soft-deleted (moved to the Recoverable Items folder). Items labeled as records can't be permanently deleted (purged from the Recoverable Items folder).	✓	✓	✓

MAILBOX ACTION	DESCRIPTION	ADMIN	DELEGATE	OWNER
RemoveFolderPermissions	**Note**: Although this value is accepted as a mailbox action, it's already included in the **UpdateFolderPermissions** action and isn't audited separately. In other words, don't use this value.			
Send	The user sends an email message, replies to an email message, or forwards an email message. This value is only available for E5 or E5 Compliance add-on subscription users. For more information, see Set up Advanced Audit for users.	✓ *	✓ *	✓ *
SendAs	A message was sent using the SendAs permission. This means another user sent the message as though it came from the mailbox owner.	✓ *	✓ *	
SendOnBehalf	A message was sent using the SendOnBehalf permission. This means another user sent the message on behalf of the mailbox owner. The message indicates to the recipient who the message was sent on behalf of and who actually sent the message.	✓ *	✓ *	
SoftDelete	A message was permanently deleted or deleted from the Deleted Items folder. Soft-deleted items are moved to the Recoverable Items folder.	✓ *	✓ *	✓ *

MAILBOX ACTION	DESCRIPTION	ADMIN	DELEGATE	OWNER
Update	A message or its properties was changed.	✓ *	✓ *	✓ *
UpdateCalendarDelegation	A calendar delegation was assigned to a mailbox. Calendar delegation gives someone else in the same organization permissions to manage the mailbox owner's calendar.	✓ *		✓ *
UpdateComplianceTag	A different retention label is applied to a mail item (an item can only have one retention label assigned to it).	✓	✓	✓
UpdateFolderPermissions	A folder permission was changed. Folder permissions control which users in your organization can access folders in a mailbox and the messages located in those folders.	✓ *	✓ *	✓ *
UpdateInboxRules	An inbox rule was added, removed, or changed. Inbox rules are used to process messages in the user's Inbox based on the specified conditions and take actions when the conditions of a rule are met, such as moving a message to a specified folder or deleting a message.	✓ *	✓ *	✓ *

IMPORTANT

If you customized the mailbox actions to audit for any logon type *before* mailbox auditing on by default was enabled in your organization, the customized settings are preserved on the mailbox and aren't overwritten by the default mailbox actions as described in this section. To revert the audit mailbox actions to their default values (which you can do at any time), see the Restore the default mailbox actions section later in this topic.

Mailbox actions for Microsoft 365 Group mailboxes

Mailbox auditing on by default brings mailbox audit logging to Microsoft 365 Group mailboxes, but you can't customize what's being logged (you can't add or remove mailbox actions that are logged for any logon type).

The following table describes the mailbox actions that are logged by default on Microsoft 365 Group mailboxes for each logon type.

Remember, an admin with Full Access permission to a Microsoft 365 Group mailbox is considered a delegate.

MAILBOX ACTION	DESCRIPTION	ADMIN	DELEGATE	OWNER
Create	Creation of a calendar Item. Creating, sending, or receiving a message isn't audited.	✓ *	✓ *	
HardDelete	A message was purged from the Recoverable Items folder.	✓ *	✓ *	✓ *
MoveToDeletedItems	A message was deleted and moved to the Deleted Items folder.	✓ *	✓ *	✓ *
SendAs	A message was sent using the SendAs permission.	✓ *	✓ *	
SendOnBehalf	A message was sent using the SendOnBehalf permission.	✓ *	✓ *	
SoftDelete	A message was permanently deleted or deleted from the Deleted Items folder. Soft-deleted items are moved to the Recoverable Items folder.	✓ *	✓ *	✓ *
Update	A message or its properties was changed.	✓ *	✓ *	✓ *

Verify that default mailbox actions are being logged for each logon type

Mailbox auditing on by defaults adds a new *DefaultAuditSet* property to all mailboxes. The value of this property indicates whether the default mailbox actions (managed by Microsoft) are being audited on the mailbox.

To display the value on user mailboxes or shared mailboxes, replace <MailboxIdentity> with the name, alias, email address, or user principal name (username) of the mailbox and run the following command in Exchange Online PowerShell:

```
Get-Mailbox -Identity <MailboxIdentity> | Format-List DefaultAuditSet
```

To display the value on Microsoft 365 group mailboxes, replace <MailboxIdentity> with the name, alias, or email address of the shared mailbox and run the following command in Exchange Online PowerShell:

```
Get-Mailbox -Identity <MailboxIdentity> -GroupMailbox | Format-List DefaultAuditSet
```

The value `Admin, Delegate, Owner` indicates:

- The default mailbox actions for all three logon types are being audited. This is the only value you'll see on Microsoft 365 Group mailboxes.
- An admin *has not* changed the audited mailbox actions for any logon type on a user mailbox or a shared mailbox. Note this is the default state after mailbox auditing on by default is initially turned on in your organization.

If an admin has ever changed the mailbox actions that are audited for a logon type (by using the *AuditAdmin*, *AuditDelegate*, or *AuditOwner* parameters on the **Set-Mailbox** cmdlet), the property value will be different.

For example, the value `Owner` for the *DefaultAuditSet* property on a user mailbox or shared mailbox indicates:

- The default mailbox actions for the mailbox owner are being audited.
- The audited mailbox actions for the `Delegate` and `Admin` logon types have been changed from the default actions.

A blank value for the *DefaultAuditSet* property indicates the mailbox actions for all three logon types have been changed on the user mailbox or a shared mailbox.

For more information, see the Change or restore mailbox actions logged by default section in this topic

Display the mailbox actions that are being logged on mailboxes

To see the mailbox actions that are currently being logged on user mailboxes or shared mailboxes, replace <MailboxIdentity> with the name, alias, email address, or user principal name (username) of the mailbox, and run one or more of the following commands in Exchange Online PowerShell.

> **NOTE**
> Although you can add the `-GroupMailbox` switch to the following **Get-Mailbox** commands for Microsoft 365 Group mailboxes, don't believe the values that are returned. The default and static mailbox actions that are audited for Microsoft 365 Group mailboxes are described in the Mailbox actions for Microsoft 365 Group mailboxes section earlier in this topic.

Owner actions

```
Get-Mailbox -Identity <MailboxIdentity> | Select-Object -ExpandProperty AuditOwner
```

Delegate actions

```
Get-Mailbox -Identity <MailboxIdentity> | Select-Object -ExpandProperty AuditDelegate
```

Admin actions

```
Get-Mailbox -Identity <MailboxIdentity> | Select-Object -ExpandProperty AuditAdmin
```

Change or restore mailbox actions logged by default

As previously explained, one of the key benefits of having mailbox auditing on by default is: you don't need to manage the mailboxes actions that are audited. Microsoft does this for you and we'll automatically add new mailbox actions to be audited by default as they're released.

However, your organization might be required to audit a different set of mailbox actions for user mailboxes and shared mailboxes. The procedures in this section show you how to change the mailbox actions that are audited for each logon type, and how to revert back to the Microsoft-managed default actions.

> **IMPORTANT**
> If you use the following procedures to customize the mailbox actions that are logged on user mailboxes or shared mailboxes, any new default mailbox actions released by Microsoft will not be automatically audited on those mailboxes. You'll need to manually add any new mailbox actions to your customized list of actions.

Change the mailbox actions to audit

You can use the *AuditAdmin*, *AuditDelegate*, or *AuditOwner* parameters on the **Set-Mailbox** cmdlet to change the mailbox actions that are audited for user mailboxes and shared mailboxes (audited actions for Microsoft 365 group mailboxes can't be customized).

You can use two different methods to specify the mailbox actions:

- *Replace* (overwrite) the existing mailbox actions by using this syntax: `action1,action2,...actionN`.
- *Add or remove* mailbox actions without affecting other existing values by using this syntax: `@{Add="action1","action2",..."actionN"}` or `@{Remove="action1","action2",..."actionN"}`.

This example changes the admin mailbox actions for the mailbox named "Gabriela Laureano" by overwriting the default actions with SoftDelete and HardDelete.

```
Set-Mailbox -Identity "Gabriela Laureano" -AuditAdmin HardDelete,SoftDelete
```

This example adds the MailboxLogin owner action to the mailbox laura@contoso.onmicrosoft.com.

```
Set-Mailbox -Identity laura@contoso.onmicrosoft.com -AuditOwner @{Add="MailboxLogin"}
```

This example removes the MoveToDeletedItems delegate action for the Team Discussion mailbox.

```
Set-Mailbox -Identity "Team Discussion" -AuditDelegate @{Remove="MoveToDeletedItems"}
```

Regardless of the method you use, customizing the audited mailbox actions on user mailboxes or shared mailboxes has the following results:

- For the logon type that you customized, the audited mailbox actions are no longer managed by Microsoft.
- The logon type that you customized is no longer displayed in the *DefaultAuditSet* property value for the mailbox as previously described.

Restore the default mailbox actions

If you customized the mailbox actions that are audited on a user mailbox or a shared mailbox, you can restore the default mailbox actions for one or all logon types by using this syntax:

```
Set-Mailbox -Identity <MailboxIdentity> -DefaultAuditSet <Admin | Delegate | Owner>
```

You can specify multiple *DefaultAuditSet* values separated by commas

Note: The following procedures don't apply to Microsoft 365 Group mailboxes (they're limited to the default actions as described here).

This example restores the default audited mailbox actions for all logon types on the mailbox mark@contoso.onmicrosoft.com.

```
Set-Mailbox -Identity mark@contoso.onmicrosoft.com -DefaultAuditSet Admin,Delegate,Owner
```

This example restores the default audited mailbox actions for the Admin logon type on the mailbox chris@contoso.onmicrosoft.com, but leaves the customized audited mailbox actions for the Delegate and Owner logon types.

```
Set-Mailbox -Identity chris@contoso.onmicrosoft.com -DefaultAuditSet Admin
```

Restoring he default audited mailbox actions for a logon type has the following results:

- The current list of mailbox actions is replaced with the default mailbox actions for the logon type.
- Any new mailbox actions that are released by Microsoft are automatically added to the list of audited actions for the logon type.
- The *DefaultAuditSet* property value for the mailbox is updated to include the restored logon type.

Turn off mailbox auditing on by default for your organization

You can turn off mailbox auditing on by default for your entire organization by running the following command in Exchange Online PowerShell:

```
Set-OrganizationConfig -AuditDisabled $true
```

Turning off mailbox auditing on by default has the following results:

- Mailbox auditing is disabled for your organization.
- From the time you disabled mailbox auditing on by default, no mailbox actions are audited, even if auditing is enabled on a mailbox (the *AuditEnabled* property on the mailbox is **True**).
- Mailbox auditing is not enabled for new mailboxes and setting the *AuditEnabled* property on a new or existing mailbox to **True** will be ignored.
- Any mailbox audit bypass association settings (configured by using the **Set-MailboxAuditBypassAssociation** cmdlet) are ignored.
- Existing mailbox audit records are retained until the audit log age limit for the record expires.

Turn on mailbox auditing on by default

To turn mailbox auditing back on for your organization, run the following command in Exchange Online PowerShell:

```
Set-OrganizationConfig -AuditDisabled $false
```

Bypass mailbox audit logging

Currently, you can't disable mailbox auditing for specific mailboxes when mailbox auditing on by default is turned on in your organization. For example, setting the *AuditEnabled* mailbox property to **False** is ignored.

However, you can still use the **Set-MailboxAuditBypassAssociation** cmdlet in Exchange Online PowerShell to

prevent *any and all* mailbox actions by the specified users from being logged, regardless where the actions occur. For example:

- Mailbox owner actions performed by the bypassed users aren't logged.
- Delegate actions performed by the bypassed users on other users' mailboxes (including shared mailboxes) aren't logged.
- Admin actions performed by the bypassed users aren't logged.

To bypass mailbox audit logging for a specific user, replace <MailboxIdentity> with the name, email address, alias, or user principal name (username) of the user and run the following command:

```
Set-MailboxAuditBypassAssociation -Identity <MailboxIdentity> -AuditByPassEnabled $true
```

To verify that auditing is bypassed for the specified user, run the following command:

```
Get-MailboxAuditBypassAssociation -Identity <MailboxIdentity> | Format-List AuditByPassEnabled
```

The value **True** indicates that mailbox audit logging is bypassed for the user.

More information

- Although mailbox audit logging on by default is enabled for all organizations, only users with E5 licenses will return mailbox audit log events in audit log searches in the Security & Compliance Center or via the Office 365 Management Activity API **by default**.

  To retrieve mailbox audit log entries for users without E5 licenses, you can:

  - Manually enable mailbox auditing on individual mailboxes (run the command, `Set-Mailbox -Identity <MailboxIdentity> -AuditEnabled $true`). After you do this, you can use audit log searches in the Security & Compliance Center or via the Office 365 Management Activity API.

    > **NOTE**
    > If mailbox auditing already appears to be enabled on the mailbox, but your searches return no results, change the value of the *AuditEnabled* parameter to `$false` and then back to `$true`.

  - Use the following cmdlets in Exchange Online PowerShell:
    - Search-MailboxAuditLog to search the mailbox audit log for specific users.
    - New-MailboxAuditLogSearch to search the mailbox audit log for specific users and to have the results sent via email to specified recipients.
  - Use the Exchange admin center (EAC) in Exchange Online to do the following actions:
    - Export mailbox audit logs
    - Run a non-owner mailbox access report
- By default, mailbox audit log records are retained for 90 days before they're deleted. You can change the age limit for audit log records by using the *AuditLogAgeLimit* parameter on the **Set-Mailbox** cmdlet in Exchange Online PowerShell. However, increasing this value doesn't allow you to search for events that are older than 90 days in the audit log.

  If you increase the age limit, you need to use the Search-MailboxAuditLog cmdlet in Exchange Online

PowerShell to search the user's mailbox audit log for records that are older than 90 days.

- If you've changed the *AuditLogAgeLimit* property for a mailbox prior to mailbox auditing on by default being turned on for organization, the mailbox's existing audit log age limit isn't changed. In other words, mailbox auditing on by default doesn't affect the current age limit for mailbox audit records.
- To change the *AuditLogAgeLimit* value on a Microsoft 365 Group mailbox, you need to include the `-GroupMailbox` switch in the **Set-Mailbox** command.
- Mailbox audit log records are stored in a subfolder (named *Audits*) in the Recoverable Items folder in each user's mailbox. Keep the following things in mind about mailbox audit records and the Recoverable Items folder:
  - Mailbox audit records count against the storage quota of the Recoverable Items folder, which is 30 GB by default (the warning quota is 20 GB). The storage quota is automatically increased to 100 GB (with a 90 GB warning quota) when:
    - A hold is placed on a mailbox.
    - The mailbox is assigned to a retention policy in the Compliance Center.
  - Mailbox audit records also count against the folder limit for the Recoverable Items folder. A maximum of 3 million items (audit records) can be stored in the Audits subfolder.

    > **NOTE**
    > It's unlikely that mailbox auditing on by default will impact the storage quota or the folder limit for the Recoverable Items folder.

    - You can run the following command in Exchange Online PowerShell to display the size and number of items in the Audits subfolder in the Recoverable Items folder:

      ```
      Get-MailboxFolderStatistics -Identity <MailboxIdentity> -FolderScope RecoverableItems |
      Where-Object {$_.Name -eq 'Audits'} | Format-List FolderPath,FolderSize,ItemsInFolder
      ```

    - You can't directly access an audit log record in the Recoverable Items folder; instead, you use the **Search-MailboxAuditLog** cmdlet or search the audit log to find and view mailbox audit records.
- If a mailbox is placed on hold or assigned to a retention policy in the Compliance Center, audit log records are still retained for the duration that's defined by the mailbox's *AuditLogAgeLimit* property (90 days by default). To retain audit log records longer for mailboxes on hold, you need to increase mailbox's *AuditLogAgeLimit* value.
- In a multi-geo environment, cross-geo mailbox auditing is not supported. For example, if a user is assigned permissions to access a shared mailbox in a different geo location, mailbox actions performed by that user are not logged in the mailbox audit log of the shared mailbox.

Advanced Audit in Microsoft 365

2/18/2021 • 12 minutes to read • Edit Online

The unified auditing functionality in Microsoft 365 provides organizations with visibility into many types of audited activities across many different services in Microsoft 365. Advanced Audit helps organizations to conduct forensic and compliance investigations by increasing audit log retention required to conduct an investigation, providing access to crucial events that help determine scope of compromise, and faster access to Office 365 Management Activity API.

> **NOTE**
> Advanced Audit is available for organizations with an Office 365 E5/G5 or Microsoft 365 Enterprise E5/G5 subscription. Additionally, a Microsoft 365 E5 Compliance or E5 eDiscovery and Audit add-on license can be assigned to users when per-user licensing is required for Advanced Audit features as is the case for long-term retention of audit logs and access to crucial events for investigations. For more information about licensing, see Microsoft 365 licensing guidance for security & compliance.

This article provides an overview of Advanced Audit capabilities and shows you how to set up users for Advanced Audit.

Long-term retention of audit logs

Advanced Audit retains all Exchange, SharePoint, and Azure Active Directory audit records for one year. This is accomplished by a default audit log retention policy that retains any audit record that contains the value of **Exchange**, **SharePoint**, or **AzureActiveDirectory** for the **Workload** property (which indicates the service in which the activity occurred) for one year. Retaining audit records for longer periods can help with on-going forensic or compliance investigations. For more information, see the "Default audit log retention policy" section in Manage audit log retention policies.

We're also releasing the capability to retain audit logs for 10 years. The 10-year retention of audit logs helps support long running investigations and respond to regulatory, legal, and internal obligations.

> **NOTE**
> Retaining audit logs for 10 years will require an additional add-on license. This new license will be available in early 2021. For more information, see the FAQs for Advanced Audit section in this article.

Audit log retention policies

All audit records generated in other services that aren't covered by the default audit log retention policy (described in the previous section) are retained for 90 days. But you can create customized audit log retention policies to retain other audit records for longer periods of time up to 10 years. You can create a policy to retain audit records based on one or more of the following criteria:

- The Microsoft 365 service where the audited activities occur.
- Specific audited activities.
- The user who performs an audited activity.

You can also specify how long to retain audit records that match the policy and a priority level so that specific policies will take priority over other policies. Also note that any custom audit log retention policy will take

precedence over the default audit retention policy in case you need retain Exchange, SharePoint, or Azure Active Directory audit records for less than a year (or for 10 years) for some or all users in your organization. For more information, see Manage audit log retention policies.

Access to crucial events for investigations

Advanced Audit helps organizations to conduct forensic and compliance investigations by providing access to crucial events such as when mail items were accessed, or when mail items were replied to and forwarded, and when and what a user searched for in Exchange Online and SharePoint Online. These crucial events can help you investigate possible breaches and determine the scope of compromise. Advanced Auditing provides the following crucial events:

- MailItemsAccessed
- Send
- SearchQueryInitiatedExchange
- SearchQueryInitiatedSharePoint

MailItemsAccessed

The MailItemsAccessed event is a mailbox auditing action and is triggered when mail data is accessed by mail protocols and mail clients. The MailItemsAccessed action can help investigators identify data breaches and determine the scope of messages that may have been compromised. If an attacker gained access to email messages, the MailItemsAccessed action will be triggered even if there is no explicit signal that messages were actually read (in other words, the type of access such as a bind or sync is recorded in the audit record).

The MailItemsAccessed mailbox action replaces MessageBind in mailbox auditing logging in Exchange Online and provides these improvements:

- MessageBind was only configurable for AuditAdmin user logon type; it did not apply to delegate or owner actions. MailItemsAccessed applies to all logon types.
- MessageBind only covered access by a mail client. It didn't apply to sync activities. MailItemsAccessed events are triggered by both bind and sync access types.
- MessageBind actions would trigger the creation of multiple audit records when the same email message was accessed, which resulted in auditing "noise". In contrast, MailItemsAccessed events are aggregated into fewer audit records.

For information about audit records for MailItemsAccessed activities, see Use Advanced Audit to investigate compromised accounts.

To search for MailItemsAccessed audit records, you can search for the **Accessed mailbox items** activity in the **Exchange mailbox activities** drop-down list in the audit log search tool in the Microsoft 365 compliance center.

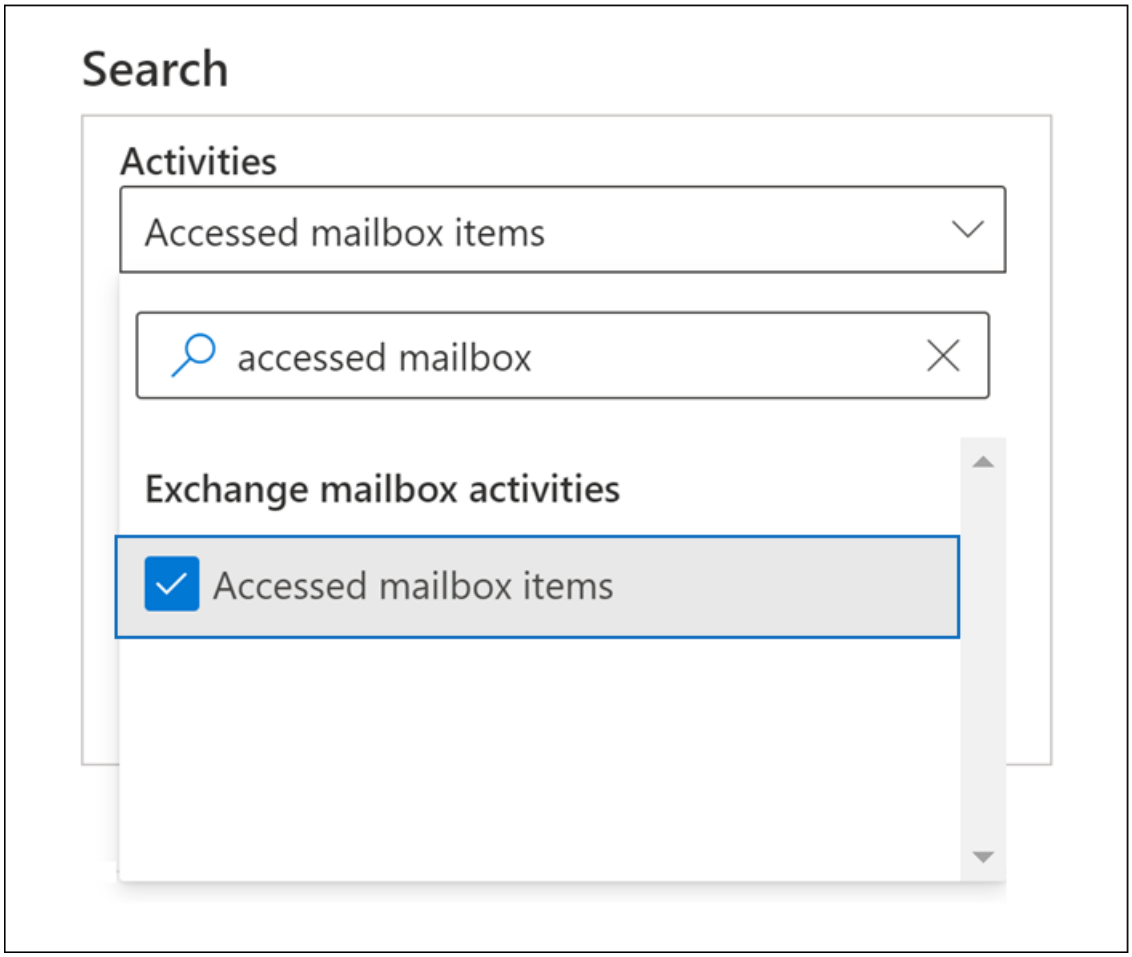

You can also run the Search-UnifiedAuditLog -Operations MailItemsAccessed or Search-MailboxAuditLog -Operations MailItemsAccessed commands in Exchange Online PowerShell.

Send

The Send event is also a mailbox auditing action and is triggered when a user performs one of the following actions:

- Sends an email message
- Replies to an email message
- Forwards an email message

Investigators can use the Send event to identify email sent from a compromised account. The audit record for a Send event contains information about the message, such as when the message was sent, the InternetMessage ID, the subject line, and if the message contained attachments. This auditing information can help investigators identify information about email messages sent from a compromised account or sent by an attacker. Additionally, investigators can use a Microsoft 365 eDiscovery tool to search for the message (by using the subject line or message ID) to identify the recipients the message was sent to and the actual contents of the sent message.

To search for Send audit records, you can search for the **Sent message** activity in the **Exchange mailbox activities** drop-down list in the audit log search tool in the Microsoft 365 compliance center.

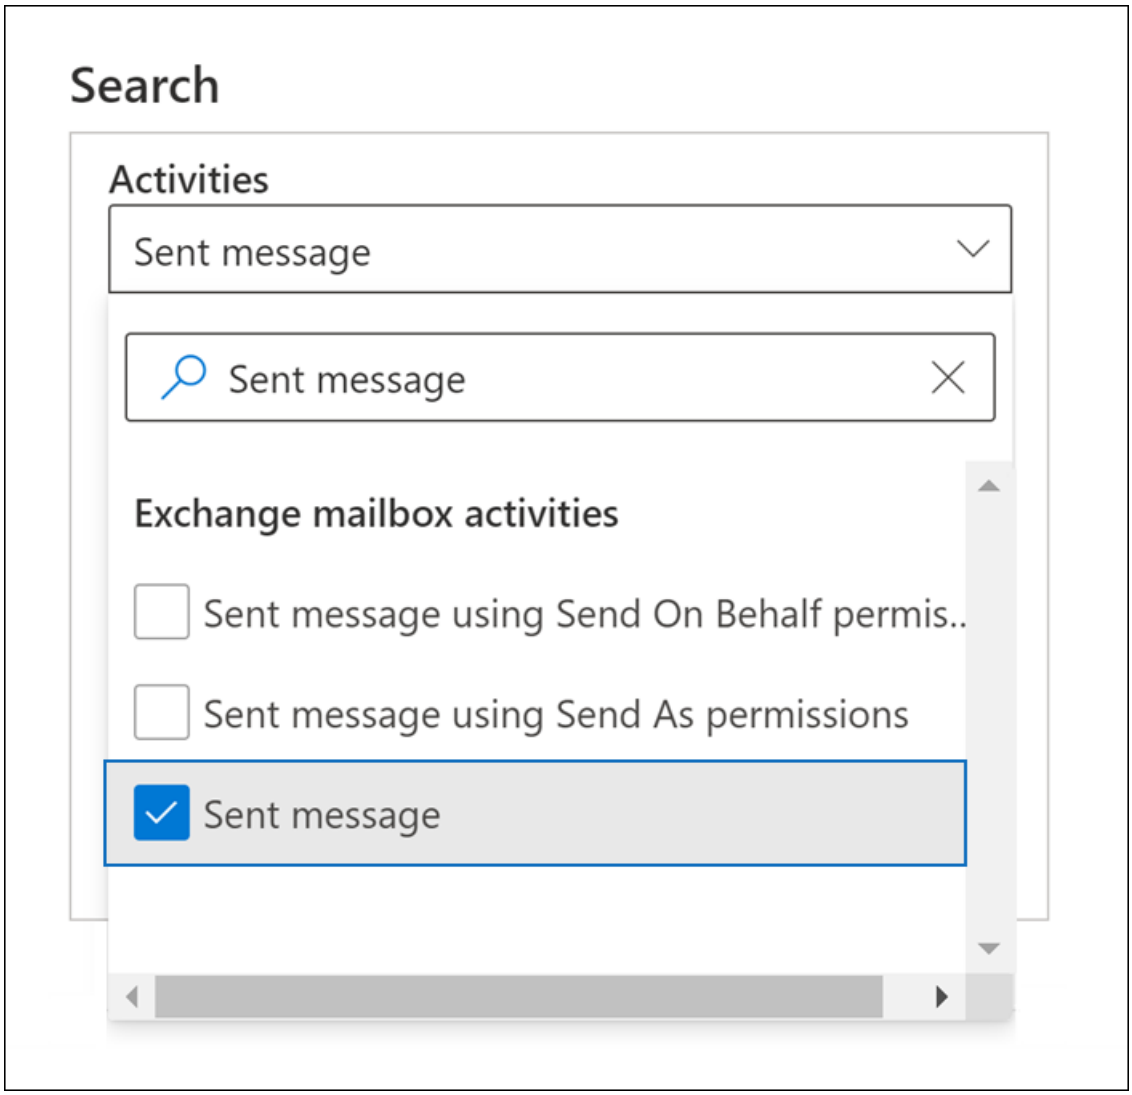

You can also run the Search-UnifiedAuditLog -Operations Send or Search-MailboxAuditLog -Operations Send commands in Exchange Online PowerShell.

SearchQueryInitiatedExchange

The SearchQueryInitiatedExchange event is triggered when a person uses Outlook to search for items in a mailbox. Events are triggered when searches are performed in the following Outlook environments:

- Outlook (desktop client)
- Outlook on the web (OWA)
- Outlook for iOS
- Outlook for Android
- Mail app for Windows 10

Investigators can use the SearchQueryInitiatedExchange event to determine if an attacker who may have compromised an account looked for or tried to access sensitive information in the mailbox. The audit record for a SearchQueryInitiatedExchange event contains information such as the actual text of the search query. The audit record also indicates the Outlook environment the search was performed in. By looking at the search queries that an attacker may have performed, an investigator can better understand the intent of the email data that was searched for.

To search for SearchQueryInitiatedExchange audit records, you can search for the **Performed email search** activity in the **Search activities** drop-down list in the audit log search tool in the compliance center.

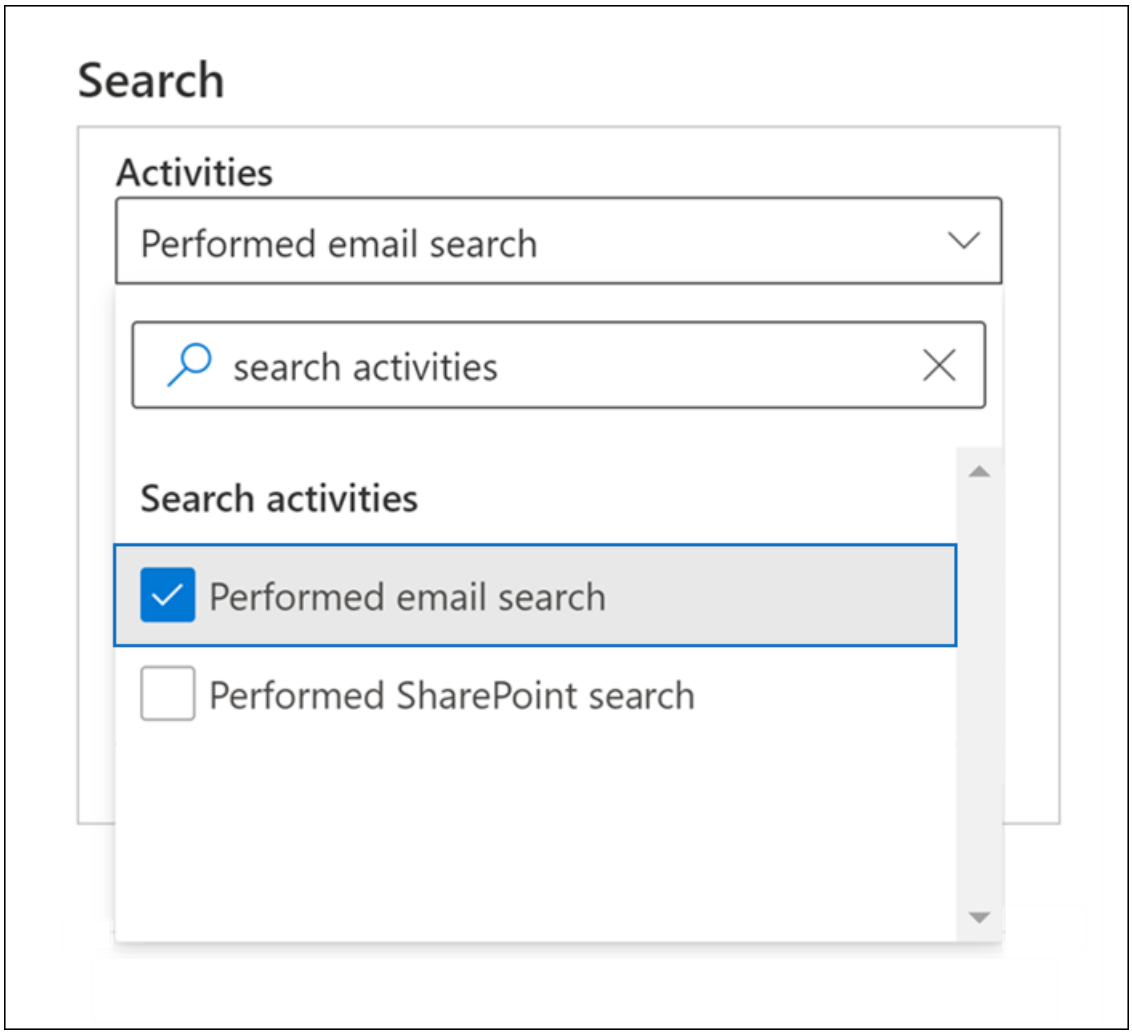

You can also run the Search-UnifiedAuditLog -Operations SearchQueryInitiatedExchange in Exchange Online PowerShell.

NOTE

You must run the following command in Exchange Online PowerShell so that SearchQueryInitiatedExchange events (performed by the specified E5 user) are included in audit log search results: `Set-Mailbox <user identity> -AuditOwner @{Add="SearchQueryInitiated"}`.

In a multi-geo environment, you must run the **Set-Mailbox** command in the forest where the user's mailbox is located. To identify the user's mailbox location, run the following command: `Get-Mailbox <user identity> | FL MailboxLocations`. If the `Set-Mailbox -AuditOwner @{Add="SearchQueryInitiated"}` command was previously run in the forest that's different than the one the user's mailbox is located in, then you must remove the SearchQueryInitiated value from the user's mailbox (by running `Set-Mailbox -AuditOwner @{Remove="SearchQueryInitiated"}`) and then add it to the user's mailbox in the forest where the user's mailbox is located.

SearchQueryInitiatedSharePoint

Similar to searching for mailbox items, the SearchQueryInitiatedSharePoint event is triggered when a person searches for items in SharePoint. Events are triggered when searches are performed in the following types of SharePoint sites:

- Home sites
- Communication sites
- Hub sites

- Sites associated with Microsoft Teams

Investigators can use the SearchQueryInitiatedSharePoint event to determine if an attacker tried to find (and possibly accessed) sensitive information in SharePoint. The audit record for a SearchQueryInitiatedSharePoint event contains also contains the actual text of the search query. The audit record also indicates the type of SharePoint site that was searched. By looking at the search queries that an attacker may have performed, an investigator can better understand the intent and scope of the file data being searched for.

To search for SearchQueryInitiatedSharePoint audit records, you can search for the **Performed SharePoint search** activity in the **Search activities** drop-down list in the audit log search tool in the compliance center.

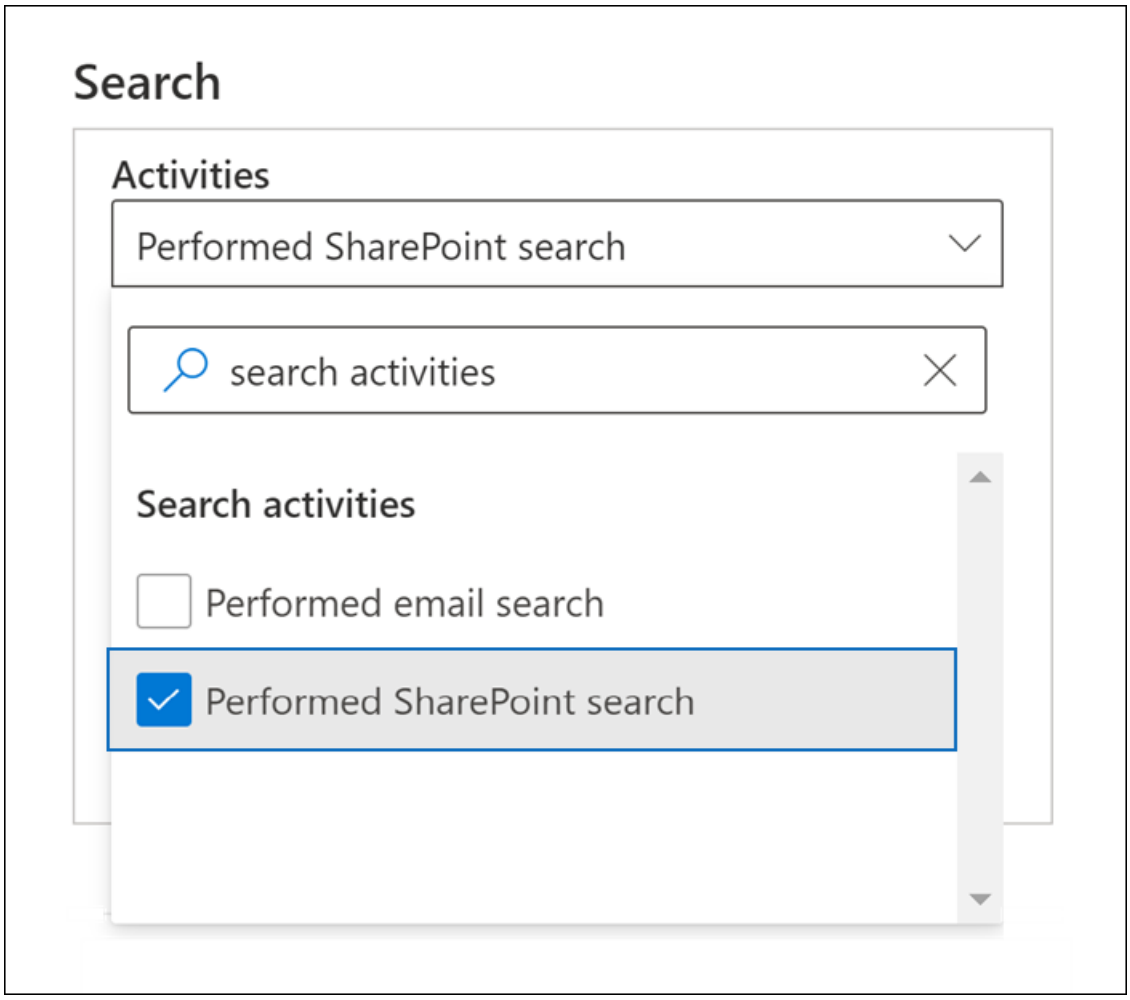

You can also run the Search-UnifiedAuditLog -Operations SearchQueryInitiatedSharePoint in Exchange Online PowerShell.

NOTE

You must run the following command in Exchange Online PowerShell so that SearchQueryInitiatedExchange events (performed by the specified E5 user) are included in audit log search results: `Set-Mailbox <user identity> -AuditOwner @{Add="SearchQueryInitiated"}`.

In a multi-geo environment, you must run the **Set-Mailbox** command in the forest where the user's mailbox is located. To identify the user's mailbox location, run the following command: `Get-Mailbox <user identity> | FL MailboxLocations`. If the `Set-Mailbox -AuditOwner @{Add="SearchQueryInitiated"}` command was previously run in the forest that's different than the one the user's mailbox is located in, then you must remove the SearchQueryInitiated value from the user's mailbox (by running `Set-Mailbox -AuditOwner @{Remove="SearchQueryInitiated"}`) and then add it to the user's mailbox in the forest where the user's mailbox is located.

High-bandwidth access to the Office 365 Management Activity API

Organizations that access auditing logs through the Office 365 Management Activity API were restricted by throttling limits at the publisher level. This means that for a publisher pulling data on behalf of multiple customers, the limit was shared by all those customers.

With the release of Advanced Audit, we're moving from a publisher-level limit to a tenant-level limit. The result is that each organization will get their own fully allocated bandwidth quota to access their auditing data. The bandwidth is not a static, predefined limit but is modeled on a combination of factors including the number of seats in the organization and that E5 organizations will get more bandwidth than non-E5 organizations.

All organizations are initially allocated a baseline of 2,000 requests per minute. This limit will dynamically increase depending on an organization's seat count and their licensing subscription. E5 organizations will get about twice as much bandwidth as non-E5 organizations. There will also be cap on the maximum bandwidth to protect the health of the service.

For more information, see the "API throttling" section in Office 365 Management Activity API reference.

Set up Advanced Audit for users

Advanced Audit features such as the ability to log crucial events such as MailItemsAccessed and Send require an appropriate E5 license assigned to users. Additionally, the Advanced Auditing app/service plan must be enabled for those users. To verify that the Advanced Auditing app is assigned to users, perform the following steps for each user:

1. In the Microsoft 365 admin center, go to **Users** > **Active users**, and select a user.
2. On the user properties flyout page, click **Licenses and apps**.
3. In the **Licenses** section, verify that the user is assigned an E5 license.
4. Expand the **Apps** section, and verify that the **Microsoft 365 Advanced Auditing** checkbox is selected.
5. If the checkbox isn't selected, select it, and then click **Save changes.**

   The logging of audit records for MailItemsAccessed, Send, and other crucial events for the user will begin within 24 hours.

For organizations that assign licenses to groups of users by using group-based licensing, you have to turn off the licensing assignment for Microsoft 365 Advanced Auditing for the group. After you save your changes, verify that Microsoft 365 Advanced Auditing is turned off for the group. Then turn the licensing assignment for the group back on. For instructions about group-based licensing, see Assign licenses to users by group membership in Azure Active Directory.

Also, if you have customized the mailbox actions that are logged on user mailboxes or shared mailboxes, new default mailbox actions such as MailItemsAccessed will not be automatically audited on those mailboxes. For information about changing the mailbox actions that are audited for each logon type, see the "Change or restore mailbox actions logged by default" section in Manage mailbox auditing.

FAQs for Advanced Audit

Does every user need an E5 license to benefit from Advanced Audit?

To benefit from user-level Advanced Audit capabilities, a user needs to be assigned an E5 license. There are some capabilities that will check for the appropriate license to expose the feature for the user. For example, if you're trying to retain the audit records for a user who isn't assigned an E5 license for longer than 90 days, the system will return an error message.

My organization has an E5 subscription, do I need to do anything to get access to audit records for crucial events?

For eligible customers and users that are assigned the appropriate license, there is no action to get access to crucial auditing events.

When will the new 10-year audit log retention add-on license be available?

The new 10-year audit log retention add-on will be available for purchase by customers with E5 subscriptions in early 2021.

What happens to my organization's audit log data if I create 10-year audit log retention policy the feature is released to general availability but before the required add-on license is available in early 2021?

Any audit log data covered by a 10-year audit log retention policy that you create after general availability will be retained for 10 years. When the 10-year audit log retention add-on license is available in early 2021, you will need to purchase add-on licenses for users who's audit data is being retained by an existing 10-year audit retention policy. Also, once the add-on license is available in early 2021, the appropriate licensing will be enforced when you create new 10-year audit log retention policies.

Are the new events in Advanced Audit available in the Office 365 Management Activity API?

Yes. As long as audit records are generated for users with the appropriate license, you'll be able to access these records via the Office 365 Management Activity API.

Does higher bandwidth mean better latency or higher SLA?

At this time, high bandwidth provides a better pipeline, especially for organizations with a high volume of auditing signals and significant consumption patterns. More bandwidth can lead to better latency. But there isn't an SLA associated with high bandwidth. Standard latencies are documented, and these latencies don't change with the release of Advanced Audit.

Manage audit log retention policies

2/18/2021 • 7 minutes to read • Edit Online

You can create and manage audit log retention policies in the Security & Compliance Center. Audit log retention policies are part of the new Advanced Audit capabilities in Microsoft 365. An audit log retention policy lets you specify how long to retain audit logs in your organization. You can retain audit logs for up to 10 years. You can create policies based on the following criteria:

- All activities in one or more Microsoft 365 services
- Specific activities (in a Microsoft 365 service) performed by all users or by specific users
- A priority level that specifies which policy takes precedence in you have multiple policies in your organization

Default audit log retention policy

Advanced Audit in Microsoft 365 provides a default audit log retention policy for all organizations. This policy retains all Exchange, SharePoint, and Azure Active Directory audit records for one year. This default policy retains audit records that contain the value of **AzureActiveDirectory**, **Exchange**, or **SharePoint** for the **Workload** property (which is the service in which the activity occurred). The default policy can't be modified. See the More information section in this article for a list of record types for each workload that are included in the default policy.

> **NOTE**
> The default audit log retention policy only applies to audit records for activity performed by users who are assigned an Office 365 or Microsoft 365 E5 license or have a Microsoft 365 E5 Compliance or E5 eDiscovery and Audit add-on license. If you have non-E5 users or guest users in your organization, their corresponding audit records are retained for 90 days.

Before you create an audit log retention policy

- You have to be assigned the Organization Configuration role in the Security & Compliance Center to create or modify an audit retention policy.
- You can have a maximum of 50 audit log retention policies in your organization.
- To retain an audit log for longer than 90 days, the user who generated the audit log must be assigned an Office 365 E5 or Microsoft 365 E5 license or have a Microsoft 365 E5 Compliance or E5 eDiscovery and Audit add-on license.
- All custom audit log retention policies (created by your organization) take priority over the default retention policy. For example, if you create an audit log retention policy for Exchange mailbox activity that has a retention period that's shorter than one year, audit records for Exchange mailbox activities will be retained for the shorter duration specified by the custom policy.

Create an audit log retention policy

1. Go to https://compliance.microsoft.com and sign in with a user account that's assigned the Organization Configuration role on the Permissions page in the Security & Compliance Center.
2. In the left pane of the Microsoft 365 compliance center, click **Show all**, and then click **Audit**.

3. Click the **Audit retention policies** tab.
4. Click **Create audit retention policy**, and then complete the following fields on the flyout page:

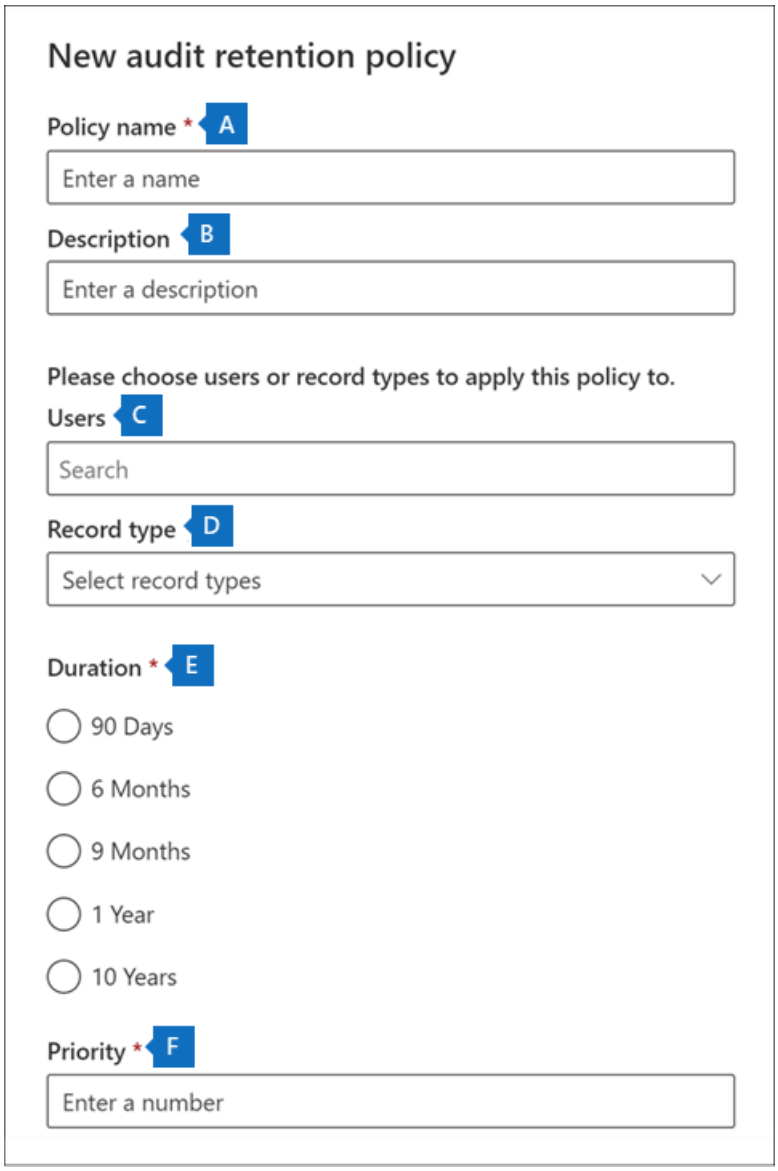

a. **Policy name:** The name of the audit log retention policy. This name must be unique in your organization, and it can't be change after the policy is created.

b. **Description:** Optional, but helpful to provide information about the policy, such as the record type or workload, users specified in the policy, and the duration.

c. **Users:** Select one or more users to apply the policy to. If you leave this box blank, then the policy will apply to all users. If you leave the **Record type** blank, then you must select a user.

d. **Record type:** The audit record type the policy applies to. If you leave this property blank, you must select a user in the **Users** box. You can select a single record type or multiple record types:

- If you select a single record type, the **Activities** field is dynamically displayed. You can use the drop-down list to select activities from the selected record type to apply the policy to. If you don't choose specific activities, the policy will apply to all activities of the selected record type.
- If you select multiple record types, you don't have the ability to select activities. The policy will apply to all activities of the selected record types.

e. **Duration:** The amount of time to retain the audit logs that meet the criteria of the policy.

f. **Priority:** This value determines the order in which audit log retention policies in your organization

are processed. A higher value indicates a higher priority. For example, a policy with a priority value of **5** would take priority over a policy with a priority value of **0**. As previously explained, any custom audit log retention policy takes priority over the default policy for your organization.

5. Click **Save** to create the new audit log retention policy.

   The new policy is displayed in the list on the **Audit retention policies** tab.

Manage audit log retention policies

Audit log retention policies are listed on the **Audit retention policies** tab (also called the *dashboard*). You can use the dashboard to view, edit, and delete audit retention policies.

View policies in the dashboard

Audit log retention policies are listed in the dashboard. One advantage of viewing policies in the dashboard is that you can click the **Priority** column to list the policies in the priority in which they are applied. As previously explained, a higher value indicates a higher priority.

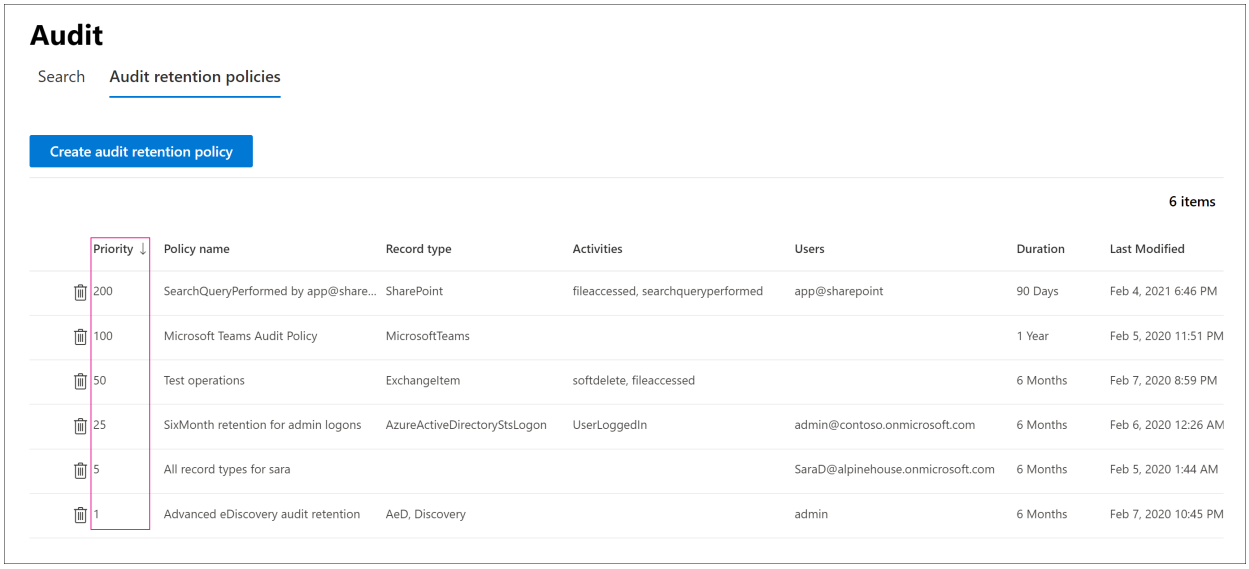

You can also select a policy to display its settings on the flyout page.

> **NOTE**
> The default audit log retention policy for your organization isn't displayed in the dashboard.

Edit policies in the dashboard

To edit a policy, select it to display the flyout page. You can modify one or more setting and then save your changes.

> **IMPORTANT**
> If you use the **New-UnifiedAuditLogRetentionPolicy** cmdlet, it's possible to create an audit log retention policy for record types or activities that aren't available in the **Create audit retention policy** tool in the dashboard. In this case, you won't be able to edit the policy (for example, change the retention duration or add and remove activities) from the **Audit retention policies** dashboard. You'll only be able to view and delete the policy in the compliance center. To edit the policy, you'll have to use the Set-UnifiedAuditLogRetentionPolicy cmdlet in Security & Compliance Center PowerShell.
>
> **Tip:** A message is displayed at the top of the flyout page for policies that have to be edited using PowerShell.

Delete policies in the dashboard

To delete a policy, click the **Delete**  icon and then confirm that you want to delete the policy. The policy is removed from the dashboard, but it might take up to 30 minutes for the policy to be removed from your organization.

Create and manage audit log retention policies in PowerShell

You can also use Security & Compliance Center PowerShell to create and manage audit log retention policies. One reason to use PowerShell is to create a policy for a record type or activity that isn't available in the UI.

Create an audit log retention policy in PowerShell

Follow these steps to create an audit log retention policy in PowerShell:

1. Connect to Security & Compliance Center PowerShell.
2. Run the following command to create an audit log retention policy.

   ```
   New-UnifiedAuditLogRetentionPolicy -Name "Microsoft Teams Audit Policy" -Description "One year
   retention policy for all Microsoft Teams activities" -RecordTypes MicrosoftTeams -RetentionDuration
   TenYears -Priority 100
   ```

   This example creates an audit log retention policy named "Microsoft Teams Audit Policy" with these settings:

   - A description of the policy.
   - Retains all Microsoft Teams activities (as defined by the *RecordType* parameter).
   - Retains Microsoft Teams audit logs for 10 years.
   - A priority of 100.

Here's another example of creating an audit log retention policy. This policy retains audit logs for the "User logged in" activity for six months for the user admin@contoso.onmicrosoft.com.

```
New-UnifiedAuditLogRetentionPolicy -Name "SixMonth retention for admin logons" -RecordTypes
AzureActiveDirectoryStsLogon -Operations UserLoggedIn -UserIds admin@contoso.onmicrosoft.com -
RetentionDuration SixMonths -Priority 25
```

For more information, see New-UnifiedAuditLogRetentionPolicy.

View policies in PowerShell

Use the Get-UnifiedAuditLogRetentionPolicy cmdlet in Security & Compliance Center PowerShell to view audit log retention policies.

Here's a sample command to display the settings for all audit log retention policies in your organization. This command sorts the policies from the highest to lowest priority.

```
Get-UnifiedAuditLogRetentionPolicy | Sort-Object -Property Priority -Descending | FL
Priority,Name,Description,RecordTypes,Operations,UserIds,RetentionDuration
```

> **NOTE**
>
> The **Get-UnifiedAuditLogRetentionPolicy** cmdlet doesn't return the default audit log retention policy for your organization.

Edit policies in PowerShell

Use the Set-UnifiedAuditLogRetentionPolicy cmdlet in Security & Compliance Center PowerShell to edit an existing audit log retention policy.

Delete policies in PowerShell

Use the Remove-UnifiedAuditLogRetentionPolicy cmdlet in Security & Compliance Center PowerShell to delete an audit log retention policy. It might take up to 30 minutes for the policy to be removed from your organization.

More information

As previously stated, audit records for operations in Azure Active Directory, Exchange, and SharePoint are retained for one year by default. The following table lists all the record types (for each of these services) included in the default audit log retention policy. This means that audit logs for any operation with this record type are retained for one year unless a custom audit log retention policy takes precedence for a specific record type, operation, or user. The Enum value (which is displayed as the value for the RecordType property in an audit record) for each record type is shown in parentheses.

AZUREACTIVEDIRECTORY	EXCHANGE	SHAREPOINT
AzureActiveDirectory (8)	ExchangeAdmin (1)	ComplianceDLPSharePoint (11)
AzureActiveDirectoryAccountLogon (9)	ExchangeItem (2)	ComplianceDLPSharePointClassification (33)
AzureActiveDirectoryStsLogon (15)	Campaign (62)	Project (35)
	ComplianceDLPExchange (13)	SharePoint (4)
	ComplianceSupervisionExchange (68)	SharePointCommentOperation (37)
	CustomerKeyServiceEncryption (69)	SharePointContentTypeOperation (55)
	ExchangeAggregatedOperation (19)	SharePointFieldOperation (56)
	ExchangeItemAggregated (50)	SharePointFileOperation (6)
	ExchangeItemGroup (3)	SharePointListOperation (36)
	InformationBarrierPolicyApplication (53)	SharePointSharingOperation (14)

Use Advanced Audit to investigate compromised accounts

2/18/2021 • 10 minutes to read • Edit Online

A compromised user account (also called an *account takeover*) is a type of attack when an attacker gains access to a user account and operates as the user. These types of attacks sometimes cause more damage than the attacker may have intended. When investigating compromised email accounts, you have to assume that more mail data was compromised than may be indicated by tracing the attacker's actual presence. Depending on the type of data in email messages, you have to assume that sensitive information was compromised or face regulatory fines unless you can prove that sensitive information wasn't exposed. For example, HIPAA-regulated organizations face significant fines if there is evidence that patient health information (PHI) was exposed. In these cases, attackers are unlikely to be interested in PHI, but organizations still must report data breaches unless they can prove otherwise.

To help you with investigating compromise email accounts, we're now auditing accesses of mail data by mail protocols and clients with the *MailItemsAccessed* mailbox auditing action. This new audited action will help investigators better understand email data breaches and help you identify the scope of compromises to specific mail items that may been compromised. The goal of using this new auditing action is forensics defensibility to help assert that a specific piece of mail data was not compromised. If an attacker gained access to a specific piece of mail, Exchange Online audits the event even though there is no indication that the mail item was actually read.

The MailItemsAccessed mailbox auditing action

The new MailItemsAccessed action is part of the new Advanced Audit functionality. It's part of Exchange mailbox auditing and is enabled by default for users that are assigned an Office 365 or Microsoft 365 E5 license or for organizations with a Microsoft 365 E5 Compliance add-on subscription.

The MailItemsAccessed mailbox auditing action covers all mail protocols: POP, IMAP, MAPI, EWS, Exchange ActiveSync, and REST. It also covers both types of accessing mail: *sync* and *bind*.

Auditing sync access

Sync operations are only recorded when a mailbox is accessed by a desktop version of the Outlook client for Windows or Mac. During the sync operation, these clients typically download a large set of mail items from the cloud to a local computer. The audit volume for sync operations is huge. So, instead of generating an audit record for each mail item that's synched, we just generate an audit event for the mail folder containing items that were synched. This makes the assumption that *all* mail items in the synched folder have been compromised. The access type is recorded in the OperationProperties field of the audit record.

See step 2 in the Use MailItemsAccessed audit records for forensic investigations section for an example of displaying the sync access type in an audit record.

Auditing bind access

A bind operation is an individual access to an email message. For bind access, the InternetMessageId of individual messages will be recorded in the audit record. The MailItemsAccessed audit action records bind operations and then aggregates into a single audit record. All bind operations that occur within a 2-minute interval are aggregated in a single audit record in the Folders field within the AuditData property. Each message that was accessed is identified by its InternetMessageId. The number of bind operations that were aggregated in the record is displayed in the OperationCount field in the AuditData property.

See step 4 in the Use MailItemsAccessed audit records for forensic investigations section for an example of displaying the bind access type in an audit record.

Throttling of MailItemsAccessed audit records

If more than 1,000 MailItemsAccessed audit records are generated in less than 24 hours, Exchange Online will stop generating auditing records for MailItemsAccessed activity. When a mailbox is throttled, MailItemsAccessed activity will not be logged for 24 hours after the mailbox was throttled. If this occurs, there's a potential that mailbox could have been compromised during this period. The recording of MailItemsAccessed activity will be resumed following a 24-hour period.

Here's a few things to keep in mind about throttling:

- Less than 1% of all mailboxes in Exchange Online are throttled
- When a mailbox is throttling, only audit records for MailItemsAccessed activity are not audited. Other mailbox auditing actions aren't affected.
- Mailboxes are throttled only for Bind operations. Audit records for sync operations are not throttled.
- If a mailbox is throttled, you can probably assume there was MailItemsAccessed activity that wasn't recorded in the audit logs.

See step 1 in the Use MailItemsAccessed audit records for forensic investigations section for an example of displaying the IsThrottled property in an audit record.

Use MailItemsAccessed audit records for forensic investigations

Mailbox auditing generates audit records for access to email messages so that you can be confident that email messages haven't been compromised. For this reason, in circumstances where we're not certain that some data has been accessed, we assume that it has by recording all mail access activity.

Using MailItemsAccessed audit records for forensics purposes is typically performed after a data breach has been resolved and the attacker has been evicted. To begin your investigation, you should identify the set of mailboxes that they have been compromised and determine the time frame when attacker had access to mailboxes in your organization. Then, you can use the **Search-UnifiedAuditLog** or **Search-MailboxAuditLog** cmdlets in Exchange Online PowerShell to search audit records that correspond to the data breach.

You can run one of the following commands to search for MailItemsAccessed audit records:

Unified audit log

```
Search-UnifiedAuditLog -StartDate 01/06/2020 -EndDate 01/20/2020 -UserIds <user1,user2> -Operations
MailItemsAccessed -ResultSize 1000
```

Mailbox audit log

```
Search-MailboxAuditLog -Identity <user> -StartDate 01/06/2020 -EndDate 01/20/2020 -Operations
MailItemsAccessed -ResultSize 1000 -ShowDetails
```

> **TIP**
> One primary difference between these two cmdlets in that you can use the **Search-UnifiedAuditLog** cmdlet to search for audit records for activity performed by one or more users. That's because *UserIds* is a multi-value parameter. The **Search-MailboxAuditLog** cmdlet searches the mailbox audit log for a single user.

Here are the steps for using MailItemsAccessed audit records to investigate a compromised user attack. Each step shows the command syntax for the **Search-UnifiedAuditLog** or **Search-MailboxAuditLog** cmdlets.

1. Check whether the mailbox has been throttled. If so, this would mean that some mailbox auditing records would not have been logged. In the case that any audit records have the "IsThrottled" is "True," you should assume that for a 24-hour period afterwards that record was generated, that any access to the mailbox was not audited and that all mail data has been compromised.

   To search for MailItemsAccessed records where the mailbox was throttled, run the following command:

   ### Unified audit log

   ```
   Search-UnifiedAuditLog -StartDate 01/06/2020 -EndDate 01/20/2020 -UserIds <user1,user2> -Operations
   MailItemsAccessed -ResultSize 1000 | Where {$_.AuditData -like '*"IsThrottled","Value":"True"*'} | FL
   ```

   ### Mailbox audit log

   ```
   Search-MailboxAuditLog -StartDate 01/06/2020 -EndDate 01/20/2020 -Identity <user> -Operations
   MailItemsAccessed -ResultSize 10000 -ShowDetails | Where {$_.OperationProperties -like
   "*IsThrottled:True*"} | FL
   ```

2. Check for sync activities. If an attacker uses an email client to downloaded messages in a mailbox, they can disconnect the computer from the Internet and access the messages locally without interacting with the server. This means that mailbox auditing would not be able to audit these activities.

   To search for MailItemsAccessed records where the mail items were accessed by a sync operation, run the following command:

   ### Unified audit log

   ```
   Search-UnifiedAuditLog -StartDate 01/06/2020 -EndDate 02/20/2020 -UserIds <user1,user2> -Operations
   MailItemsAccessed -ResultSize 1000 | Where {$_.AuditData -like '*"MailAccessType","Value":"Sync"*'} |
   FL
   ```

   ### Mailbox audit log

   ```
   Search-MailboxAuditLog -StartDate 01/06/2020 -EndDate 01/20/2020 -Identity <user> -Operations
   MailItemsAccessed -ResultSize 10000 -ShowDetails | Where {$_.OperationProperties -like
   "*MailAccessType:Sync*"} | FL
   ```

3. Check sync activities to determine in any of them have occurred in the same context as the one used by the attacker access the mailbox. Context is identified and differentiated by the IP address of the client computer used to access the mailbox and the mail protocol. For more information, see the Identifying the access contexts of different audit records section.

   Use the properties listed below to investigate. These properties are located in the AuditData or OperationProperties property. If any of the syncs occur in the same context as the attacker activity, assume the attacker has synced all mail items to their client, which means the entire mailbox has probably been compromised.

PROPERTY	DESCRIPTION
ClientInfoString	Describes protocol, client (includes version)

PROPERTY	DESCRIPTION
ClientIPAddress	IP address of the client machine.
SessionId	Session ID helps to differentiate attacker actions vs day-to-day user activities on the same account (in the case of a compromised account)
UserId	UPN of the user reading the message.

4. Check for bind activities. After performing steps 2 and step 3, you can be confident that all other access to email messages by the attacker will be captured in the MailItemsAccessed audit records that have a MailAccessType property with a value of "Bind".

   To search for MailItemsAccessed records where the mail items were accessed by a Bind operation, run the following command.

   ### Unified audit log

   ```
   Search-UnifiedAuditLog -StartDate 01/06/2020 -EndDate 01/20/2020 -UserIds <user1,user2> -Operations
   MailItemsAccessed -ResultSize 1000 | Where {$_.AuditData -like '*"MailAccessType","Value":"Bind"*'} |
   FL
   ```

   ### Mailbox audit log

   ```
   Search-MailboxAuditLog -StartDate 01/06/2020 -EndDate 01/20/2020 -Identity <user> -Operations
   MailItemsAccessed -ResultSize 10000 -ShowDetails | Where {$_.OperationProperties -like
   "*MailAccessType:Bind*"} | FL
   ```

   Email messages that were accessed are identified by their internet message Id. You can also check to see if any audit records have the same context as the ones for other attacker activity. For more information, see the Identifying the access contexts of different audit records section.

   You can use the audit data for bind operations in two different ways:

   - Access or collect all email messages the attacker accessed by using the InternetMessageId to find them and then checking to see if any of those messages contains sensitive information.
   - Use the InternetMessageId to search audit records related to a set of potentially sensitive email messages. This is useful if you're concerned only about a small number of messages.

Filtering of duplicate audit records

Duplicate audit records for the same bind operations that occur within an hour of each other are filtered out to remove auditing noise. Sync operations are also filtered out at one-hour intervals. The exception to this de-duplication process occurs if, for the same InternetMessageId, any of the properties described in the following table are different. If one of these properties is different in a duplicate operation, a new audit record is generated. This process is described in more detail in the next section.

PROPERTY	DESCRIPTION	
ClientIPAddress	IP address of the client computer.	

PROPERTY	DESCRIPTION	
ClientInfoString	The client protocol, client used to access the mailbox.	
ParentFolder	The full folder path of the mail item that was accessed.	
Logon_type	The logon type of the user who performed the action. The logon types (and their corresponding Enum value) are Owner (0), Admin (1), or Delegate (2).	
MailAccessType	Whether the access is a bind or a sync operation.	
MailboxUPN	The UPN of the mailbox where the message being read is located.	
User	The UPN of the user reading the message.	
SessionId	The Session Id helps to differentiate attacker actions and day-to-day user activities in the same mailbox (in the case of account compromise) For more information about sessions, see Contextualizing attacker activity within sessions in Exchange Online.	

Identifying the access contexts of different audit records

It's common that an attacker may access a mailbox at the same time the mailbox owner is accessing it. To differentiate between access by the attacker and the mailbox owner, there are audit record properties that define the context of the access. As previously explained, when the values for these properties are different, even when the activity occurs within the aggregation interval, separate audit records are generated. In the following example, there are three different audit records. Each one is differentiated by the Session Id and ClientIPAddress properties. The messages that were accessed are also identified.

AUDIT RECORD 1	AUDIT RECORD 2	AUDIT RECORD 3
ClientIPAddress**1** SessionId**2**	ClientIPAddress**2** SessionId**2**	ClientIPAddress**1** SessionId**3**
InternetMessageId**A** InternetMessageId**D** InternetMessageId**E** InternetMessageId**F**	InternetMessageId**A** InternetMessageId**C**	InternetMessageId**B**

If any of the properties listed in the table in the previous section are different, a separate audit record is generated to track the new context. Accesses will be sorted into the separate audit records depending on the context in which the activity took place.

For example, in audit records shown in the following screenshot, though we are accessing mail from EWSEditor and OWA simultaneously, the access activity is collated in different audit records depending on the context in which the access took place. In this case, the context is defined by different values for the ClientInfoString property.

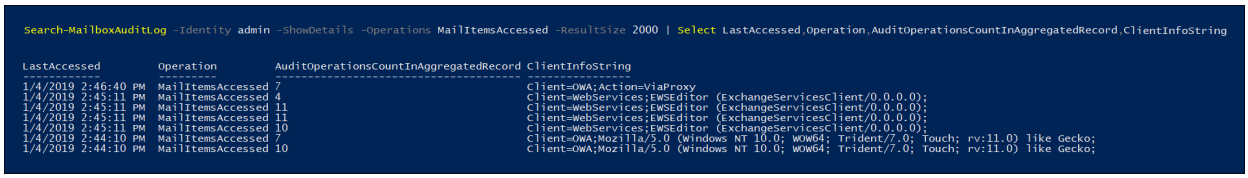

Here is the syntax for the command shown in the previous screenshot:

```
Search-MailboxAuditLog -Identity admin -ShowDetails -Operations MailItemsAccessed -ResultSize 2000 | Select LastAccessed,Operation,AuditOperationsCountInAggregatedRecord,ClientInfoString
```

Alert policies in the security and compliance center

2/18/2021 • 30 minutes to read • Edit Online

You can use the alert policy and alert dashboard tools in the Microsoft 365 security and compliance centers to create alert policies and then view the alerts generated when users perform activities that match the conditions of an alert policy. There are several default alert policies that help you monitor activities such as assigning admin privileges in Exchange Online, malware attacks, phishing campaigns, and unusual levels of file deletions and external sharing.

Alert policies let you categorize the alerts that are triggered by a policy, apply the policy to all users in your organization, set a threshold level for when an alert is triggered, and decide whether to receive email notifications when alerts are triggered. There's also a **View alerts** page in the security and compliance center where you can view and filter alerts, set an alert status to help you manage alerts, and then dismiss alerts after you've addressed or resolved the underlying incident.

> **NOTE**
> Alert policies are available for organizations with a Microsoft 365 Enterprise, Office 365 Enterprise, or Office 365 US Government E1/F1/G1, E3/F3/G3, or E5/G5 subscription. Advanced functionality is only available for organizations with an E5/G5 subscription, or for organizations that have an E1/F1/G1 or E3/F3/G3 subscription and a Microsoft Defender for Office 365 P2 or a Microsoft 365 E5 Compliance or an E5 eDiscovery and Audit add-on subscription. The functionality that requires an E5/G5 or add-on subscription is highlighted in this topic. Also note that alert policies are available in Office 365 GCC, GCC High, and DoD US government environments.

How alert policies work

Here's a quick overview of how alert policies work and the alerts that are triggers when user or admin activity matches the conditions of an alert policy.

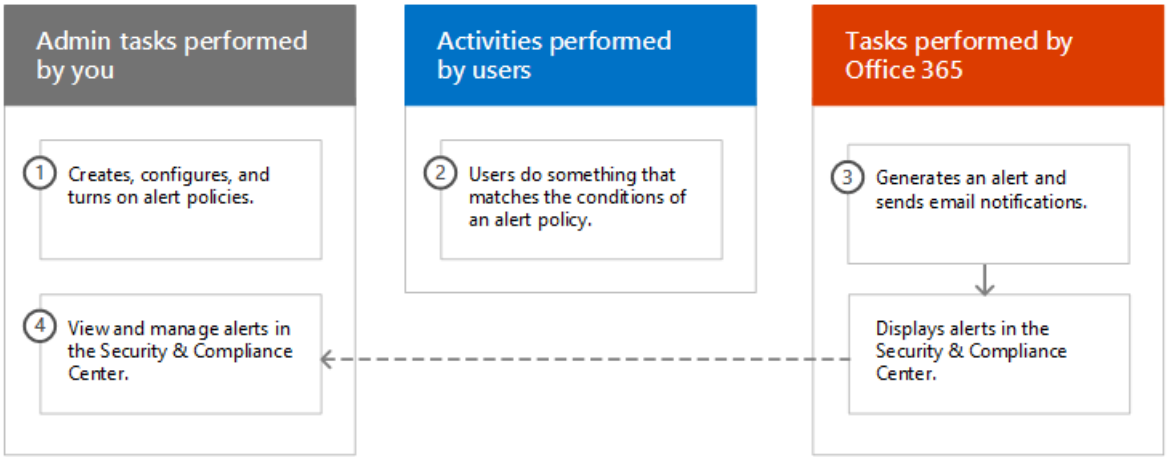

1. An admin in your organization creates, configures, and turns on an alert policy by using the **Alert policies** page in the security and compliance center. You can also create alert policies by using the New-ProtectionAlert cmdlet in Security & Compliance Center PowerShell.

   To create alert policies, you have to be assigned the Manage Alerts role or the Organization Configuration role in the security and compliance center.

> **NOTE**
> It takes up to 24 hours after creating or updating an alert policy before alerts can be triggered by the policy. This is because the policy has to be synced to the alert detection engine.

2. A user performs an activity that matches the conditions of an alert policy. In the case of malware attacks, infected email messages sent to users in your organization trigger an alert.
3. Microsoft 365 generates an alert that's displayed on the **View alerts** page in the Security & Compliance Center. Also, if email notifications are enabled for the alert policy, Microsoft sends a notification to a list of recipients. The alerts that an admin or other users can see that on the View alerts page is determined by the roles assigned to the user. For more information, see RBAC permissions required to view alerts.
4. An admin manages alerts in the security and compliance center. Managing alerts consists of assigning an alert status to help track and manage any investigation.

Alert policy settings

An alert policy consists of a set of rules and conditions that define the user or admin activity that generates an alert, a list of users who trigger the alert if they perform the activity, and a threshold that defines how many times the activity has to occur before an alert is triggered. You also categorize the policy and assign it a severity level. These two settings help you manage alert policies (and the alerts that are triggered when the policy conditions are matched) because you can filter on these settings when managing policies and viewing alerts in the security and compliance center. For example, you can view alerts that match the conditions from the same category or view alerts with the same severity level.

To view and create alert policies, go to https://protection.office.com and then select **Alerts** > **Alert policies**.

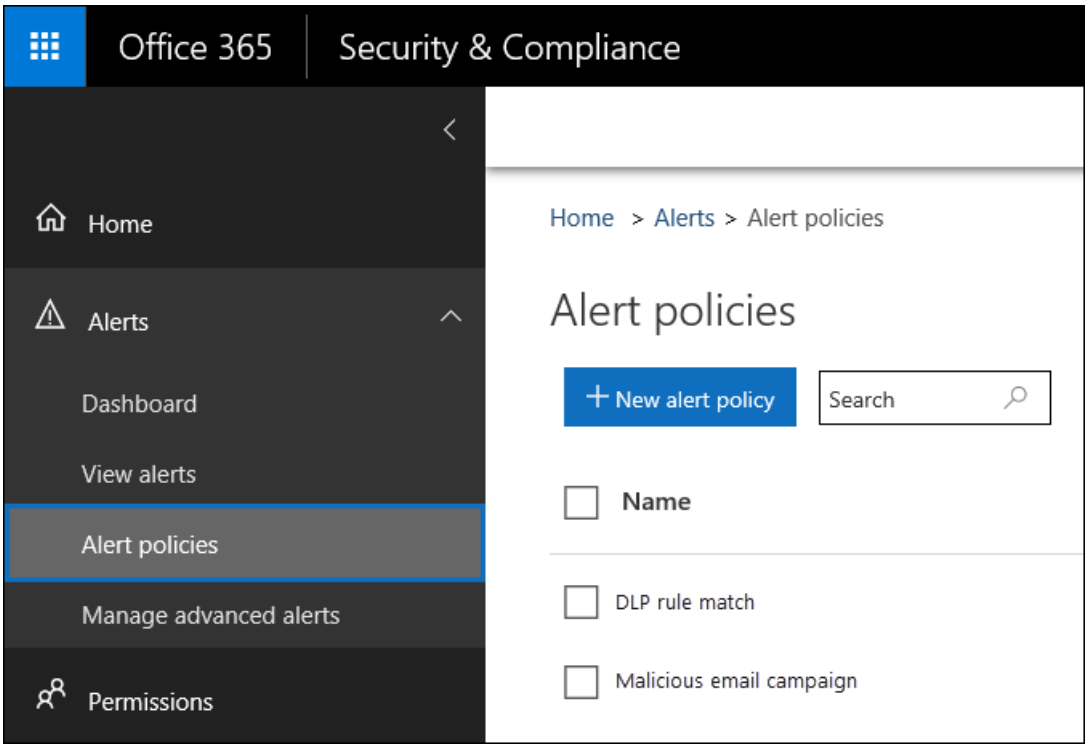

An alert policy consists of the following settings and conditions.

- **Activity the alert is tracking** - You create a policy to track an activity or in some cases a few related activities, such a sharing a file with an external user by sharing it, assigning access permissions, or creating an anonymous link. When a user performs the activity defined by the policy, an alert is triggered based on the alert threshold settings.

> **NOTE**
> The activities that you can track depend on your organization's Office 365 Enterprise or Office 365 US Government plan. In general, activities related to malware campaigns and phishing attacks require an E5/G5 subscription or an E1/F1/G1 or E3/F3/G3 subscription with an Defender for Office 365 Plan 2 add-on subscription.

- **Activity conditions** - For most activities, you can define additional conditions that must be met to trigger an alert. Common conditions include IP addresses (so that an alert is triggered when the user performs the activity on a computer with a specific IP address or within an IP address range), whether an alert is triggered if a specific user or users perform that activity, and whether the activity is performed on a specific file name or URL. You can also configure a condition that triggers an alert when the activity is performed by any user in your organization. The available conditions are dependent on the selected activity.

- **When the alert is triggered** - You can configure a setting that defines how often an activity can occur before an alert is triggered. This allows you to set up a policy to generate an alert every time an activity matches the policy conditions, when a certain threshold is exceeded, or when the occurrence of the activity the alert is tracking becomes unusual for your organization.

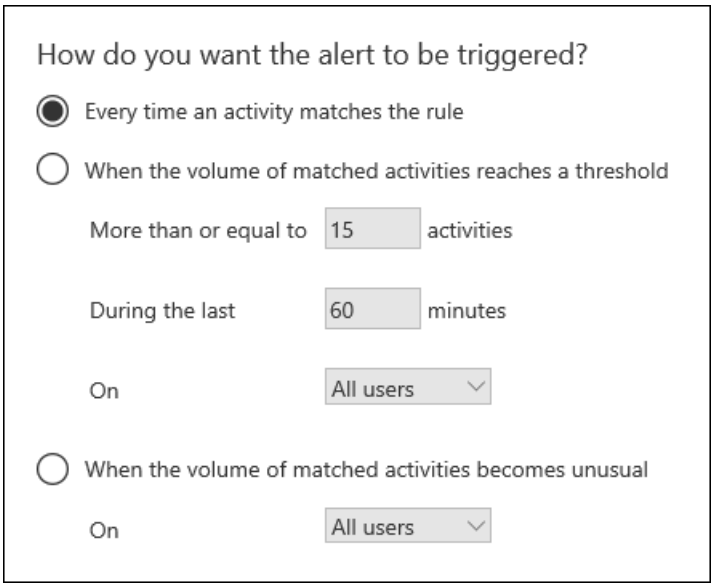

  If you select the setting based on unusual activity, Microsoft establishes a baseline value that defines the normal frequency for the selected activity. It takes up to seven days to establish this baseline, during which alerts won't be generated. After the baseline is established, an alert is triggered when the frequency of the activity tracked by the alert policy greatly exceeds the baseline value. For auditing-related activities (such as file and folder activities), you can establish a baseline based on a single user or based on all users in your organization; for malware-related activities, you can establish a baseline based on a single malware family, a single recipient, or all messages in your organization.

  > **NOTE**
  > The ability to configure alert policies based on a threshold or based on unusual activity requires an E5/G5 subscription, or an E1/F1/G1 or E3/F3/G3 subscription with a Microsoft Defender for Office 365 P2, Microsoft 365 E5 Compliance, or Microsoft 365 eDiscovery and Audit add-on subscription. Organizations with an E1/F1/G1 and E3/F3/G3 subscription can only create alert policies where an alert is triggered every time that an activity occurs.

- **Alert category** - To help with tracking and managing the alerts generated by a policy, you can assign one of the following categories to a policy.

  - Data loss prevention

- Information governance
- Mail flow
- Permissions
- Threat management
- Others

When an activity occurs that matches the conditions of the alert policy, the alert that's generated is tagged with the category defined in this setting. This allows you to track and manage alerts that have the same category setting on the **View alerts** page in the security and compliance center because you can sort and filter alerts based on category.

- **Alert severity** - Similar to the alert category, you assign a severity attribute (**Low**, **Medium**, **High**, or **Informational**) to alert policies. Like the alert category, when an activity occurs that matches the conditions of the alert policy, the alert that's generated is tagged with the same severity level that's set for the alert policy. Again, this allows you to track and manage alerts that have the same severity setting on the **View alerts** page. For example, you can filter the list of alerts so that only alerts with a **High** severity are displayed.

  > **TIP**
  > When setting up an alert policy, consider assigning a higher severity to activities that can result in severely negative consequences, such as detection of malware after delivery to users, viewing of sensitive or classified data, sharing data with external users, or other activities that can result in data loss or security threats. This can help you prioritize alerts and the actions you take to investigate and resolve the underlying causes.

- **Email notifications** - You can set up the policy so that email notifications are sent (or not sent) to a list of users when an alert is triggered. You can also set a daily notification limit so that once the maximum number of notifications has been reached, no more notifications are sent for the alert during that day. In addition to email notifications, you or other administrators can view the alerts that are triggered by a policy on the **View alerts** page. Consider enabling email notifications for alert policies of a specific category or that have a higher severity setting.

Default alert policies

Microsoft provides built-in alert policies that help identify Exchange admin permissions abuse, malware activity, potential external and internal threats, and information governance risks. On the **Alert policies** page, the names of these built-in policies are in bold and the policy type is defined as **System**. These policies are turned on by default. You can turn off these policies (or back on again), set up a list of recipients to send email notifications to, and set a daily notification limit. The other settings for these policies can't be edited.

The following table lists and describes the available default alert policies and the category each policy is assigned to. The category is used to determine which alerts a user can view on the View alerts page. For more information, see RBAC permissions required to view alerts.

The table also indicates the Office 365 Enterprise and Office 365 US Government plan required for each one. Some default alert policies are available if your organization has the appropriate add-on subscription in addition to an E1/F1/G1 or E3/F3/G3 subscription.

DEFAULT ALERT POLICY	DESCRIPTION	CATEGORY	ENTERPRISE SUBSCRIPTION

DEFAULT ALERT POLICY	DESCRIPTION	CATEGORY	ENTERPRISE SUBSCRIPTION
A potentially malicious URL click was detected	Generates an alert when a user protected by Safe Links in your organization clicks a malicious link. This event is triggered when URL verdict changes are identified by Microsoft Defender for Office 365 or when users override the Safe Links pages (based on your organization's Microsoft 365 for business Safe Links policy). This alert policy has a **High** severity setting. For Defender for Office 365 P2, E5, G5 customers, this alert automatically triggers automated investigation and response in Office 365. For more information on events that trigger this alert, see Set up Safe Links policies.	Threat management	E5/G5 or Defender for Office 365 P2 add-on subscription
Admin Submission result completed	Generates an alert when an Admin Submission completes the rescan of the submitted entity. An alert will be triggered every time a rescan result is rendered from an Admin Submission. These alerts are meant to remind you to review the results of previous submissions, submit user reported messages to get the latest policy check and rescan verdicts, and help you determine if the filtering policies in your organization are having the intended impact. This policy has a **Low** severity setting.	Threat management	E1/F1, E3/F3, or E5

DEFAULT ALERT POLICY	DESCRIPTION	CATEGORY	ENTERPRISE SUBSCRIPTION
Admin triggered manual investigation of email	Generates an alert when an admin triggers the manual investigation of an email from Threat Explorer. For more information, see [Example: A security administrator triggers an investigation from Threat Explorer](https://docs.microsoft.com/microsoft-365/security/office-365-security/automated-investigation-response-office#example-a-security-administrator-triggers-an-investigation-from-threat-explorer). This alert notifies your organization that the investigation was started. The alert provides information about who triggered it and includes a link to the investigation. This policy has an **Informational** severity setting.	Threat management	E5/G5 or Microsoft Defender for Office 365 P2 add-on subscription
Creation of forwarding/redirect rule	Generates an alert when someone in your organization creates an inbox rule for their mailbox that forwards or redirects messages to another email account. This policy only tracks inbox rules that are created using Outlook on the web (formerly known as Outlook Web App) or Exchange Online PowerShell. This policy has a **Low** severity setting. For more information about using inbox rules to forward and redirect email in Outlook on the web, see Use rules in Outlook on the web to automatically forward messages to another account.	Threat management	E1/F1/G1, E3/F3/G3, or E5/G5

DEFAULT ALERT POLICY	DESCRIPTION	CATEGORY	ENTERPRISE SUBSCRIPTION
eDiscovery search started or exported	Generates an alert when someone uses the Content search tool in the Security and compliance center. An alert is triggered when the following content search activities are performed: * A content search is started * The results of a content search are exported * A content search report is exported Alerts are also triggered when the previous content search activities are performed in association with an eDiscovery case. This policy has a **Medium** severity setting. For more information about content search activities, see Search for eDiscovery activities in the audit log.	Threat management	E1/F1/G1, E3/F3/G3, or E5/G5
Elevation of Exchange admin privilege	Generates an alert when someone is assigned administrative permissions in your Exchange Online organization. For example, when a user is added to the Organization Management role group in Exchange Online. This policy has a **Low** severity setting.	Permissions	E1/F1/G1, E3/F3/G3, or E5/G5
Email messages containing malware removed after delivery	Generates an alert when any messages containing malware are delivered to mailboxes in your organization. If this event occurs, Microsoft removes the infected messages from Exchange Online mailboxes using Zero-hour auto purge. This policy has an **Informational** severity setting and automatically triggers automated investigation and response in Office 365.	Threat management	E5/G5 or Microsoft Defender for Office 365 P2 add-on subscription

DEFAULT ALERT POLICY	DESCRIPTION	CATEGORY	ENTERPRISE SUBSCRIPTION
Email messages containing phish URLs removed after delivery	Generates an alert when any messages containing phish are delivered to mailboxes in your organization. If this event occurs, Microsoft removes the infected messages from Exchange Online mailboxes using Zero-hour auto purge. This policy has an **Informational** severity setting and automatically triggers automated investigation and response in Office 365.	Threat management	E5/G5 or Defender for Office 365 P2 add-on subscription
Email reported by user as malware or phish	Generates an alert when users in your organization report messages as phishing email using the Report Message add-in. This policy has an **Informational** severity setting. For more information about this add-in, see Use the Report Message add-in. For Defender for Office 365 P2, E5, G5 customers, this alert automatically triggers automated investigation and response in Office 365.	Threat management	E1/F1/G1, E3/F3/G3, or E5/G5
Email sending limit exceeded	Generates an alert when someone in your organization has sent more mail than is allowed by the outbound spam policy. This is usually an indication the user is sending too much email or that the account may be compromised. This policy has a **Medium** severity setting. If you get an alert generated by this alert policy, it's a good idea to check whether the user account is compromised.	Threat management	E1/F1/G1, E3/F3/G3, or E5/G5
Form blocked due to potential phishing attempt	Generates an alert when someone in your organization has been restricted from sharing forms and collecting responses using Microsoft Forms due to detected repeated phishing attempt behavior. This policy has a **High severity** setting.	Threat management	E1, E3/F3, or E5

DEFAULT ALERT POLICY	DESCRIPTION	CATEGORY	ENTERPRISE SUBSCRIPTION
Form flagged and confirmed as phishing	Generates an alert when a form created in Microsoft Forms from within your organization has been identified as potential phishing through Report Abuse and confirmed as phishing by Microsoft. This policy has a **High** severity setting.	Threat management	E1, E3/F3, or E5
Messages have been delayed	Generates an alert when Microsoft can't deliver email messages to your on-premises organization or a partner server by using a connector. When this happens, the message is queued in Office 365. This alert is triggered when there are 2,000 messages or more that have been queued for more than an hour. This policy has a **High** severity setting.	Mail flow	E1/F1/G1, E3/F3/G3, or E5/G5
Malware campaign detected after delivery	Generates an alert when an unusually large number of messages containing malware are delivered to mailboxes in your organization. If this event occurs, Microsoft removes the infected messages from Exchange Online mailboxes. This policy has a **High** severity setting.	Threat management	E5/G5 or Microsoft Defender for Office 365 P2 add-on subscription
Malware campaign detected and blocked	Generates an alert when someone has attempted to send an unusually large number of email messages containing a certain type of malware to users in your organization. If this event occurs, the infected messages are blocked by Microsoft and not delivered to mailboxes. This policy has a **Low** severity setting.	Threat management	E5/G5 or Defender for Office 365 P2 add-on subscription
Malware campaign detected in SharePoint and OneDrive	Generates an alert when an unusually high volume of malware or viruses is detected in files located in SharePoint sites or OneDrive accounts in your organization. This policy has a **High** severity setting.	Threat management	E5/G5 or Defender for Office 365 P2 add-on subscription

DEFAULT ALERT POLICY	DESCRIPTION	CATEGORY	ENTERPRISE SUBSCRIPTION
Malware not zapped because ZAP is disabled	Generates an alert when Microsoft detects delivery of a malware message to a mailbox because Zero-Hour Auto Purge for Phish messages is disabled. This policy has an **Informational** severity setting.	Threat management	E5/G5 or Defender for Office 365 P2 add-on subscription
Phish delivered because a user's Junk Mail folder is disabled	Generates an alert when Microsoft detects a user's Junk Mail folder is disabled, allowing delivery of a high confidence phishing message to a mailbox. This policy has an **Informational** severity setting.	Threat management	E5/G5 or Defender for Office 365 P1 or P2 add-on subscription
Phish delivered due to an ETR override	Generates an alert when Microsoft detects an Exchange Transport Rule (ETR) that allowed delivery of a high confidence phishing message to a mailbox. This policy has an **Informational** severity setting. For more information about Exchange Transport Rules (Mail flow rules), see Mail flow rules (transport rules) in Exchange Online.	Threat management	E5/G5 or Defender for Office 365 P1 or P2 add-on subscription
Phish delivered due to an IP allow policy	Generates an alert when Microsoft detects an IP allow policy that allowed delivery of a high confidence phishing message to a mailbox. This policy has an **Informational** severity setting. For more information about the IP allow policy (connection filtering), see Configure the default connection filter policy - Office 365.	Threat management	E5/G5 or Defender for Office 365 P1 or P2 add-on subscription

DEFAULT ALERT POLICY	DESCRIPTION	CATEGORY	ENTERPRISE SUBSCRIPTION
Phish not zapped because ZAP is disabled	Generates an alert when Microsoft detects delivery of a high confidence phishing message to a mailbox because Zero-Hour Auto Purge for Phish messages is disabled. This policy has an **Informational** severity setting.	Threat management	E5/G5 or Defender for Office 365 P2 add-on subscription
Phish delivered due to tenant or user override[1]	Generates an alert when Microsoft detects an admin or user override allowed the delivery of a phishing message to a mailbox. Examples of overrides include an inbox or mail flow rule that allows messages from a specific sender or domain, or an anti-spam policy that allows messages from specific senders or domains. This policy has a **High** severity setting.	Threat management	E5/G5 or Defender for Office 365 P2 add-on subscription
Suspicious email forwarding activity	Generates an alert when someone in your organization has autoforwarded email to a suspicious external account. This is an early warning for behavior that may indicate the account is compromised, but not severe enough to restrict the user. This policy has a **Medium** severity setting. Although it's rare, an alert generated by this policy may be an anomaly. It's a good idea to check whether the user account is compromised.	Threat management	E1/F1/G1, E3/F3/G3, or E5/G5

DEFAULT ALERT POLICY	DESCRIPTION	CATEGORY	ENTERPRISE SUBSCRIPTION
Suspicious email sending patterns detected	Generates an alert when someone in your organization has sent suspicious email and is at risk of being restricted from sending email. This is an early warning for behavior that may indicate that the account is compromised, but not severe enough to restrict the user. This policy has a **Medium** severity setting. Although it's rare, an alert generated by this policy may be an anomaly. However, it's a good idea to check whether the user account is compromised.	Threat management	E1/F1/G1, E3/F3/G3, or E5/G5
Tenant restricted from sending email	Generates an alert when most of the email traffic from your organization has been detected as suspicious and Microsoft has restricted your organization from sending email. Investigate any potentially compromised user and admin accounts, new connectors, or open relays, and then contact Microsoft Support to unblock your organization. This policy has a **High** severity setting. For more information about why organizations are blocked, see Fix email delivery issues for error code 5.7.7xx in Exchange Online.	Threat management	E1/F1/G1, E3/F3/G3, or E5/G5
Unusual external user file activity	Generates an alert when an unusually large number of activities are performed on files in SharePoint or OneDrive by users outside of your organization. This includes activities such as accessing files, downloading files, and deleting files. This policy has a **High** severity setting.	Information governance	E5/G5, Microsoft Defender for Office 365 P2, or Microsoft 365 E5 add-on subscription
Unusual volume of external file sharing	Generates an alert when an unusually large number of files in SharePoint or OneDrive are shared with users outside of your organization. This policy has a **Medium** severity setting.	Information governance	E5/G5, Defender for Office 365 P2, or Microsoft 365 E5 add-on subscription

DEFAULT ALERT POLICY	DESCRIPTION	CATEGORY	ENTERPRISE SUBSCRIPTION
Unusual volume of file deletion	Generates an alert when an unusually large number of files are deleted in SharePoint or OneDrive within a short time frame. This policy has a **Medium** severity setting.	Information governance	E5/G5, Defender for Office 365 P2, or Microsoft 365 E5 add-on subscription
Unusual increase in email reported as phish	Generates an alert when there's a significant increase in the number of people in your organization using the Report Message add-in in Outlook to report messages as phishing mail. This policy has a **High** severity setting. For more information about this add-in, see Use the Report Message add-in.	Threat management	E5/G5 or Defender for Office 365 P2 add-on subscription
User impersonation phish delivered to inbox/folder[1,2]	Generates an alert when Microsoft detects that an admin or user override has allowed the delivery of a user impersonation phishing message to the inbox (or other user-accessible folder) of a mailbox. Examples of overrides include an inbox or mail flow rule that allows messages from a specific sender or domain, or an anti-spam policy that allows messages from specific senders or domains. This policy has a **Medium** severity setting.	Threat management	E5/G5 or Defender for Office 365 P2 add-on subscription

DEFAULT ALERT POLICY	DESCRIPTION	CATEGORY	ENTERPRISE SUBSCRIPTION
User restricted from sending email	Generates an alert when someone in your organization is restricted from sending outbound mail. This typically results when an account is compromised, and the user is listed on the **Restricted Users** page in the Security & Compliance Center. (To access this page, go to **Threat management** > **Review** > **Restricted Users**). This policy has a **High** severity setting. For more information about restricted users, see Removing a user, domain, or IP address from a block list after sending spam email.	Threat management	E1/F1/G1, E3/F3/G3, or E5/G5
User restricted from sharing forms and collecting responses	Generates an alert when someone in your organization has been restricted from sharing forms and collecting responses using Microsoft Forms due to detected repeated phishing attempt behavior. This policy has a **High** severity setting.	Threat management	E1, E3/F3, or E5

NOTE

[1] We've temporarily removed this default alert policy based on customer feedback. We're working to improve it, and will replace it with a new version in the near future. Until then, you can create a custom alert policy to replace this functionality by using the following settings:

* Activity is Phish email detected at time of delivery
* Mail is not ZAP'd
* Mail direction is Inbound
* Mail delivery status is Delivered
* Detection technology is Malicious URL retention, URL detonation, Advanced phish filter, General phish filter, Domain impersonation, User impersonation, and Brand impersonation

For more information about anti-phishing in Office 365, see Set up anti-phishing and anti-phishing policies.

[2] To recreate this alert policy, follow the guidance in the previous footnote, but choose User impersonation as the only Detection technology.

The unusual activity monitored by some of the built-in policies is based on the same process as the alert threshold setting that was previously described. Microsoft establishes a baseline value that defines the normal frequency for "usual" activity. Alerts are then triggered when the frequency of activities tracked by the built-in alert policy greatly exceeds the baseline value.

Viewing alerts

When an activity performed by users in your organization matches the settings of an alert policy, an alert is generated and displayed on the **View alerts** page in the security and compliance center. Depending on the settings of an alert policy, an email notification is also sent to a list of specified users when an alert is triggered. For each alert, the dashboard on the **View alerts** page displays the name of the corresponding alert policy, the severity and category for the alert (defined in the alert policy), and the number of times an activity has occurred that resulted in the alert being generated. This value is based on the threshold setting of the alert policy. The dashboard also shows the status for each alert. For more information about using the status property to manage alerts, see Managing alerts.

To view alerts, go to https://protection.office.com and then select **Alerts** > **View alerts**.

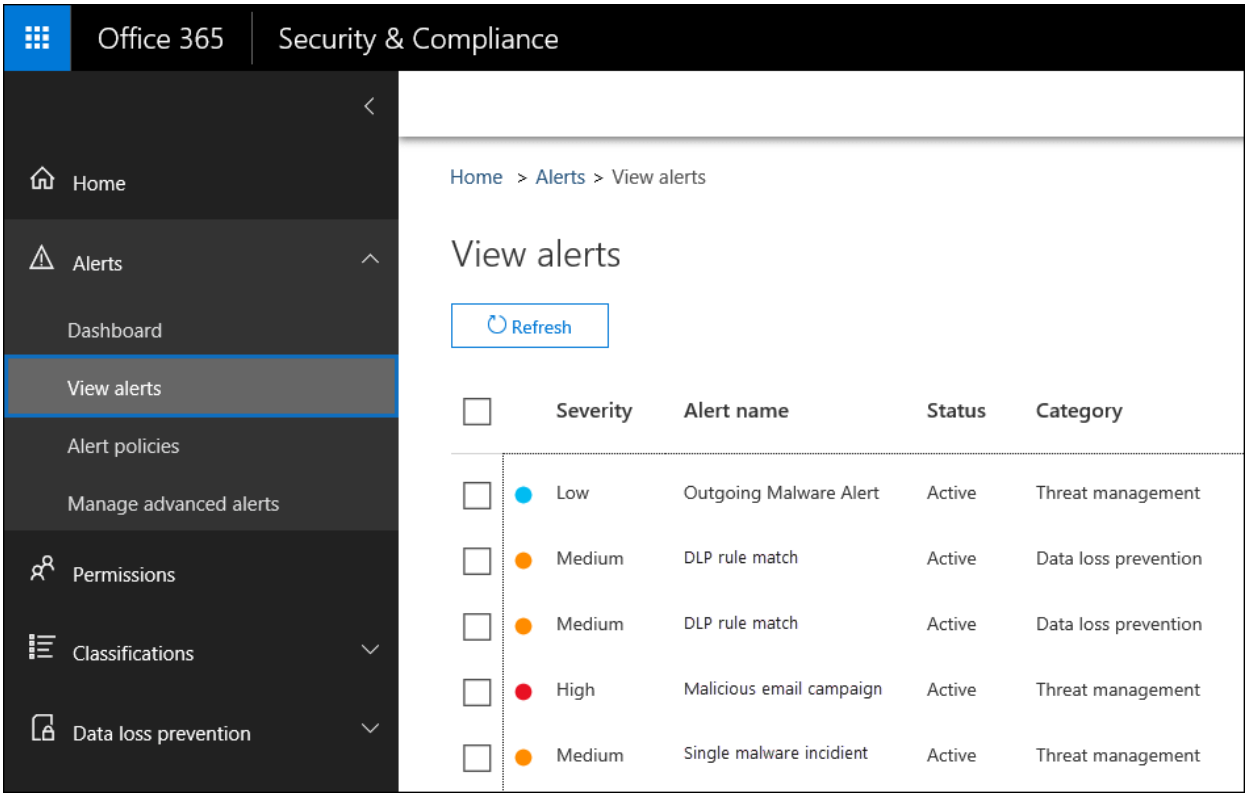

You can use the following filters to view a subset of all the alerts on the **View alerts** page.

- **Status.** Use this filter to show alerts that are assigned a particular status. The default status is **Active**. You or other administrators can change the status value.
- **Policy.** Use this filter to show alerts that match the setting of one or more alert policies. Or you can display all alerts for all alert policies.
- **Time range.** Use this filter to show alerts that were generated within a specific date and time range.
- **Severity.** Use this filter to show alerts that are assigned a specific severity.
- **Category.** Use this filter to show alerts from one or more alert categories.
- **Tags.** Use this filter to show alerts from one or more user tags. Tags are reflected based on tagged mailboxes or users that appear in the alerts. See User tags in Office 356 ATP to learn more.
- **Source.** Use this filter to show alerts triggered by alert policies in the security and compliance center or alerts triggered by Office 365 Cloud App Security policies, or both. For more information about Office 365 Cloud App Security alerts, see Viewing Cloud App Security alerts.

IMPORTANT

Filtering and sorting by user tags is currently in public preview. It may be substantially modified before it's commercially released. Microsoft makes no warranties, express or implied, with respect to the information provided about it.

Alert aggregation

When multiple events that match the conditions of an alert policy occur with a short period of time, they are added to an existing alert by a process called *alert aggregation*. When an event triggers an alert, the alert is generated and displayed on the **View alerts** page and a notification is sent. If the same event occurs within the aggregation interval, then Microsoft 365 adds details about the new event to the existing alert instead of triggering a new alert. The goal of alert aggregation is to help reduce alert "fatigue" and let you focus and take action on fewer alerts for the same event.

The length of the aggregation interval depends on your Office 365 or Microsoft 365 subscription.

SUBSCRIPTION	AGGREGATION INTERVAL
Office 365 or Microsoft 365 E5/G5	1 minute
Defender for Office 365 Plan 2	1 minute
E5 Compliance add-on or E5 Discovery and Audit add-on	1 minute
Office 365 or Microsoft 365 E1/F1/G1 or E3/F3/G3	15 minutes
Defender for Office 365 Plan 1 or Exchange Online Protection	15 minutes

When events that match the same alert policy occur within the aggregation interval, details about the subsequent event are added to the original alert. For all events, information about aggregated events is displayed in the details field and the number of times an event occurred with the aggregation interval is displayed in the activity/hit count field. You can view more information about all aggregated events instances by viewing the activity list.

The following screenshot shows an alert with four aggregated events. The activity list contains information about the four email messages relevant to the alert.

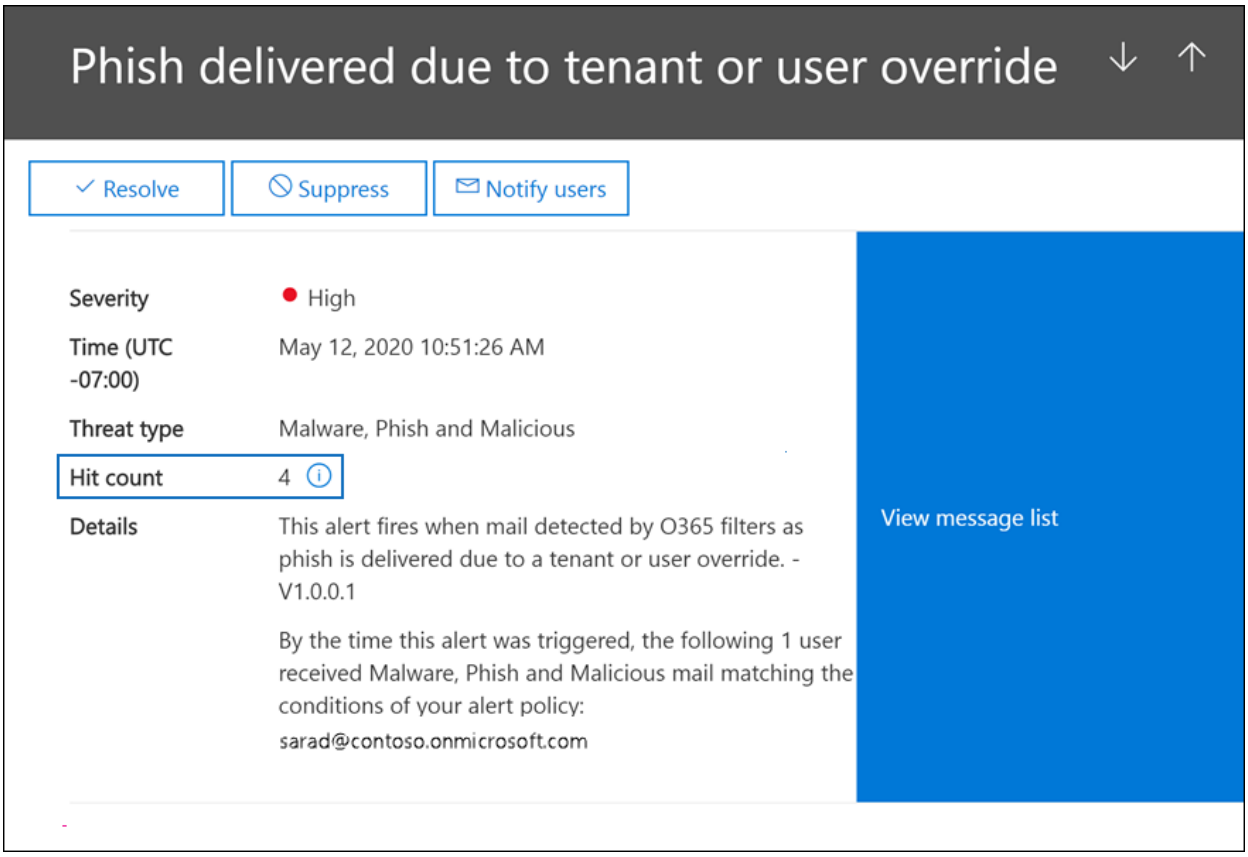

Keep the following things in mind about alert aggregation:

- Alerts triggered by the **A potentially malicious URL click was detected** default alert policy are not aggregated. This is because alerts triggered by this policy are unique to each user and email message.
- At this time, the **Hit count** alert property doesn't indicate the number of aggregated events for all alert policies. For alerts triggered by these alert policies, you can view the aggregated events by clicking **View message list** or **View activity** on the alert. We're working to make the number of aggregated events listed in the **Hit count** alert property available for all alert policies.

RBAC permissions required to view alerts

The Role Based Access Control (RBAC) permissions assigned to users in your organization determine which alerts a user can see on the **View alerts** page. How is this accomplished? The management roles assigned to users (based on their membership in role groups in the Security & Compliance Center) determine which alert categories a user can see on the **View alerts** page. Here are some examples:

- Members of the Records Management role group can view only the alerts that are generated by alert policies that are assigned the **Information governance** category.
- Members of the Compliance Administrator role group can't view alerts that are generated by alert policies that are assigned the **Threat management** category.
- Members of the eDiscovery Manager role group can't view any alerts because none of the assigned roles provide permission to view alerts from any alert category.

This design (based on RBAC permissions) lets you determine which alerts can be viewed (and managed) by users in specific job roles in your organization.

The following table lists the roles that are required to view alerts from the six different alert categories. The first column in the tables lists all roles in the Security & Compliance Center. A check mark indicates that a user who is assigned that role can view alerts from the corresponding alert category listed in the top row.

To see which category a default alert policy is assigned to, see the table in Default alert policies.

ROLE	INFORMATION GOVERNANCE	DATA LOSS PREVENTION	MAIL FLOW	PERMISSIONS	THREAT MANAGEMENT	OTHERS
Audit Logs						
Case Management						
Compliance Administrator	✓	✓		✓		✓
Compliance Search						
Device Management						
Disposition Management						
DLP Compliance Management		✓				
Export						
Hold						
Manage Alerts						✓
Organization Configuration						✓
Preview						
Record Management	✓					
Retention Management	✓					
Review						
RMS Decrypt						
Role Management				✓		
Search And Purge						
Security Administrator		✓		✓	✓	✓

ROLE	INFORMATION GOVERNANCE	DATA LOSS PREVENTION	MAIL FLOW	PERMISSIONS	THREAT MANAGEMENT	OTHERS
Security Reader		✓		✓	✓	✓
Service Assurance View						
Supervisory Review Administrator						
View-Only Audit Logs						
View-Only Device Management						
View-Only DLP Compliance Management		✓				
View-Only Manage Alerts						✓
View-Only Recipients			✓			
View-Only Record Management	✓					
View-Only Retention Management	✓					

TIP

To view the roles that are assigned to each of the default role groups, run the following commands in Security & Compliance Center PowerShell:

```
$RoleGroups = Get-RoleGroup
```

```
$RoleGroups | foreach {Write-Output -InputObject `r`n,$_.Name,"----------------------"; Get-RoleGroup
$_.Identity | Select-Object -ExpandProperty Roles}
```

You can also view the roles assigned to a role group in the Security & Compliance Center. Go to the **Permissions** page, and select a role group. The assigned roles are listed on the flyout page.

Managing alerts

After alerts have been generated and displayed on the **View alerts** page in the security and compliance center, you can triage, investigate, and resolve them. Here are some tasks you can perform to manage alerts.

- **Assign a status to alerts.** You can assign one of the following statuses to alerts: **Active** (the default value), **Investigating**, **Resolved**, or **Dismissed**. Then, you can filter on this setting to display alerts with the same status setting. This status setting can help track the process of managing alerts.
- **View alert details.** You can select an alert to display a flyout page with details about the alert. The detailed information depends on the corresponding alert policy, but it typically includes the following: name of the actual operation that triggered the alert (such as a cmdlet), a description of the activity that triggered the alert, the user (or list of users) who triggered the alert, and the name (and link to) of the corresponding alert policy.
    - The name of the actual operation that triggered the alert, such as a cmdlet or an audit log operation.
    - A description of the activity that triggered the alert.
    - The user who triggered the alert. This is included only for alert policies that are set up to track a single user or a single activity.
    - The number of times the activity tracked by the alert was performed. This number may not match that actual number of related alerts listed on the View alerts page because more alerts may have been triggered.
    - A link to an activity list that includes an item for each activity that was performed that triggered the alert. Each entry in this list identifies when the activity occurred, the name of actual operation (such as "FileDeleted"), and the user who performed the activity, the object (such as a file, an eDiscovery case, or a mailbox) that the activity was performed on, and the IP address of the user's computer. For malware-related alerts, this links to a message list.
    - The name (and link to) of the corresponding alert policy.
- **Suppress email notifications.** You can turn off (or suppress) email notifications from the flyout page for an alert. When you suppress email notifications, Microsoft won't send notifications when activities or events that match the conditions of the alert policy. But alerts will be triggered when activities performed by users match the conditions of the alert policy. You can also turn off email notifications by editing the alert policy.
- **Resolve alerts.** You can mark an alert as resolved on the flyout page for an alert (which sets the status of the alert to **Resolved**). Unless you change the filter, resolved alerts aren't displayed on the **View alerts** page.

Viewing Cloud App Security alerts

Alerts that are triggered by Office 365 Cloud App Security policies are now displayed on the **View alerts** page in the security and compliance center. This includes alerts that are triggered by activity policies and alerts that are triggered by anomaly detection policies in Office 365 Cloud App Security. This means you can view all alerts in the security and compliance center. Office 365 Cloud App Security is only available for organizations with an Office 365 Enterprise E5 or Office 365 US Government G5 subscription. For more information, see Overview of Cloud App Security.

Organizations that have Microsoft Cloud App Security as part of an Enterprise Mobility + Security E5 subscription or as a standalone service can also view Cloud App Security alerts that are related to Office 365 apps and services in the Security & Compliance Center.

To display only Cloud App Security alerts in the security and compliance center, use the **Source** filter and select **Cloud App Security**.

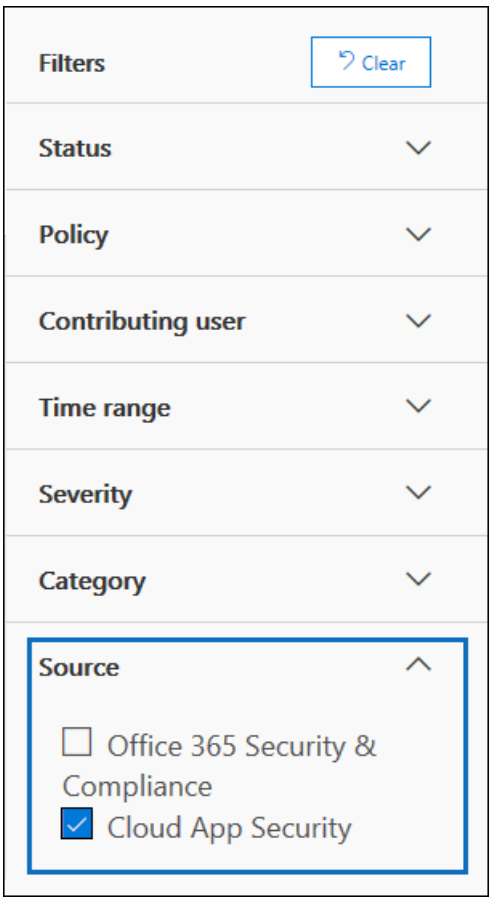

Similar to an alert triggered by an alert policy in the security and compliance center, you can select a Cloud App Security alert to display a flyout page with details about the alert. The alert includes a link to view the details and manage the alert in the Cloud App Security portal and a link to the corresponding Cloud App Security policy that triggered the alert. See Monitor alerts in Cloud App Security.

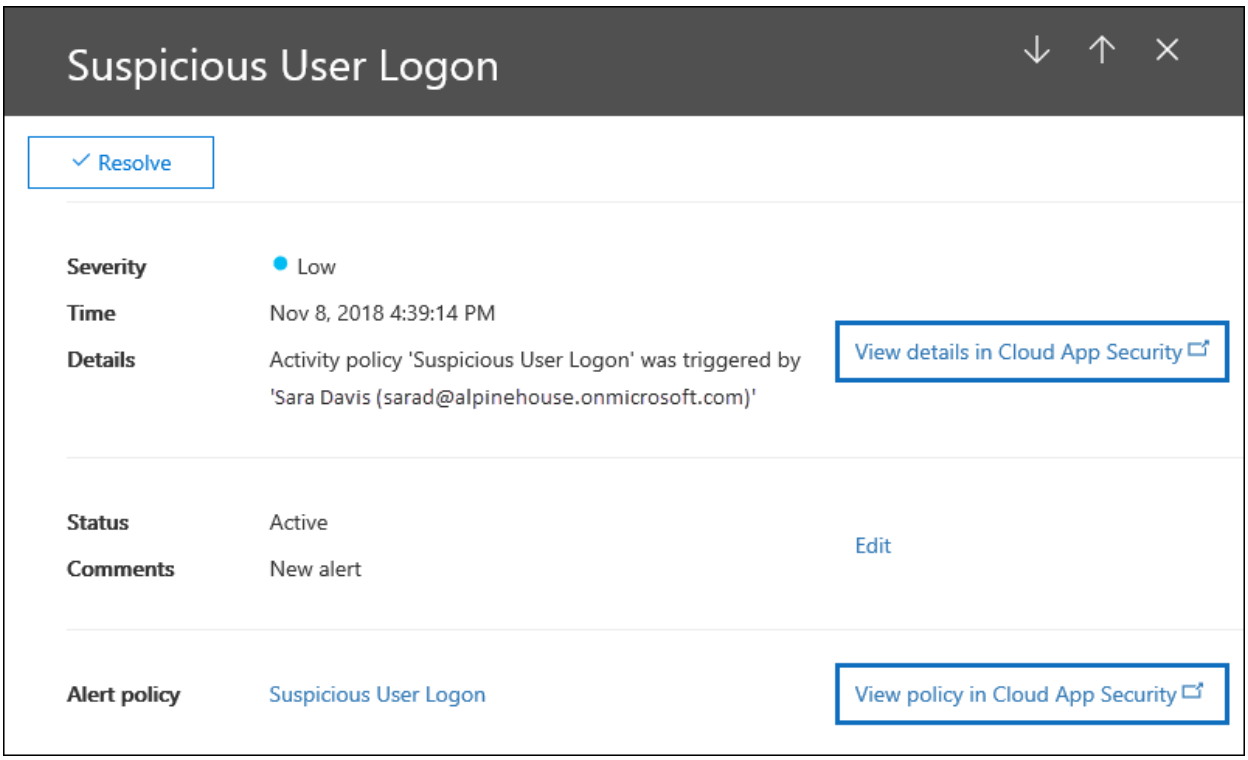

IMPORTANT

Changing the status of a Cloud App Security alert in the security and compliance center won't update the resolution status for the same alert in the Cloud App Security portal. For example, if you mark the status of the alert as **Resolved** in the security and compliance center, the status of the alert in the Cloud App Security portal is unchanged. To resolve or dismiss a Cloud App Security alert, manage the alert in the Cloud App Security portal.

Microsoft Compliance Manager

2/18/2021 • 5 minutes to read • Edit Online

In this article: Learn what Compliance Manager is, how it helps simplify compliance and reduce risk, and its key components.

What's new: the GA release of Compliance Manager

Compliance Manager is now generally available (GA) as an end-to-end compliance management solution inside the Microsoft 365 compliance center. With this release, Compliance Manager completes the transition from its previous location in the Microsoft Service Trust Portal. Compliance Manager is also now available to US Government Community (GCC) Moderate and GCC High customers.

What began as the public preview of Compliance Score has evolved into a centralized tool with enhanced compliance management capabilities and greater ease of use. The GA release brings a larger collection of pre-built assessments to help you scale your compliance activities.

Learn more about the GA release:

- Our frequently asked questions walk you through the evolution in greater detail.
- Read about GA feature enhancements in this blog post.

Watch the video below to learn how Compliance Manager can help simplify how your organization manages compliance:

What is Compliance Manager

Microsoft Compliance Manager is a feature in the Microsoft 365 compliance center that helps you manage your organization's compliance requirements with greater ease and convenience. Compliance Manager can help you throughout your compliance journey, from taking inventory of your data protection risks to managing the complexities of implementing controls, staying current with regulations and certifications, and reporting to auditors.

Compliance Manager helps simplify compliance and reduce risk by providing:

- Pre-built assessments for common industry and regional standards and regulations, or custom assessments to meet your unique compliance needs (available assessments depend on your licensing agreement; learn more).
- Workflow capabilities to help you efficiently complete your risk assessments through a single tool.
- Detailed step-by-step guidance on suggested improvement actions to help you comply with the standards and regulations that are most relevant for your organization. For actions that are managed by Microsoft, you'll see implementation details and audit results.
- A risk-based compliance score to help you understand your compliance posture by measuring your progress in completing improvement actions.

Your Compliance Manager dashboard shows your current compliance score, helps you see what needs attention, and guides you to key improvement actions. Below is an example of what your Compliance Manager

dashboard will look like:

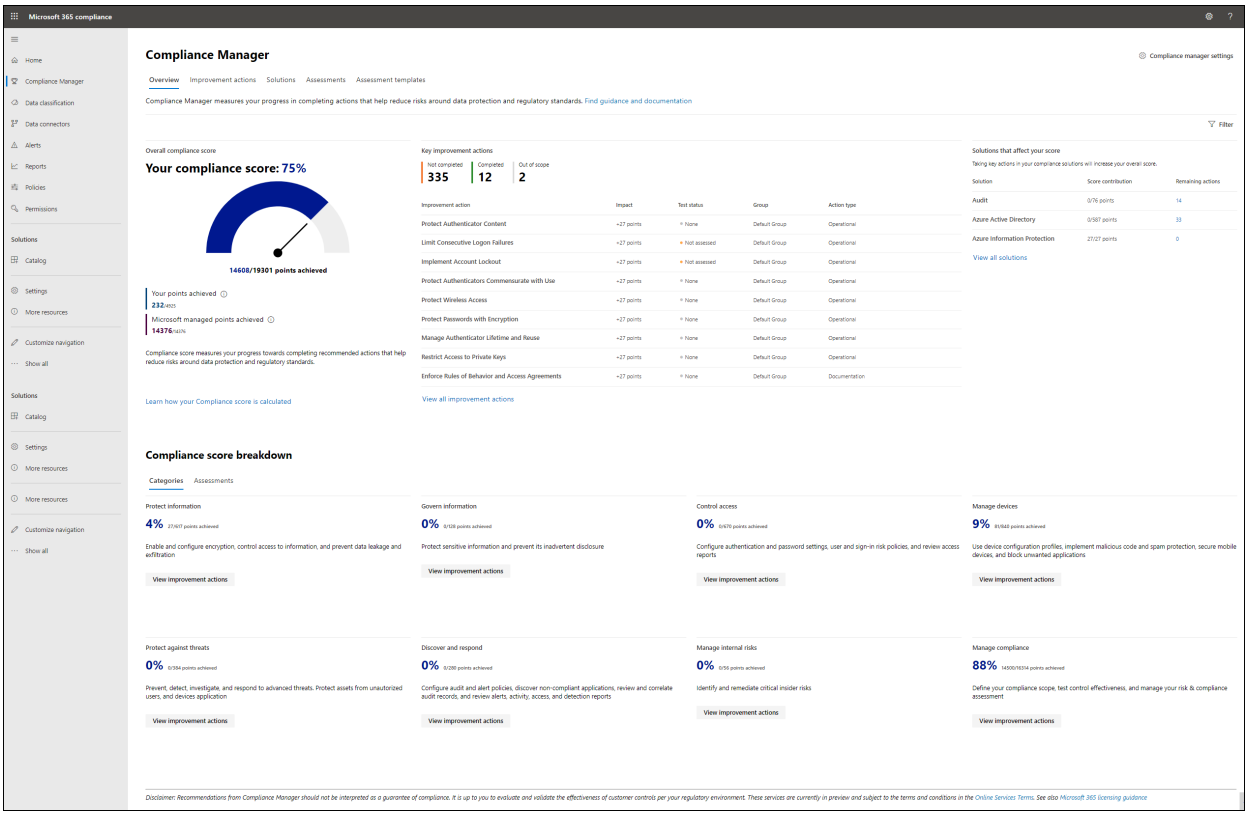

Understanding your compliance score

Compliance Manager awards you points for completing improvement actions taken to comply with a regulation, standard, or policy, and combines those points into an overall compliance score. Each action has a different impact on your score depending on the potential risks involved. Your compliance score can help prioritize which action to focus on to improve your overall compliance posture.

Compliance Manager gives you an initial score based on the Microsoft 365 data protection baseline. This baseline is a set of controls that includes key regulations and standards for data protection and general data governance.

Learn more

Understand how your compliance score is calculated.

Learn how to work with improvement actions.

Key elements: controls, assessments, templates, improvement actions

Compliance Manager uses several data elements to help you manage your compliance activities. As you use Compliance Manager to assign, test, and monitor compliance activities, it's helpful to have a basic understanding of the key elements: controls, assessments, templates, and improvement actions.

Controls

A control is a requirement of a regulation, standard, or policy. It defines how you assess and manage system configuration, organizational process, and people responsible for meeting a specific requirement of a regulation, standard, or policy.

Compliance Manager tracks the following types of controls:

1. **Microsoft managed controls**: controls for Microsoft cloud services, which Microsoft is responsible for implementing
2. **Your controls**: sometimes referred to as customer managed controls, these are controls implemented and

managed by your organization

3. **Shared controls**: these are controls that both your organization and Microsoft share responsibility for implementing

Learn more

Monitor progress of your controls.

Learn how Compliance Manager continuously assesses controls.

Assessments

An assessment is grouping of controls from a specific regulation, standard, or policy. Completing the actions within an assessment help you meet the requirements of a standard, regulation, or law. For example, you may have an assessment that, when you complete all actions within it, helps to bring your Microsoft 365 settings in line with ISO 27001 requirements.

Assessments have several components:

- **In-scope services**: the specific set of Microsoft services applicable to the assessment
- **Microsoft managed controls**: controls for Microsoft cloud services, which Microsoft implements on your behalf
- **Your controls**: sometimes referred to as customer managed controls, these are controls implemented and managed by your organization
- **Shared controls**: these are controls that both your organization and Microsoft share responsibility for implementing
- **Assessment score**: shows your progress in achieving total possible points from actions within the assessment that are managed by your organization and by Microsoft

When creating assessments, you'll assign them to a group. You can configure groups in whatever way is most logical for your organization. For example, you may group assessments by audit year, region, solution, teams within your organization, or some other way. Once you create groups, you can filter your Compliance Manager dashboard to view your score by one or more groups.

Learn more

Build and manage assessments in Compliance Manager.

Templates

Compliance Manager provides templates to help you quickly create assessments. You can modify these templates to create an assessment optimized for your needs. You can also build a custom assessment by creating a template with your own controls and actions. For example, you may want a template to cover an internal business process control, or a regional data protection standard that isn't covered by one of our 150+ pre-built assessment templates.

Learn more

View the list of assessment templates provided by Compliance Manager.

Get detailed instructions for creating and modifying templates for assessments.

Improvement actions

Improvement actions help centralize your compliance activities. Each improvement action provides recommended guidance that's intended to help you align with data protection regulations and standards. Improvement actions can be assigned to users in your organization to perform implementation and testing work. You can also store documentation, notes, and record status updates within the improvement action.

Learn more

Use improvement actions to manage your compliance workflow.

Learn how actions impact your compliance score.

Supported languages

Compliance Manager is available in the following languages:

- English
- Bahasa Indonesian
- Bahasa Malay
- Chinese (Simplified)
- Chinese (Traditional)
- Czech
- Danish
- Dutch
- Finnish
- French
- German
- Hebrew
- Hungarian
- Italian
- Japanese
- Korean
- Norwegian
- Polish
- Portuguese (Brazilian)
- Russian
- Spanish
- Swedish
- Thai
- Turkish

Next steps: set up and customize

Learn how to sign in, assign permissions and roles, configure settings, and personalize your dashboard view at Get started with Compliance Manager.

Then start customizing Compliance Manager to help you comply with industry standards that matter most to your organization by setting up assessments.

Compliance Manager quickstart

2/18/2021 • 2 minutes to read • Edit Online

In this article: Use this quickstart guide to help you along your journey of using Microsoft Compliance Manager to manage your organization's compliance with regulations, policies, and standards.

Compliance Manager provides intelligent and actionable data upon your first visit. Compliance Manager also has advanced capabilities for scaling your compliance when you're ready. Available assessments depend on your licensing agreement; learn more.

Whether you're coming to Compliance Manager for the first time, or are ready to use some of the advanced features, this guide can support you along your journey.

First visit: get to know Compliance Manager

Compliance Manager is located in the Microsoft 365 compliance center at https://compliance.microsoft.com. Your organization's global administrator will need to set up user permissions and assign roles before you start using Compliance Manager.

The first time you visit Compliance Manager, you'll see a compliance score for your organization. Compliance Manager is already assessing your current Microsoft 365 environment against the data protection baseline. The best way to start getting familiar with Compliance Manager is to understand what it's showing you, its key elements, and how to customize your dashboard.

Our Compliance Manager overview page is the best first stop for a comprehensive review of what Compliance Manager is and how it works. You may also want to jump right to key sections of our documentation using the links below:

- Understand your compliance score
- Overview of key elements: controls, assessments, templates, and improvement actions
- Understand the Compliance Manager dashboard
- Filter your dashboard view
- Learn about improvement actions
- Understand assessments
- Do a quick scan of your environment using the Microsoft Compliance Configuration Manager

Ramping up: configure Compliance Manager to manage your compliance activities

Once you're familiar with the basics, it's time to set things up to meet your organization's needs. You can start working with assessments and taking improvement actions to implement controls and improve your compliance score. Knowing how to perform all the activities at this stage can help your organization comply and demonstrate compliance with regulations across your industry and region. Visit the links below to dive in:

- Choose a pre-built assessment to create and manage your first assessment
- Understand how to use templates for building assessments
- Perform implementation and testing work on improvement actions to complete controls in your assessments
- Better understand how different actions impact your compliance score

Scaling up: use advanced functionality to meet your custom needs

When you're comfortable managing assessments in Compliance Manager, you can work with templates to modify a Compliance Manager assessment with your own actions and controls. You can also create your own custom assessment. Custom assessments are helpful for:

- Managing compliance for non-Microsoft 365 products such as third-party apps and services, on-premises applications, and other assets.
- Managing your own custom or business-specific compliance controls.

You can also set up automated testing of all or a subset of improvement actions. Visit the links below to understand more advanced functionality in Compliance Manager:

- Extend a Compliance Manager assessment by adding your own controls and improvement actions
- Create your own custom assessment
- Modify an existing template to add or remove controls and actions
- Set up automated testing of improvement actions
- Reassign improvement actions to another user

Get started with Compliance Manager

2/18/2021 • 18 minutes to read • Edit Online

In this article: This article helps you set up Compliance Manager. Learn how to **access** Compliance Manager, **set roles and permissions**, and configure **automatic testing of improvement actions**. Walk through **your Compliance Manager dashboard** and understand the main pages: the improvement actions page, the solutions page, the assessments page, and the assessment templates page.

Who can access Compliance Manager

Compliance Manager is available to organizations with Office 365 and Microsoft 365 licenses, and to US Government Community Cloud (GCC) Moderate and GCC High customers. Assessment availability and management capabilities depend on your licensing agreement. View service description details.

Before you begin

The Microsoft 365 global administrator for your organization will likely be the first user to access Compliance Manager. We recommend the global admin sign in and set user permissions as outlined below when visiting Compliance Manager for the first time.

Sign in

1. Go to the Microsoft 365 compliance center and **sign in** with your Microsoft 365 global administrator account.
2. Select **Compliance Manager** on the left navigation pane. You'll arrive at your Compliance Manager dashboard.

The direct link to access Compliance Manager is https://compliance.microsoft.com/compliancemanager.

Set user permissions and assign roles

Compliance Manager uses a role-based access control (RBAC) permission model. Only users who are assigned a role may access Compliance Manager, and the actions allowed by each user are restricted by role type.

Where to set permissions

The person holding the global admin role for your organization can set user permissions for Compliance Manager. Permissions can be set in the Office 365 Security & Compliance center as well as in Azure Active Directory (Azure AD).

> **NOTE**
> Customers in US Government Community (GCC) High environments can only set user permissions and roles for Compliance Manager in Azure AD. See below for Azure AD instructions and role type definitions.

To set permissions and assign roles in the Office 365 Security & Compliance center, follow the steps below:

1. Go to the Office 365 Security & Compliance Center and select **Permissions** on the left navigation.
2. Find the role group to which you want to add one or more users, and check the box to the left of the group name. (See the list of roles and related functions below. The role group names mimic the role

name.)

3. On the flyout pane for that group, select **Edit** under the **Members** header.
4. Select **Choose members**. Another flyout window will appear.
5. Select **+ Add** to choose one or more users to add to the group.
6. Select the checkbox next to the names you want to add, then select the **Add** button at the bottom.
7. When you're done assigning users, select **Done**, then select **Save**, then **Close**.

More about the Office 365 Security & Compliance Center

Learn more about permissions in the Office 365 Security & Compliance Center.

If you don't have access to the Office 365 Security and Compliance Center, or if you need to access the classic version of Compliance Manager in the Microsoft Service Trust Portal, the Admin settings in the Service Trust Portal provides another way to assign roles (view instructions). Be aware that such roles are more limited in their functionality.

More about Azure AD

To assign roles and set permissions in Azure AD, see Assign administrator and non-administrator roles to users with Azure Active Directory.

Users with Azure AD identities who don't have Office 365 or Microsoft 365 subscriptions won't be able to access Compliance Manager in the Microsoft 365 compliance center. To seek assistance in accessing Compliance Manager, contact cmresearch@microsoft.com.

Role types

The table below shows the functions allowed by each role in Compliance Manager. The table also shows how each Azure AD role maps to Compliance Manager roles. Users will need at least the Compliance Manager reader role, or Azure AD global reader role, to access Compliance Manager.

USER CAN:	COMPLIANCE MANAGER ROLE	AZURE AD ROLE
Read but not edit data	Compliance Manager Reader	Azure AD Global reader, Security reader
Edit data	Compliance Manager Contribution	Compliance Administrator
Edit test results	Compliance Manager Assessment	Compliance Administrator
Manage assessments, and template and tenant data	Compliance Manager Administration	Compliance Administrator, Compliance Data Administrator, Security Administrator
Assign users	Global Administrator	Global Administrator

Settings for automated testing and user history

The Compliance Manager settings in the Microsoft 365 compliance center allow you to enable and disable automatic testing of improvement actions. The settings also allow you to manage the data of users associated to improvement actions, including the ability to reassign improvement actions to a different user. Only people with a global administrator or Compliance Manager Administrator role can access the Compliance Manager settings.

> **NOTE**
>
> The automated testing feature is not available to customers in GCC High environments because Secure Score isn't available in these environments. GCC High customers will need to manually implement and test their improvement actions.

Set up automated testing

Some improvement actions in Compliance Manager are also monitored by Microsoft Secure Score. You can set up automated testing of actions that are jointly monitored, which means that when an action is tested and updated in Secure Score, those results synch with the same actions in Compliance Manager and count toward your compliance score.

Automatic testing is turned on by default for organizations new to Compliance Manager. When you first deploy Microsoft 365 or Office 365, it takes approximately seven days for Secure Score to fully collect data and factor it into your compliance score. When automated testing is turned on, the action's test date won't be updated, but its test status will update. When new assessments are created, scores automatically include Microsoft control scores and Secure Score integration.

The global administrator for your organization can change the settings for automated testing at any time. You can turn off automated testing for common improvement actions, or turn it on for individual actions. Follow the instructions below to change your automated testing settings.

To manage your automated testing settings:

1. Select **Settings** on the left navigation from anywhere in the Microsoft 365 compliance center.
2. On the settings page, select **Compliance Manager**.
3. Select **Automated testing** from the left navigation.
4. Select the applicable button to turn on automatic testing for all improvement actions, turn it off for all actions, or turn on by individual action.
5. If you select **Turn on per improvement action**, a list will show all the available improvement actions to choose from. Check the box next to any action you want automatically tested.
6. Select **Save** to save your settings. You'll receive a confirmation message at the top of your screen that your selection was saved. If you receive a failure notice, try again.

Note: Only the global administrator can turn on or off automatic updates for all actions. The Compliance Manager Administrator can turn on automatic updates for individual actions, but not for all actions globally.

Manage user history

The **Manage user history** settings help you quickly identify which users have worked with improvement actions in Compliance Manager. The identifiable user data associated with improvement actions includes any implementation and testing work done, documents they uploaded, and any notes they entered. Understanding and retrieving this type of data may be necessary for your organization's own compliance needs.

The user history settings also allow you to reassign all improvement actions from one user to another.

To find the user history settings:

1. Select Settings on the left navigation from anywhere in the Microsoft 365 compliance center.
2. On the settings page, select **Compliance Manager**.
3. Select **Manage user history** from the left navigation.

The **manage user history** page shows a list of all users by email address who are assigned to an improvement

action. Use the **Search** button to quickly find a specific user by typing in their email address.

To the right of each user's email address, the **Select** drop-down menu provides options to export a report, reassign improvement actions, or delete history. See each section below for details about each option.

Export a report of user history data

You can export an Excel file containing a list of improvement actions currently assigned to a user. The report also lists any evidence files uploaded by that user. This information can help you reassign open improvement actions.

The report reflects the improvement action's status as of its creation date. It's not a historical report of all previous changes to its status or assignment (learn how to export a report from your improvement actions page).

Follow the steps below to export a report by user:

1. Select **Settings** on the left navigation from anywhere in the Microsoft 365 compliance center.
2. On the settings page, select **Compliance Manager**.
3. Select **Manage user history** from the navigation at left.
4. Find your intended user by searching the list email addresses, or by selecting **Search** and entering the user's email address.
5. From the **Select** drop-down menu, choose **Export report**.
6. Once the Excel file of your report is generated, you can open it and save it to your local machine.

Reassign improvement actions to another user

You can reassign improvement actions from one user to another. When you reassign an action, the document upload history doesn't change, but the name of the user who originally uploaded the documentation no longer appears within the improvement action.

Follow the steps below to reassign improvement actions to another user:

1. Select **Settings** on the left navigation from anywhere in the Microsoft 365 compliance center.
2. On the settings page, select **Compliance Manager**.
3. Select **Manage user history** from the navigation at left.
4. Find a user by searching the list email addresses, or by selecting **Search** and entering that user's email address.
5. From the **Select** drop-down menu, choose **Reassign improvement actions**. The **Reassign improvement actions** flyout pane will appear.
6. In the **Search users** field, enter the name or email address of the user you want assign the improvement actions *to*.
7. When you see the name of your intended user under **Improvement actions will be assigned to**, select the user, then select **Assign actions**.
8. When the reassignment is complete, you'll see a confirmation message in the flyout pane confirming that all improvement actions from the previous user have been reassigned to the new user. If you receive a reassignment failure notice, close the window and try again. To close the flyout pane, select **Done**.

The new assignee receives an email that they've been assigned to an improvement action. The email contains a direct link into the improvement action's details page.

> **NOTE**
>
> If you reassign an action that has a pending update, the direct link to the action in the reassignment email will break if the update is accepted after reassignment. You can fix this by re-assigning the action to the user after the update is accepted. Learn more about updates to improvement actions.

Delete user history

Deleting a user's history will remove them as an owner of improvement actions, and will remove their name from all other fields in Compliance Manager. When you delete a user's history, the improvement actions they owned will not display an **Assigned to** value until a new user is assigned. Any documents uploaded to the improvement action will show **User removed** in place of the deleted user's name. Deleting user history is permanent.

To delete a user's history, follow the steps below:

1. Select **Settings** on the left navigation from anywhere in the Microsoft 365 compliance center.
2. On the settings page, select **Compliance Manager**.
3. Select **Manage user history** from the navigation at left.
4. Find a user by searching the list email addresses, or by selecting **Search** and entering that user's email address.
5. From the **Select** drop-down menu, choose **Delete history**.
6. A window appears asking you to confirm the permanent deletion of the user's history. To continue with deletion, select **Delete history**. To leave without deleting the history, select **Cancel**.
7. You'll arrive back at the **Manage user history** page with a confirmation message at the top that the history for the user was deleted.

Understand the Compliance Manager dashboard

The Compliance Manager dashboard is designed to provide you an at-a-glance view of your current compliance posture.

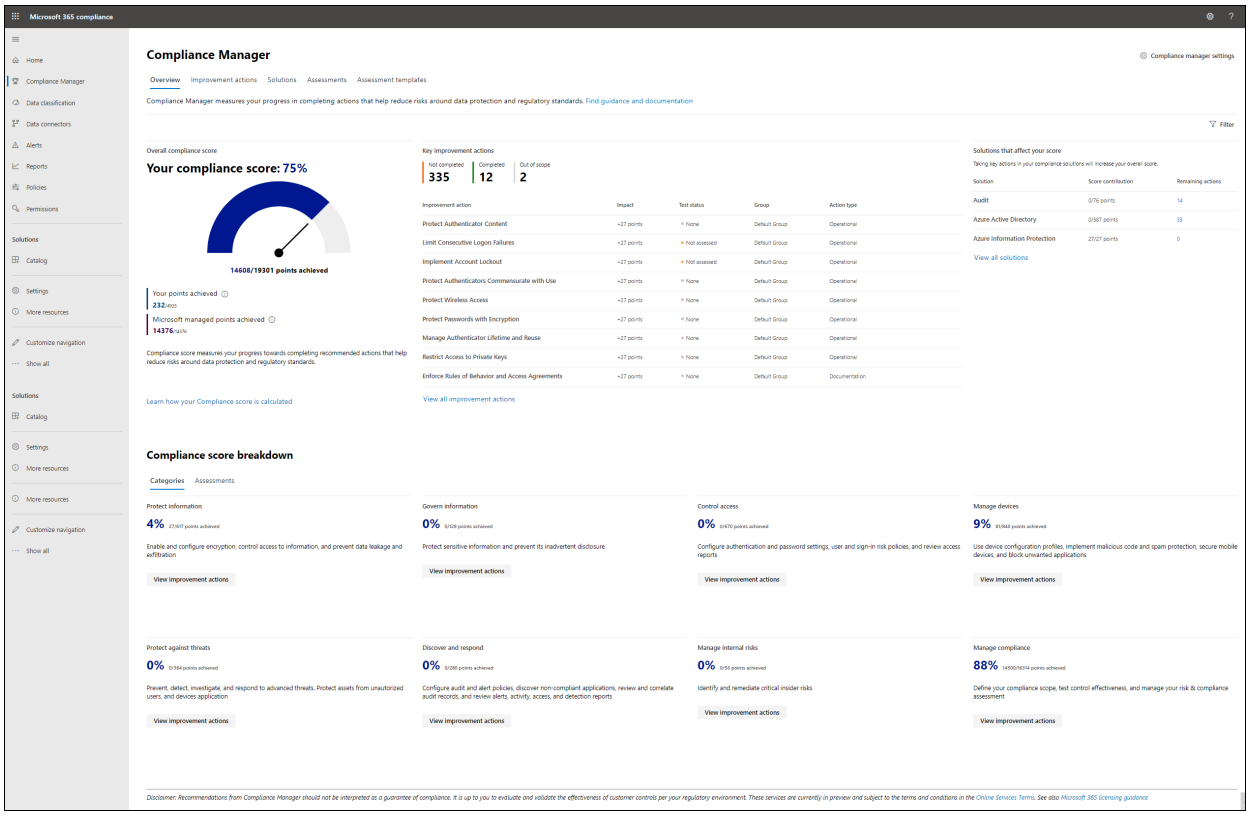

Overall compliance score

Your compliance score is featured prominently at the top. It shows a percentage based on points achievable for completing improvement actions that address key data protection standards and regulations. Points from Microsoft actions, which are managed my Microsoft, also count toward your compliance score.

When you come to Compliance Manager for the first time, your initial score is based on the Microsoft 365 data protection baseline. This baseline assessment, which is available to all organizations, is a set of controls that includes common industry regulations and standards. Compliance Manager scans your existing Microsoft 365 solutions and gives you an initial assessment based on your current privacy and security settings. As you add assessments that are relevant to your organization, your score becomes more meaningful for you.

Learn more: Understand how your compliance score is calculated.

Key improvement actions

This section lists the top improvement actions you can take right now to make the largest positive impact on your overall compliance score. Select **View all improvement actions** to go to your improvement actions page.

Solutions that affect your score

This section highlights solutions containing improvement actions that can positively impact your score, and the number of outstanding improvement actions in those solutions. Select **View all solutions** to visit your solutions page.

Compliance score breakdown

This section gives you a more detailed view of your score in two different ways:

- **Categories**: shows the percentage of your overall score within data protection categories, such as "protect information" or "manage devices."
- **Assessments**: shows the percentage of your progress in managing assessments for particular compliance and data protection standards, regulations, or laws, such as GDPR or NIST 800-53.

Filtering your dashboard view

You can filter your dashboard view to see only the items related to particular regulations and standards,

solutions, type of action, assessment groups, or data protection categories. Filtering your view in this way will also filter the score on your dashboard, showing how many points you've achieved out of total possible points based on your filter criteria.

To apply filters:

1. Select **Filter** on the upper-right side of the dashboard.
2. Select your filter criteria from the **Filters** flyout pane, then select **Apply**.

After you apply a filter, you'll see your score adjusted in real time. The compliance score percentage and breakdown information, and the improvement actions and solutions, now only pertain to data covered by your filter criteria. If you sign out of Compliance Manager, your filtered view remains when you sign back in.

To remove filters:

- At the **Applied filters** heading above your compliance score, select the **X** next to the individual filter you want to remove; or
- Select **Filter** on the upper-right side of your dashboard, then on the **Filters** flyout pane, select **Clear filters**.

Improvement actions page

Improvement actions are actions managed by your organization. Working with improvement actions helps to centralize your compliance activities and align with data protection regulations and standards. Each improvement action gives detailed implementation guidance and a link to launch you into the appropriate solution. Improvement actions can be assigned to users in your organization to perform implementation and testing work. You can also store documentation, notes, and record status updates within the improvement action.

View your improvement actions

The Compliance Manager dashboard shows your **key improvement actions.** To view all of your improvement actions, select the Improvement actions tab on your dashboard, which brings you to your improvement actions page. You can also select View all improvement actions underneath the list of key improvement actions on your dashboard to get to your improvement actions page.

The improvement actions page shows all of the improvement actions that are managed by your organization. Actions that are managed by Microsoft can be viewed within each assessment (learn more about Microsoft actions).

If you have a long list of actions on your improvement actions page, it may be helpful to filter your view. Select **Filter** at the upper-right corner of the actions list. When the **Filters** flyout pane appears, select your criteria based on regulations and standards, solution, and group. You can also customize your view by selecting **Group** in the upper-right corner. From the drop-down menu, select to view by group, solution, category, action type, or status.

The default view for this page does not show improvement actions with a test status of **Passed**. To view actions that have passed testing, check the **Passed** box in the Filters flyout pane. Only actions with a test status of **Passed** count toward your score. Some actions may show a **pending update label.** Learn more about updates to improvement actions.

The improvement actions page shows the following data points for each improvement action:

- **Points achieved**: the number of points achieved out of the total available by completing the action
- **Regulations**: the regulations or standards pertaining to the action
- **Group**: the group to which you assigned the action
- **Solutions**: the solution where you can go to perform the action
- **Assessments**: the assessments that contain the action

- **Categories**: the related data protection category (such as, protect information, manage devices, etc.)
- **Test status**:
  - **None** – no status update recorded
  - **Not assessed** - testing hasn't started
  - **Passed** - implementation successfully tested
  - **Failed low risk** - testing failed, low risk
  - **Failed medium risk** - testing failed, medium risk
  - **Failed high risk** - testing failed, high risk
  - **Out of scope** – the action is not in scope for the assessment and doesn't impact your score
  - **To be detected** - for manual test, indicates an action has been implemented but not tested; for automated test, indicates an action is waiting for automation result
  - **Could not be detected** - automated status can't be determined
  - **Partially tested** – automated scoring that awards partial points

Learn more: See how to assign and perform work on improvement actions.

Solutions page

The solutions page shows the share of earned and potential points as organized by solution. Viewing your remaining points and improvement actions from this view helps you understand which solutions need more immediate attention.

Find the solutions page by selecting the **Solutions** tab on your Compliance Manager dashboard. You can also select **View all solutions** underneath **Solutions that affect your score** in the upper-right section of your dashboard.

Filtering your solutions view

To filter your view of solutions:

1. Select **Filter** at the top-left corner of your assessments list.
2. On the **Filters** flyout pane, place a check next to the desired criteria (standards and regulations, solution, action type, Compliance Manager group, category).
3. Select the **Apply** button. The filter pane will close and you'll see your filtered view.

You can also modify your view to see assessments by group, product, or regulation by selecting the type of grouping from the **Group** drop-down menu above your assessments list.

Taking action from the solution page

The solutions page displays your organization's solutions that are connected to improvement actions. The table lists each solution's contribution to your overall score, the points achieved and possible within that solution, and the remaining number of improvement actions grouped in that solution that can increase your score.

There are two ways you can take action from this screen:

1. On the row of your intended solution, under the **Remaining actions** column, select the hyperlinked number. You'll see a filtered view of the improvement actions screen showing untested improvement actions for that solution.

2. On the row of your intended solution, under the **Open solution** column, select **Open**. You'll see the solution or location in the Microsoft 365 and Office 365 security and compliance centers where you can take the recommended action.

Assessments page

The assessments page lists all the assessments you set up for your organization. Your compliance score denominator is determined by all your tracked assessments. As you add more assessments, you'll see more improvement actions listed on your improvement actions page, and your compliance score denominator increases.

The assessments page summarizes key information about each assessment:

- **Assessment**: name of the assessment
- **Status**:
  - **Complete** - all controls have a status of "passed," or at least one is passed and the rest are "out of scope"
  - **Incomplete** – at least one control has a status of "failed"
  - **None** - all controls have have not been tested
  - **In progress** - improvement actions have any other status, including "in progress," "partial credit," or "undetected
- **Assessment progress**: the percentage of the work done toward completion, as measured by the number of controls successfully tested
- **Your improvement actions**: the number of completed actions to satisfy implementation of your controls
- **Microsoft actions**: the number of completed actions to satisfy implementation of Microsoft controls
- **Group**: name of the group the assessment belongs to
- **Product**: associated Microsoft 365 service
- **Regulation**: the regulatory standard, policy, or law that applies to the assessment

Filtering your assessments view

To filter you view of assessments:

1. Select **Filter** at the top-left corner of your assessments list.
2. On the **Filters** flyout pane, check your desired criteria.
3. Select the **Apply** button. The filter pane will close and you will see your filtered view.

You can also modify your view to see assessments by group, product, or regulation by selecting the type of grouping from the **Group** drop-down menu above your assessments list.

Default assessment

By default, you'll see the Data Protection Baseline assessment on the assessments page. Compliance Manager also provides several pre-built templates for building assessments.

Assessment templates page

A template is a framework for creating an assessment in Compliance Manager. The assessment templates page displays a list of templates and key details. The list includes templates provided by Compliance Manager as well as any templates your organization has modified or created. You can apply filters to find a template based on certification, product scope, country, industry, and who created it.

Select a template from its row to bring up its details page, which contains a description of the template and further information about certification, scope, and controls details. From this page you can select the appropriate buttons to create an assessment, export the template data to Excel, or modify the template.

Learn more: Read how to work with assessment templates.

Next step

Customize Compliance Manager by setting up assessments.

Build and manage assessments in Compliance Manager

11/2/2020 • 18 minutes to read • Edit Online

In this article: Learn how to customize Compliance Manager for your organization by creating and managing **assessments**. This article walks you through how to create assessments, how to organize them into **groups**, working with **controls**, accepting **updates**, and exporting assessment **reports**.

> **IMPORTANT**
>
> The assessments available to your organization depend on your licensing agreement. Review the details.

Introduction to assessments

Compliance Manager helps you manage compliance with assessments for the regulations and certifications that apply to your organization. Assessments are groupings of controls from a specific regulation, standard, or policy. Compliance Manager makes it easy to start tracking your compliance by providing pre-built assessments that cover a variety of industry and regional regulations and certifications.

Each assessment is created from an assessment template. Templates serve as a framework containing the necessary controls, improvement actions, and Microsoft actions for completing the assessment. Setting up the most relevant assessments for your organization can help you implement policies and operational procedures to limit your compliance risk.

All of your assessments are listed on the assessments page. Learn more about how to filter your view of your assessments and interpret status states.

Data Protection Baseline default assessment

To get you started, Microsoft provides a **default** assessment in Compliance Manager for the **Microsoft 365 data protection baseline**. This baseline assessment has a set of controls for key regulations and standards for data protection and general data governance. This baseline draws elements primarily from NIST CSF (National Institute of Standards and Technology Cybersecurity Framework) and ISO (International Organization for Standardization), as well as from FedRAMP (Federal Risk and Authorization Management Program) and GDPR (General Data Protection Regulation of the European Union).

This assessment is used to calculate your initial compliance score the first time you come to Compliance Manager, before you configure any other assessments. Compliance Manager collects initial signals from your Microsoft 365 solutions. You'll see at a glance how your organization is performing relative to key data protection standards and regulations, and see suggested improvement actions to take.

Compliance Manager becomes more helpful as you build and manage your own assessments to meet your organization's particular needs.

Assessment creation overview

There are three ways you can set up assessments:

1. Use a pre-built assessment.
2. Extend a pre-built assessment to suit your own needs.

3. Create your own custom assessment.

> **NOTE**
> Only users who hold a Global Administrator or Compliance Manager Administration role can create and modify assessments. Learn more about roles and permissions.

Use a pre-built assessment

Kickstart your compliance journey by choosing an assessment already set up by Compliance Manager. We provide a wide selection of templates for regulations and certifications that align to industries, regions, and common data protection standards, such as GDPR and ISO 27001. Templates contain the controls and improvement actions for helping you meet the requirements of a particular certification. You'll be asked to choose a template when you start building an assessment.

Extend a pre-built assessment to suit your needs

You can modify a Compliance Manager assessment—a process we refer to as "extending"—by adding your own controls and actions to better suit your organization's needs. For example, if you generally need to comply with HIPAA but require additional data protection or security controls, you can extend our HIPAA template by adding your own controls to it. See the instructions for extending a pre-built assessment.

Create your own custom assessment

You can create your own assessment entirely from scratch to track precisely what your organization needs. Creating your own assessment requires you to first create your own template for the assessment in Compliance Manager. See the instructions for creating your own custom assessment.

Understand groups before creating assessments

Before you create or modify assessments, it's important to understand how groups work. When you create an assessment, you'll need to assign it to a group during the process. That's why we recommend planning a grouping strategy for your assessments before you create assessments.

What are groups

Groups are containers that allow you to organize assessments. You can group assessments in a way that is logical to you, such as by year or regulation, or based on your organization's divisions or geographies. Below are examples of two groups and their underlying assessments:

- FFIEC IS assessment 2020
  - FFIEC IS
- Data security and privacy assessments
  - ISO 27001:2013
  - ISO 27018:2014

When two different assessments in the same group share improvement actions that are managed by you, any updates you make to an action's implementation details or status will automatically synchronize to the same action in any other assessment in the group. This synchronization allows you to implement one improvement action and meet several requirements across multiple regulations.

How to create a group

You create a group during the process of creating a new assessment.

Groups can't be created as standalone entities. A group must contain at least one assessment. In order to create a group, you need to first create an assessment to put in the group.

What to know when working with groups

- Group names must be unique within your organization.
- Groups don't have security properties. All permissions are associated with assessments.
- Once you add an assessment to a group, the grouping can't be changed.
- Related assessment controls in different assessments within the same group automatically update when completed.
- If you add a new assessment to an existing group, common information from assessments in that group are copied to the new assessment.
- Groups can contain assessments for the same certification or regulation, but each group can only contain one assessment for a specific product-certification pair. For example, a group can't contain two assessments for Office 365 and NIST CSF. A group can contain multiple assessments for the same product only if the corresponding certification or regulation for each one is different.
- Deleting an assessment breaks the relationship between that assessment and the group.
- Groups can't be deleted.
- When a change is made to an improvement that appears in multiple groups, that change is reflected in all instances of that improvement action.

Use a pre-built assessment

There are two starting points for creating an assessment from a Compliance Manager template.

You can begin the process from your assessments page by selecting the **Add assessment** button and then working through the assessment creation wizard. The steps for this process are below.

You can also start from your assessment templates page by finding the template you want and selecting it from the list to arrive at its details page. On the template details page, select **Create assessment**. You'll then enter the wizard with your template already selected.

To create an assessment

1. Know which group you'll assign your assessment to, or be prepared to create a new one for this assessment. Learn more about groups.

2. Go to your **assessments** page in Compliance Manager and select **Add assessment**. An assessment wizard will appear in a large flyout pane.

3. **Select a template**: Choose a template to serve as the basis for your assessment. Select the radio button next to your chosen template, then select **Next**.

4. **Name and group:** Enter a name for your assessment in the **Assessment name** field. Assessment names must be unique within groups. If the name of your assessment matches the name of another assessment in any given group, you'll receive an error asking you to create a different name.

5. Assign your assessment to a group. You can either:

   - Select **Use existing group** to assign it to a group you've already created; or
   - Select **Create new group** to create a new group and assign this assessment to it:
     - Determine a name for your group and enter it in the field beneath the radio button.
     - You can **copy data from an existing group**, such as implementation and testing details and documents, by selecting the appropriate boxes.

   When finished, select **Next**.

6. **Review and finish:** The last screen of the wizard shows the template, name, and group chosen for the assessment. You can edit any of these settings from the links on the screen, which take you back to the relevant steps in the wizard. When you're ready, select **Create assessment**.

7. The next screen confirms that you've successfully created your new assessment. Select **Done** to close the wizard, and your new assessment's details page will appear on the screen.

If you see an **Assessment failed** screen after selecting **Create assessment**, select **Try again** to re-create your assessment.

You can change the name of your assessment after you create it by selecting the **Edit name** button in the upper-right corner of the assessment's details page.

Extend a pre-built assessment

You can modify a pre-built assessment by adding your own controls and improvement actions to the assessment's template. This process is called "extending a Microsoft template" in Compliance Manager. When you extend the template of an assessment, it will receive any updates released by Microsoft, which may happen when there are changes to the related regulation or product (see Accepting updates to assessments).

You'll complete this process by starting at your **assessment templates** page rather than your **assessments** page.

Before you begin

To prepare for this process, you'll first need to assemble a specially formatted Excel spreadsheet to import the necessary template data. There are special requirements for the formatted Excel files used in the extension process. See these additional points to help prevent errors in the import process:

- Your spreadsheet should contain only the actions and controls you want to add to the assessment.
- The spreadsheet can't contain any of the controls or actions that already exist in the assessment you want to modify.
- Consider including "extension" in your template's title, for example, "GDPR – [your company name] extension." This makes it easier to identify in the list on your **assessment templates** page as distinct from the standard Microsoft-provided template or a custom template with a similar name.

After you format your spreadsheet, follow the steps below.

Steps for extending a Compliance Manager template

1. Go to your **Assessment templates** page and select **Create new template**. A template creation wizard will open.
2. Choose the type of template you want to create. In this case, select **Extend a Microsoft template**, then **Select**.
3. A template selection flyout pane appears on the right side of your screen. Use **Search** to apply filters for locating the template you want
4. Once you locate your template, select the radio button to the left of its name, then select **Save**.
5. The next screen shows the template you selected. If correct, select **Next**. (If incorrect, choose **Select a different template** to choose again.)
6. At the **Upload file** screen, select **Browse** to find and upload your formatted Excel file containing all the required template data.
7. If there are no problems with your file, the next screen shows the name of the file uploaded. Select **Next** to continue (if you need to change the file, select **Upload a different file**).
    - If there's a problem with your file, an error message at the top explains what's wrong. You'll need to fix and re-upload your file. Errors will result if your spreadsheet is formatted improperly, or if there's invalid information in certain fields.

8. The **Review and finish** screen shows the number of improvement actions and controls and the maximum score for the template. When ready to approve, select **Next**. (If you need to make changes, select **Upload a different file**.)
9. The last screen confirms a new template has been created. Select **Done** to exit the wizard.
10. You'll arrive at your new template's details page. From here you can create your assessment by selecting **Create assessment**. For guidance, start at step #4 in the assessment creation instructions above.

Create your own custom assessment

Creating a custom assessment in Compliance Manager requires you to create your own template. To create your own template, you'll first assemble a formatted Excel spreadsheet to import the necessary template data. It also helps to decide ahead of time which group you'll assign your assessment to when you create it (learn more about groups).

Follow the steps below to create your custom assessment:

1. **Format your Excel file.** Begin by formatting your template data into an Excel spreadsheet using these instructions.
2. **Create your template** by following these instructions.
3. **Create your assessment** from the template. You can begin by opening the template's details page and selecting **Create assessment**, or go to your **assessments** page and select **Create assessment**.
4. An assessment creation wizard will appear in a large flyout pane. From here, you can follow the guidance starting at step #3 of the assessment creation instructions, using your new custom template for your assessment.

Delete an assessment

Deleting an assessment removes it from the list on your assessments page. Note these important points about deleting assessments:

- **Deleting an assessment is permanent; you cannot get it back.** If you want to use the same assessment again, you'll need to re-create it.
- If the improvement actions in the assessment don't appear in any other assessment, they'll be deleted when the assessment is deleted.
- We recommend exporting a report of the assessment before you permanently delete it.

To delete an assessment, follow the steps below:

1. From your **assessments** page, select the assessment you wish to delete to open that assessment's details page.
2. Select **Delete assessment** in the upper-right corner of your screen.
3. A window will appear asking you to confirm that you want to permanently delete the assessment. Select **Delete assessment** to close the window. You'll get a confirmation window that your assessment was deleted from Compliance Manager.

If you delete the only assessment in a group, then that group is also deleted from Compliance Manager.

NOTE

You can't delete all of your assessments. Organizations need at least one assessment for Compliance Manager to function properly. If the assessment you want to delete is the only one, add another assessment before deleting the other assessment.

Monitor assessment progress and controls

Each assessment has a details page that gives an at-a-glance view of your progress in completing the assessment. The page shows your progress in completing controls, and the test status of key improvement actions within those controls.

Overview tab

The overview tab contains a graph showing your percentage toward completion of the assessment. This graph contains a breakdown of points from actions you own, and points from actions owned by Microsoft, so you can see how many more points you need to complete the assessment.

The key improvement actions for controls in the assessment are listed in order of greatest potential impact to earn points. The associated graph details the aggregated test status of your improvement actions so you can quickly gauge what has been tested and what still needs to be done.

To access individual improvement actions, visit the **Controls** tab or the **Your improvement actions** tab.

Controls tab

The controls tab displays detailed information for each control mapped to the assessment. A **control status breakdown** chart shows the status of controls by family, so you can see at a glance which groupings of controls need attention.

Beneath the chart, a table lists detailed information about each control within the assessment. Controls are grouped by control family. Expand each family name to reveal the individual controls it contains. The information listed for each control includes:

- **Control title**
- **Status**: reflects the test status of the improvement actions within the control
  - **Passed** - all improvement actions have a test status of "passed," or at least one is passed and the rest are "out of scope"
  - **Failed** - at least one improvement action has a test status of "failed"
  - **None** - all improvement actions have not been tested
  - **Out of scope** - all improvement actions are out of scope for this assessment
  - **In progress** - improvement actions have a status other than the ones listed above, which could include "in progress," "partial credit," or "undetected"
- **Control ID**: the control's identification number, assigned by its corresponding regulation, standard, or policy
- **Points achieved**: the number of points earned by completing actions, out of the total number of achievable points
- **Your actions**: the number of your actions completed out of the total number of actions to be done
- **Microsoft actions**: the number of actions completed by Microsoft

To view a control's details, select it from its row in the table. The control details page shows a graph indicating the test status of the actions within that control. A table below the graph shows key improvement actions for that control.

Select an improvement action from the list to drill into the improvement action's details page. The details pages shows test status, implementation notes, and launch into the recommended solution.

Your improvement actions tab

The tab for your improvement actions lists all the controls in the assessment that are managed by your organization. The status bar details the aggregated test status of your improvement actions in the assessment so you can quickly gauge what has been tested and what still needs to be done. Beneath the bar is the full list of improvement actions and key details, including: test status, the number of potential and earned points, associated regulations and standards, applicable solution, action type, and control family. Learn more about how actions contribute to your compliance score.

Select an improvement action to view its details page, and select the **Launch now** link to open the solution to take action.

Microsoft actions tab

The Microsoft actions tab lists all the actions in the assessment that are managed by Microsoft. The list shows key action details, including: test status, points that contribute to your overall compliance score, associated regulations and standards, applicable solution, action type, and control family. Select an improvement action to view its details page.

Learn more about how controls and improvement actions are tracked and scored.

Accepting updates to assessments

When an update is available for an assessment, you'll see a notification and have the option to accept the update or defer it for a later time.

What causes an update

An assessment update occurs when there are underlying template changes that impact scoring. Changes may involve adjusting control mapping or other guidance based on regulatory changes or product changes. Assessment updates can originate from your organization (such as, when a custom template is modified) as well as from Microsoft.

If Microsoft updates a Compliance Manager template that you extended, your assessment will inherit those updates once you accept them. Your assessment will retain the additional attributes you applied to the assessment when you extended it.

Custom assessments that you create do not receive any template updates from Microsoft. Custom assessments can receive improvement action updates, but any Microsoft updates to control mapping between assessments and improvement actions don't apply to custom templates.

> **NOTE**
> Updates to assessments apply only at the group level. If you have two assessments built from the same template that exist in two different groups, each assessment will have a pending update notification, and you'll need to accept the update to each assessment in its respective group individually.

Where you'll see assessment update notifications

The assessment details page also shows a **Pending update** label next to the assessment with an update. Select that assessment to get to its details page.

A message near the top of the assessment details page shows that an update is available for that assessment. Select the **Review update** button in the banner to review the specific changes and accept or defer the update.

The assessment details page may also list improvement actions that have a **Pending update** label next to them. Those updates are for specific changes to the improvement actions themselves and need to be accepted separately. Visit Accepting updates to improvement actions to learn more.

Review update to accept or defer

After selecting **Review update** from the assessment details page, a flyout pane appears on the right side of your screen. The flyout pane provides the key details below about the pending update:

- The template title
- Source of the update (Microsoft, your organization, or a specific user)
- The date the update was created
- An overview explaining the update
- Specific details about the changes, including the impact to your compliance score, the amount of progress toward completion of the assessment, and the specific number of changes to improvement actions and controls.

Selecting the **Updated template** link will download an Excel file containing control data for the version of the template with the pending updates. Selecting the **Current template** link downloads a file of the existing template without the changes.

To accept the update and make the changes to your assessment, select **Accept update**. Accepted changes are permanent.

If you select **Cancel**, the update won't be applied to the assessment. However, you'll continue to see the **Pending update** notification until you accept the update.

Why we recommend accepting updates

Accepting updates helps ensure you have the most updated guidance on using solutions and taking appropriate improvement actions to help you meet the requirements of the certification at hand.

Why you might want to defer an update

If you're in the middle of completing an assessment, you may want to ensure you've finished work on it before you accept an update to the assessment that could disrupt control mapping. You can defer the update for a later time by selecting **Cancel** on the review update flyout pane.

Export an assessment report

You can export an assessment to an Excel file for compliance stakeholders in your organization or for external auditors and regulators. On your assessment details page, select the **Generate report** button near the top of the page, which creates an Excel file you can save and share.

The report is a snapshot of the assessment as of the date and time of the export. It contains the details for controls managed by both you and Microsoft, including implementation status, test date, and test results.

Working with assessment templates in Compliance Manager

2/18/2021 • 13 minutes to read • Edit Online

In this article: Understand **how templates work** and **how to manage them** from your assessment templates page. Get instructions for **creating** new templates, **modifying** existing templates, **formatting your template data with Excel**, and exporting template **reports**.

> **IMPORTANT**
>
> The assessment templates that are available to your organization depends on your licensing agreement. Review the details.

Templates overview

A template is a framework for creating an assessment in Compliance Manager. They contain the controls for meeting the requirements of a certification using a certain product. Compliance Manager provides a comprehensive set of templates to help your organization comply with national, regional, and industry-specific requirements governing the collection and use of data.

List of pre-built templates for assessments

Compliance Manager provides templates for building assessments to help you comply with various regulations and standards. View the list of templates provided by Compliance Manager. New templates are added regularly, so check the list often.

Viewing and managing templates from the assessment templates page

The assessment templates page in Compliance Manager displays a list of templates and key details. The list includes templates provided by Compliance Manager as well as any templates your organization has modified or created. You can apply filters to find a template based on certification, product scope, country, industry, who created it, and whether the template is enabled for assessment creation.

Select a template from its row to bring up its details page. This page contains a description of the template and further information about certification, scope, and controls details. From this page you can select the appropriate buttons to create an assessment, export the template data to Excel, or modify the template.

Creating and modifying templates overview

To modify an existing template or to create your own new template, you'll use a specially formatted Excel spreadsheet (download an example) to assemble the necessary control data. After completing the spreadsheet, you import it into Compliance Manager during the process of creating or modifying a template.

> **NOTE**
>
> The spreadsheet has a specific format and schema that must be used, or it will not import correctly into Compliance Manager. The formatting instructions are below.

Required roles

Only users who hold a Global Administrator or Compliance Manager Administration role can create and modify templates. Learn more about roles and permissions.

Create a new template

To create your own new template (used for building custom assessments), follow the steps below.

1. Go to your **assessment templates** page in Compliance Manager.
2. Select **Create new template**. A template creation wizard will open.
3. Choose the type of template you want to create. In this case, select **Create a custom template**, then select **Next**.
4. At the **Upload file** screen, select **Browse** to find and upload your formatted Excel file containing all the required template data (see instructions for properly formatting your file).
5. If there are no problems with your file, the name of the file uploaded will be displayed. Select **Next** to continue. (If you need to change the file, select **Upload a different file**).
   - If there's an error with your file, an error message at the top explains what's wrong. You'll need to fix your file and upload it again. Errors will result if your spreadsheet is formatted improperly, or if there's invalid information in certain fields (refer again to the formatting instructions).
6. The **Review and finish** screen shows the number of improvement actions and controls and the maximum score for the template. When ready to approve, select **Create template.** (If you need to make changes, select **Back**.)
7. The last screen confirms a new template has been created. Select **Done** to exit the wizard.
8. You'll arrive at your new template's details page, where you can create your assessment.

Formatting your template data with Excel

The Excel spreadsheet used to create templates contains four tabs, three of which are required:

1. Template (required)
2. ControlFamily (required)
3. Actions (required)
4. Dimensions (optional)

When filling out your spreadsheet with template data, the spreadsheet **must include the tabs in the order listed above**, otherwise your data won't successfully import to a template.

Template tab

The **Template** tab is required. The information in this tab provides metadata about the template. There are four required columns. The columns must retain the order on the Excel sheet as listed below. You can add your own column **after** the four columns to provide your own dimensions. If you do this, be sure to add them to the **Dimensions** tab using the instructions below.

- **title**: This is the title for your template, which must be unique. It can't share a name with another template you have in Compliance Manager, including your own templates or a Compliance Manager template.
- **product**: This is a required dimension. List the product associated with the template.
- **certification**: This is the regulation you're using for the template.
- **inScopeServices**: These are the services within the product that this assessment addresses (for

example, if you listed Office 365 as the product, Microsoft Teams could be an in-scope service). You can list multiple services separated by two semi-colons.

> **NOTE**
>
> The data you insert in the **product** and **certification** cells can't be edited after you import the spreadsheet to create or customize a template. Also, a group can't contain two assessments that have the same **product/certification** combination. You can have multiple templates with the same product/certification combination.

ControlFamily tab

The **ControlFamily** tab is required. The required columns in this tab, which must follow the order provided in the sample spreadsheet, are:

- **controlName**: This is the control name from the certification, standard, or regulation, which is typically some type of ID. Control names must be unique within a template. You can't have multiple controls with the same name in the spreadsheet.
- **controlFamily**: Provide a word or phrase for the controlFamily, which identifies a broad grouping of controls. A controlFamily doesn't have to be unique; it can be listed more than once in a spreadsheet. The same controlFamily can also be listed in multiple templates, though they have no relation to each other. Every controlFamily must be mapped to at least one control.
- **controlTitle**: Provide a title for the control. Whereas the controlName is a reference code, the title is a rich text format typically seen in the regulations.
- **controlDescription**: Provide a description of the control.
- **controlActionTitle**: This is the title of an action that you want to relate to this control. You can add multiple actions by separating by two semi-colons with no space in between. Every control you list must include at least one action, and the action must exist (which means you can list an action that you list on the **Actions** tab of the same spreadsheet, an action that exists in a different template, or an action created by Microsoft). Different controls can reference the same action.

Actions tab

The **Actions** tab is required. It designates improvement actions managed by your organization and not those of Microsoft, which already exist in Compliance Manager. The required columns for this tab, which must follow the order provided in the sample spreadsheet, are:

- **actionTitle**: This is the title for your action and is a required field. The title you provide must be unique. **Important**: if you reference an action you own that already exists (such as in another template) and you modify any of its elements in the subsequent columns, those changes will propagate to the same action in other templates.
- **implementationType**: In this required field, list one of the three implementation types below:
  - **Operational** - actions implemented by people and processes to protect the confidentiality, integrity, and availability of organizational systems, assets, data, and personnel (example: security awareness and training)
  - **Technical** - actions completed through the use of technology and mechanisms contained in the hardware, software, or firmware components of the information system to protect the confidentiality, integrity, and availability of organizational systems and data (example: multi-factor authentication)
  - **Documentation** - actions implemented through documented policies and procedures establishing and defining the controls required to protect the confidentiality, integrity, and availability of organizational systems, assets, data, and personnel (example: an information security policy)
- **actionScore**: In this required field, provide a numeric score value for your action. It must be a whole number ranging from 1 to 99; it cannot be 0, null, or blank. The higher the number, the greater its value

toward improving your compliance posture. The image below demonstrates how Compliance Manager scores controls:

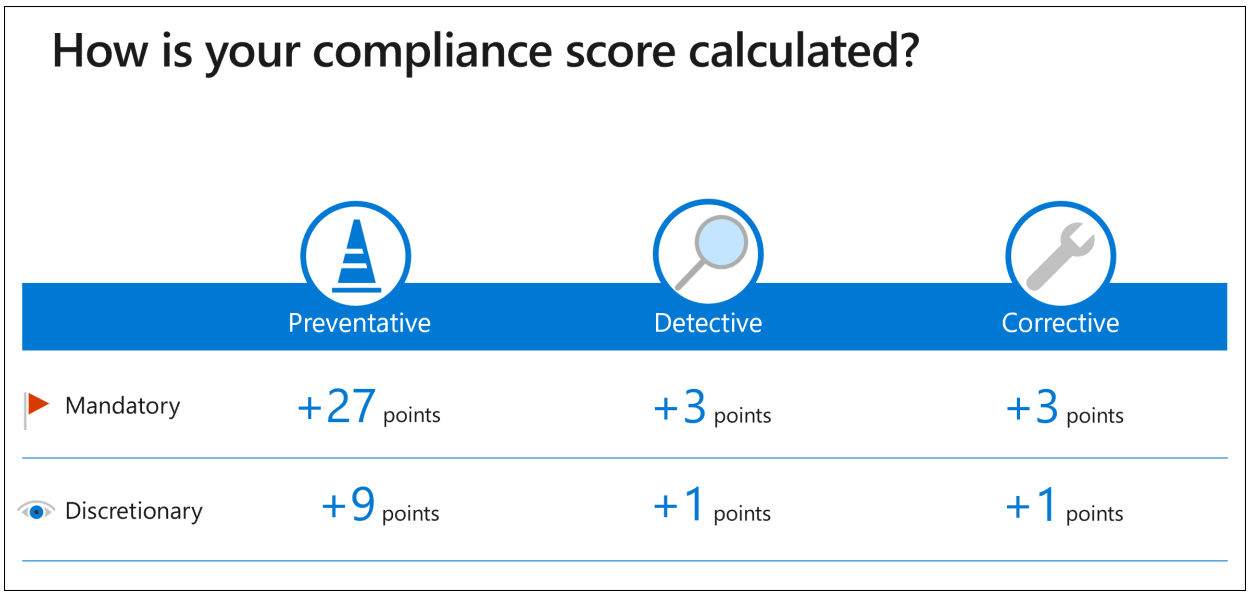

- **actionDescriptionTitle**: This is the title of the description and is required. This description title allows you to have the same action in multiple templates and surface a different description in each template. This field helps you clarify what template the description is referencing. In most cases, you can put the name of the template you're creating in this field.
- **actionDescription**: Provide a description of the action. You can apply formatting such as bold text and hyperlinks. This is required field.
- **dimension-Action Purpose**: This is an optional field. If you include it, the header must include the "dimension-" prefix. Any dimensions you include here will be used as filters in Compliance Manager and appear on the improvement actions details page in Compliance Manager.

Dimensions tab

The **Dimensions** tab is optional. However, if you reference a dimension elsewhere, you need to specify it here if it does not exist in a template you've already created or in a Microsoft template. The columns for this tab are listed below:

- **dimensionKey**: list as "product", "certifications," "action purpose"
- **dimensionValue**: examples: Office 365, HIPPA, Preventative, Detective

You can view your existing dimensions by going to **Tenant Management** and selecting the **Dimensions** tab. Also, anytime you export an existing template, the exported spreadsheet will have the **Dimensions** tab, which lists all the dimensions used in the template.

Modify a template

You may want to modify a template you've already created, such as to add controls, or add or remove improvement actions. The process is similar to the template creation process in that you'll upload formatted Excel file with your template data.

However, there are particular details to be aware of as you format your file with changes to existing template data. **We recommend you review these instructions carefully to ensure you don't overwrite any existing data that you want to retain.**

Template modification process steps

To modify a template, follow the steps below:

1. From your **assessment templates** page, select the template you want to modify, which will bring up its

details page.

2. Select **Export to Excel**. An Excel file with all your template data will download. Save the file to your local machine.
3. Make your template changes by modifying the Excel file using the instructions below.
4. When you're done making changes to your Excel file, save the file.
5. At your template's details page, select **Modify template** to initiate the modification wizard.
6. At the **Upload file** screen, select **Browse** to find and upload your Excel file.
7. If there are no problems with your file, the next screen shows the name of the file uploaded. Select **Next** to continue (if you need to change the file, select **Upload a different file**).
    - If there's a problem with your file, an error message at the top explains what's wrong. You'll need to fix your file and upload it again. Errors will result if your spreadsheet is formatted improperly, or if there's invalid information in certain fields.
8. The **Review and finish** screen shows the number of improvement actions and controls and the maximum score for the template. When ready to approve, select **Next**.
9. The last screen confirms that the template has been modified. Select **Done** to exit the wizard.

Your template will now include the changes you made. Any assessments that use this modified template will now show pending updates, and you'll need to accept the updates to the assessments to reflect the changes made in the template. Learn more about updates to assessments.

> **NOTE**
> If you use Compliance Manager in a language other than English, you'll notice that some text appears in English when you export a template to Excel. The titles of actions (both your improvement actions and Microsoft actions) must be in English to be recognized by controls. If you make changes to an action title, be sure to write it in English so that the file imports correctly.

Formatting your Excel file to modify a template

Jump to a section below to quickly find the instructions you need:

- Edit the main template attributes
- Add an improvement action
- Edit an improvement action's information
- Change an improvement action's name
- Remove an improvement action
- Remove a control

Edit the main template attributes

On the **Templates** tab, you can edit anything in the **title** column, the **inScopeServices** column, and in any other column you may have added. However, you can't edit anything in the **product** or **certification** columns.

Add an improvement action

1. Go to the **Actions** tab. Add your information in the required fields in the first empty row underneath your existing actions.
2. Go to your **ControlFamily** tab. Find the row containing the control your improvement action maps to. Add your new action to the **controlActionTitle** column in that row (remember to separate multiple actions in this field with two semi-colons, no space in between).
3. Save your spreadsheet.

Edit an improvement action's information

You can change any improvement action's information *except for its title*. You can edit any cell from columns B onward, and when you import the file back into the template, the improvement actions in that template will now contain the updated data.

You cannot edit the **actionTitle** (column A) because if you do, Compliance Manager considers this to be a new improvement action. If you want to change an improvement action's name, see the instructions immediately below.

Change an improvement action's name

If you want to change the name of an improvement action, you have to explicitly designate in the spreadsheet that you are replacing an existing name with a new name. Follow these steps:

1. In the **Actions** tab of your spreadsheet, add a new column to the spreadsheet after column A.
2. In this new column, which is now column B, put as its header in row 1: **oldActionTitle**.
3. Copy the contents of column A and paste them into column B. This puts your existing improvement action titles, which are what you want to change, into column B.
4. In column A, **actionTitle**, delete the old name and replace it with the new name for your improvement action.

Note that action titles, both for your improvement actions and for Microsoft actions, must be written in English in order to be recognized when referenced in controls.

Remove an improvement action

To remove an improvement action from a template, you'll need to remove it from every control that references it. Follow the steps below to modify your spreadsheet:

1. On the **ControlFamily** tab, search for for the title of the improvement action you want to remove.
2. Delete the improvement action's title in the cells where it appears. If the improvement action is the only action on that row, delete the entire row (which removes the control).
3. On the **Actions** tab, delete the row that contains the improvement action you're deleting.
4. Save your spreadsheet.

When you import your spreadsheet back into the template, your improvement action will be removed from the template.

Removing an improvement action from a template does not completely remove the improvement action from Compliance Manager. That action can still be referenced by another template.

Remove a control

To remove a control, modify your spreadsheet by following the steps below, then re-import your spreadsheet:

1. On the **ControlFamily** tab, find the control you want to remove in the **controlName** column.
2. Delete the row for that control.
   - If this deleted control contains improvement actions that aren't referenced by any other control, you'll need to remove those improvement actions from the **Actions** tab. Otherwise, you'll receive a validation error.
3. Save your spreadsheet.

When you import your spreadsheet back into the template, your control will be removed from the template.

Export a template

You can export an Excel file that contains all of a template's data. You'll need to export a template in order to modify the template, as this will be the Excel file you edit and upload in the modification process.

To export your template, go to your template details page and select the **Export to Excel** button.

Note that when exporting a template you extended from a Compliance Manager template, the exported file will only contain the attributes you added to the template. The exported file won't include the original template data provided by Microsoft. To get such a report, see the instructions for exporting an assessment report.

Compliance Manager templates list

2/18/2021 • 11 minutes to read • Edit Online

In this article: View the comprehensive list of **templates** available for creating assessments in Compliance Manager.

> **IMPORTANT**
> The assessment templates that are available to your organization depends on your licensing agreement. Review the details.

Overview

Microsoft Compliance Manager provides a comprehensive set of templates for creating assessments. These templates can help your organization comply with national, regional, and industry-specific requirements governing the collection and use of data.

Templates are added to Compliance Manager as new laws and regulations are enacted. Compliance Manager updates its templates when the underlying laws or regulations change. Learn more about how review and accept updates.

View detailed guidance on working with templates to create your assessments.

List of templates and where to find them

Below is the complete list of templates in Compliance Manager. Each template details page provides information about the regulation or standard to which it applies. The links in the template names below take you to related documentation about that standard, regulation, or law.

Where to find your templates

In Compliance Manager, go to your **Assessment templates** page. You'll see a list of all the templates available to your organization.

- **Included templates** are templates included as part of your organization's licensing agreement.
- **Premium templates** displays additional templates your organization may choose to obtain (refer to the service terms).

Read more about how to view and manage your templates.

Included templates

- Microsoft Data Protection Baseline
- European Union GDPR (Microsoft 365, Office 365, Intune)
- ISO 27001:2013
- NIST 800-53 Rev.4

Premium templates

- AICPA/CICA Generally Accepted Privacy Principles (GAPP) (Microsoft 365)
- Alabama - Policy 621: Data Breach Notification (Microsoft 365)

- Alaska - Chapter 48 - Personal Information Protection Act (Microsoft 365)
- Albania - The Law on the Protection of Personal Data No. 9887
- Antigua and Barbuda-Data Protection Act /2013 (Microsoft 365)
- Appendix III to OMB Circular No. A-130 - Security of Federal Automated Information Resources
- Argentina - Personal Data Protection Act 25.326 (Microsoft 365)
- Arkansas - Personal Information Protection Act (Microsoft 365)
- Asia Pacific Economic Cooperation (APEC) Privacy Framework
- Australia - ASD Essential 8 (Microsoft 365)
- Australia - National Archives Act
- Australia - Public Records Office Victoria Recordkeeping Standards (Microsoft 365)
- Australia - Spam Act 2003 (Microsoft 365)
- Australia Privacy (Credit Reporting) Code 2014 (Version 2.1)
- Australian Energy Sector Cyber Security Framework (AESCSF) (Microsoft 365)
- Australian Information Security Registered Assessor Program (IRAP) Version 2 (Microsoft 365)
- Australian Prudential Regulation Authority CPS (Microsoft 365)
- Austrian Telecommunications Act 2003 (Microsoft 365)
- Bahamas - Data Protection Act (Microsoft 365)
- Barbados - Data Protection Bill 2019 (Microsoft 365)
- Belarus Law On Information, Informatization and Protection of information (Microsoft 365)
- Belgium NBB Dec 2015 (Microsoft 365)
- Bermuda - Electronic Transaction Act (Microsoft 365)
- Bosnia and Herzegovina Law on the Protection of Personal Data
- Brazil - Consumer Protection Code Law No. 8078 (Office 365)
- Brazil - General Data Protection Law (LGPD) (Microsoft 365)
- Bulgaria Law for Protection of Personal Data 2002 (Microsoft 365)
- California - Civil Code Section 1798
- California - Database Breach Act (California SB 1386)
- California - Education Code-EDC, Title 3, Division 14, Part 65, Chapter 2.5- Social Media Privacy
- California - SB-327 Information privacy: connected devices (Microsoft 365)
- California Consumer Credit Reporting Agencies Act (Microsoft 365)
- California Consumer Privacy Act (CCPA) (Microsoft 365)
- Canada - Breach of Security Safeguards Regulations (Microsoft 365)
- Canada - British Columbia - Information Privacy & Security - FOIPPA (Microsoft 365)
- Canada - Office of the Superintendent of Financial Institutions (Microsoft 365)
- Canada - Personal Health Information Protection Act (PHIPA) (Microsoft 365)
- Canada - Personal Information Protection and Electronic Documents Act (PIPEDA) (Microsoft 365)
- Canada - Protected B
- Canada Cybersecure (Microsoft 365)
- CAN-SPAM Act (Microsoft 365)
- CDSA Content Protection & Security Standard (Microsoft 365)
- CFR - Code of Federal Regulations Title 21 (Microsoft 365)
- Chemical Facility Anti-Terrorism Standards (CFATS) (Microsoft 365)
- Children's Online Privacy Protection Rule (COPPA) (Microsoft 365)
- China - Personal Information Security Specification (Microsoft 365)
- CIS Implementation Group 1, Group 2, Group 3
- Cloud Security Alliance (CSA) Cloud Controls Matrix (CCM)

- CMMC Level 1, Level 2, Level 3, Level 4, Level 5 (Microsoft 365)
- COBIT 5 (Microsoft 365)
- Colombia - Decree No. 1377/2013 (used to be the Colombia Law 1581/2012)
- Colombia - External Circular Letter 007 of 2018 (Microsoft 365)
- Colombia - Law 1266/2008- Habeas Data Act (Microsoft 365)
- Colombia - Law 1581/2012 (Microsoft 365)
- Commission Statement and Guidance on Public Company Cybersecurity Disclosures - US
- Computer Fraud and Abuse Act (CFAA) (Microsoft 365)
- Connecticut General Statutes - General Provisions for state contractors who receive confidential information (Microsoft 365)
- Connecticut State Law - Breach of security re computerized data containing personal information (Microsoft 365)
- Consumer Personal Information Security Breach Notification Act (Microsoft 365)
- Criminal Justice Information Services (CJIS) Security Policy (Microsoft 365)
- Croatia - Personal Data Protection Act (Microsoft 365)
- Cybersecurity Law of the People's Republic of China (Microsoft 365)
- Cyprus The Processing of Personal Data Law (Microsoft 365)
- Czech - Act No. 110/2019 Coll. on Personal Data Processing - 2019 (Microsoft 365)
- Czech - On Cyber Security and Change of Related Acts (Act on Cyber Security) - Act No. 181 (Microsoft 365)
- Delaware Computer Security Breaches- Commerce and Trade Subtitle II - 12B-100 to 12B-104
- Denmark - The Data Protection Act
- Denmark - Executive Order on Information and Consent Required in Case of Storing and Accessing Information in End-User Terminal Equipment
- DFARS (Microsoft 365)
- Directive 2013/40/EU Of The European Parliament And Of The Council (Microsoft 365)
- Dubai - Health Data Protection Regulation (Microsoft 365)
- Dubai Consumer Protection Regulations (Telecommunications Regulatory Authority)(Microsoft 365)
- Dubai ISR (Microsoft 365)
- Electronic Code of Federal Regulations - Part 748.0 and Appendix A (Microsoft 365)
- ENISA Information Assurance Framework
- Estonia - Personal Data Protection Act (Microsoft 365)
- Estonia - The system of security measures for information systems (Microsoft 365)
- EU - ePrivacy Directive 2002 58 EC (Microsoft 365)
- EU - EudraLex Volume 4 — GMP Guidelines, Annex 11
- EU Directive 2006/24/EC
- FDIC Privacy Rules (Microsoft 365)
- Federal Financial Institutions Examination Council (FFIEC) Information Security Booklet (Microsoft 365, Intune)
- FedRamp High Security Controls (Office 365)
- FedRamp High Security Controls_NIST 800-53 (Microsoft 365)
- FedRAMP Moderate
- Finland - Data Protection Act
- Finnish Criteria for Assessment of Information Security of Cloud Services (Microsoft 365)
- FINRA Cybersecurity Checklist
- France - Act 78-17 Of 6 January 1978 On Information Technology, Data Files and Civil Liberties (Microsoft 365)
- Freedom of Information Act (FOIA) (Microsoft 365)

- FTC Privacy of Consumer Financial Information (Microsoft 365)
- Ghana Data Protection Act
- Generally Accepted Recordkeeping Principles (Microsoft 365)
- Germany - Cloud Computing Compliance Controls Catalog (C5) (Microsoft 365)
- Germany - Federal Data Protection Act (Microsoft 365)
- Gramm-Leach-Bliley Act, Title V, Subtitle A, Financial Privacy (Microsoft 365)
- Greece - Law 2472/1997 on the Protection of individuals with regard to the processing of personal data (Microsoft 365)
- Hawaii - Security Breach of Personal Information Chapter 487N
- HIPAA/HITECH (Microsoft 365, Intune)
- HITRUST (Microsoft 365)
- Hong Kong - Personal Data (Privacy) Ordinance (Microsoft 365)
- India - IT Act of 2000 (Microsoft 365)
- India Information Technology (Reasonable Security Practices and Procedures and Sensitive Personal Data or Information) Rules
- Indonesia - Law 11/2008 (Microsoft 365)
- Indonesia - Peraturan Pemerintah No.82 Tahun 2012 - Government Regulation - Data Protection Regulation (Microsoft 365)
- IRAP v3
- IRS-P1075 (Microsoft 365)
- IRS - Revenue Procedure 98-25 Automated Records
- ISO 15489 (Microsoft 365)
- ISO 22301:2019 (Microsoft 365)
- ISO 27005:2018 (Microsoft 365)
- ISO 27017:2015 (Microsoft 365)
- ISO 27799 Health informatics — Information security management in health using ISO/IEC 27002 (Microsoft 365)
- ISO 31000:2018 (Microsoft 365)
- ISO 80001-1 Application of risk management for IT-networks incorporating medical devices (Microsoft 365)
- ISO/IEC 27018:2014
- ISO/IEC 27701:2019 (Microsoft 365)
- Israel - Privacy Protection (Transfer of Data to Databases Abroad) Regulations (Microsoft 365)
- ITU X.1052 Information Security Management Framework (Microsoft 365)
- Japan - Act on Prohibition of Unauthorized Computer Access (Microsoft 365)
- Japan - Common Model of Information Security Measures for Government Agencies and Related Agencies (Microsoft 365)
- Japan - Common Standards for Information Security Measures for Government Agencies and Related Agencies
- Japan Privacy Mark
- Japanese Act on the Protection of Personal Information (Law No. 57 of 2003) (Microsoft 365)
- Joint Commission AHO Information Management Standard (Microsoft 365)
- Kenya Data Protection Act (Microsoft 365)
- Korea - The Act on Promotion of Information and Communications Network Utilization and Data - Protection (Microsoft 365)
- Korea - Use and Protection of Credit Information Act (Special Law) (Microsoft 365)
- Korea Personal Information Protection Act (Microsoft 365)
- Kuwait - CSF (Microsoft 365)

- Luxembourg Act (Microsoft 365)
- Maine - Act to Protect the Privacy of Online Consumer Information
- Maine - Notice of Risk to Personal Data (Microsoft 365)
- Malaysia - Personal Data Protection Act (PDPA) (Microsoft 365)
- Malaysia Risk Management in Technology (RMiT) (Microsoft 365)
- Massachusetts - 201 CMR 17.00: Standards For The Protection Of Personal Information Of - Residents Of The Commonwealth (Microsoft 365)
- Mauritius Data Protection Act 2004 (Microsoft 365)
- Mexico - Federal Consumer Protection Law (Microsoft 365)
- Mexico Federal Data Protection Law (Microsoft 365)
- Minimum Acceptable Risk Standards for Exchanges (MARS-E) 2.0 (Microsoft 365)
- Motion Picture Association (MPA) Content Security Best Practices (Microsoft 365)
- Myanmar - Law Protecting the Privacy and Security of Citizens
- NAIC - Standards for Safeguarding Customer Information Model Regulation MDL-673 (Microsoft 365)
- Nepal - Right to Information Act
- NERC CIP (Microsoft 365)
- Netherlands - Personal Data Protection Act / 1999 (Microsoft 365)
- Nevada Chapter 603A - Security and Privacy of Personal Information (Microsoft 365)
- New York Privacy Act - DRAFT (Microsoft 365)
- New Zealand Health Data Retention Policy (Office 365)
- New Zealand Health Information Privacy Code 1994 (Microsoft 365)
- New Zealand Health Information Security Framework (HISF) -2015 (Microsoft 365)
- New Zealand Privacy Act 2020 (Microsoft 365)
- New Zealand Public Records Act (Microsoft 365)
- New Zealand Telecommunications Information Privacy Code 2003
- Nigeria Data Protection Regulation (Microsoft 365)
- NIST 800-37 (Microsoft 365)
- NIST 800-53
- NIST 800-63 Digital Identity Guidelines (Microsoft 365)
- NIST 800-171 (Microsoft 365)
- NIST CSF (Microsoft 365)
- NIST Privacy Framework
- NIST Special Publication 800-128 (Microsoft 365)
- NIST Special Publication 1800-1 Securing Electronic Health Records on Mobile Devices (Microsoft 365)
- NIST Special Publication 1800-5 IT Asset Management
- Norway - Personal Data Act (Microsoft 365)
- NYDFS (Microsoft 365)
- Oman - Electronic Transactions Law (Microsoft 365)
- OWASP ProActive Controls for Developers 2018 v3.0 (Microsoft 365)
- Pakistan Electronic Data Protection Act 2005 -Draft (Microsoft 365)
- PCI DSS v3.2.1 (Microsoft 365)
- Peruvian Legislation Law 29733 Law of Data Privacy Protection
- Philippines BSP Information Security Management Guidelines (Microsoft 365)
- Philippines Data Privacy Act of 2012 (Microsoft 365)
- Privacy of Consumer Financial and Health Information Regulation, NAIC MDL-672, Q2 2017 (Microsoft 365)
- Puerto Rico - Citizen Information on Data Banks Security Act (Microsoft 365)
- Qatar Cloud Security Policy

- RBNZ BS11 Outsourcing Policy (Microsoft 365)
- Republic of Moldova Law on Personal Data Protection (Microsoft 365)
- Reserve Bank of India Cyber Security Framework (Microsoft 365)
- Romania - Data Protection Law 190/2018 (Microsoft 365)
- Russia - Federal Law 149-FZ On Information, Information Technology and Information Security
- Russian Federation Federal Law Regarding Personal Data (Microsoft 365)
- Saint Lucia Data Protection Act (Microsoft 365)
- SEC 17-4(a) (Microsoft 365)
- SIG (Microsoft 365)
- Singapore - Banking Act (Cap.19)
- Singapore - Cybersecurity 2018 (Microsoft 365)
- Singapore - IMDA IoT Cyber Security Guide (Microsoft 365)
- Singapore - Monetary Authority of Singapore Technology Risk Management Framework (Microsoft 365)
- Singapore - Multi-Tier Cloud Security (MTCS) Standard (Microsoft 365)
- Singapore - Outsourced Service Provider Audit Report (OSPAR) (Microsoft 365)
- Singapore - Personal Data Protection Act / 2012 (Microsoft 365)
- Singapore Spam Control Act (Microsoft 365)
- SOC 1 (Microsoft 365)
- SOC 2 (Microsoft 365)
- South Africa Consumer Protection ACT 68 2008 (Microsoft 365)
- South Africa Consumer Protection ACT 68 2008 (Microsoft 365)
- South Africa Electronic Communications and Transactions Act, 2002 (Microsoft 365)
- South African POPIA (Microsoft 365)
- Spain - Nation Security Framework (Microsoft 365)
- SWIFT Customer Security Controls (Microsoft 365)
- Switzerland - Federal Act on Data Protection (FADP) (Microsoft 365)
- Taiwan - Implementation Rules for the Internal Audit and Internal Control System of Electronic Payment Institutions - 2015 (Microsoft 365)
- Taiwan - Regulations Governing Approval and Administration of Financial Information Service Enterprises Engaging in Interbank Funds Transfer and Settlement (Microsoft 365)
- Taiwan - Regulations Governing the Standards for Information System and Security Management of Electronic Payment Institutions (Microsoft 365)
- Taiwan- Implementation Rules of Internal Audit and Internal Control System of Financial Holding Companies and Banking Industries (Microsoft 365)
- Taiwan Personal Data Protection Act (PDPA) (Microsoft 365)
- Texas - Identity Theft Enforcement and Protection Act (Microsoft 365)
- Thailand PDPA (Microsoft 365)
- Trade Secrets Act of The Republic of China (Microsoft 365)
- Trinidad and Tobago Data Protection (Act 13 of 2011) (Microsoft 365)
- Trusted Information Security Assessment Exchange
- Turkey - KVKK Protection of Personal Data 6698 (Microsoft 365)
- UAE - Federal Law No 2 of 2019 On the Use of the Information and Communication Technology (ICT) in Health Fields
- UK - The Offshore Petroleum Activities Regulations / 2011 (Microsoft 365)
- UK Cyber Essentials (Microsoft 365)
- UK- Cyber Security for Defense Suppliers Standard (Microsoft 365)
- UK Privacy and Electronic Communications (Microsoft 365)

- Ukraine - Protection of Personal Data Law (Microsoft 365)
- US DoE 10 CFR Part 810 (Microsoft 365)
- US - Federal Information Security Modernization Act of 2014 (FISMA) (Microsoft 365)
- US FERPA (Microsoft 365)
- US-Cloud Act (Microsoft 365)
- Utah Consumer Credit Protection Act (Microsoft 365)
- Uzbekistan Law on Personal Data
- Victorian Protective Data Security Standards V2.0 (VPDSS 2.0) (Microsoft 365)
- Vietnam - Consumer Rights Protection Law (Microsoft 365)
- Vietnam - Law of Cybersecurity (Microsoft 365)
- Vietnam - Law of Network Information Security
- Vietnam - Law on Information Technology (Microsoft 365)
- Yemen - Yemen Law of the Right of Access to Information (Microsoft 365)

Assign and complete improvement actions in Compliance Manager

2/18/2021 • 8 minutes to read • Edit Online

In this article: This article explains how to **manage your compliance workflow** with improvement actions. Learn how to **assign improvement actions** for implementation and testing, **manage updates**, and export **reports**.

Manage compliance workflows with improvement actions

Improvement actions centralize your compliance activities. Each improvement action gives detailed implementation guidance to help you align with data protection regulations and standards. Actions can be assigned to users in your organization to perform implementation and testing work. You can also store documentation, notes, and record status updates within the action.

All of your improvement actions are listed on the improvement actions page. Learn more about viewing your improvement actions.

Improvement actions details page

Each improvement action has a details page showing its current status, the related standards and regulatory requirements, and recommended implementation guidance. Technical actions include a **Launch now** link that takes you to the appropriate solution for implementation. You can attach implementation and testing documentation directly into an improvement action's details page.

To view an improvement action's details page:

1. Go to your improvement actions page.
2. Select the row of your intended improvement action, which opens its details page.

You can easily view the next or previous improvement action in the list by selecting the up or down arrow in the upper-right corner of the screen. If you filtered your list on the improvement actions page, moving up or down takes you to the next item within that filtered list.

Assign improvement actions

To begin implementation work on an improvement action, you can do the work yourself or assign it to another user. The assigned person could be:

- A business policy owner
- An IT implementer
- Another employee with responsibility to perform the task

Once you identify the appropriate assignee, be sure they hold a sufficient Compliance Manager role to perform the work. Then follow the steps below to assign the improvement action:

1. From the improvement actions details page, select **Edit status** near the upper-left section of the screen.
2. In the edit status flyout pane, select the **Assigned to** box to show a **Suggested people** list of users. You can select the user from the list, or type the email address of the person you want to assign it to.
3. Select **Save and close**. The assigned user will receive an email explaining that the improvement action

has been assigned to them, with a direct link to the improvement action. (Note: US Government Community (GCC) High customers won't receive an email when actions are assigned to them.)

The assigned user can then perform the recommended actions.

Assign multiple improvement actions to a single user

You can assign multiple improvement actions to one user by following these steps:

1. Go to your Improvement actions page.
2. Select the area to the left of the improvement action's name. A round check icon will appear indicating you've selected that action. Check all the actions you want to assign.
3. Select the **Assign to user** link at the top of the improvement actions table.
4. A pop-up window appears. In the **Assign to** field, start typing the name of the person you want to assign the actions to. You can also select from the list of suggested people.
5. After you populate the **Assign to** field with the assignee's name, select **Assign**.
6. You'll then see your Improvement actions page with the new assignee listed for the actions you just assigned.

Perform work and store documentation

You can upload files and notes related to implementation and testing work directly to the **Notes and documentation** section. This environment is a secure, centralized repository to help you demonstrate satisfaction of controls to meet compliance standards and regulations. Any user with read-only access can read content in this section. Only users with editing rights can upload and download files and enter or edit notes.

The **Notes and documentation** section contains fields for uploaded documents, implementation notes, test notes, and additional notes.

Uploaded documents

- Select **Manage documents** to upload any relevant files.
- When the manage documents flyout pane opens, select **Add document**, then select your file from your system. Accepted file types:
  - Documents (.doc, .xls, .ppt, .txt, .pdf)
  - Images (.jpg, .png)
  - Video (.mkv)
  - Compressed files (.zip, .rar)
- Once your file resolves in the pane select **Close**, which automatically saves the file attachment. You'll then see the file listed underneath **Uploaded documents**.
- To download or delete the document, select **Manage documents** from underneath the list of documents. On the flyout pane, select the document row to highlight it, then select **Download** or **Delete**.

Implementation notes, test notes, and additional notes

- To add notes in any of these three fields, select **Edit implementation notes** underneath any of these fields.
- When the flyout pane opens, enter notes in the text field, then select **Save and close**.
- To edit notes, select **Edit implementation notes**, make your edits, then select **Save and close**.

There's no character limit in the notes fields. We recommend keeping notes brief so that you can easily view and edit them from the improvement actions details page.

Change improvement action status

You can record the implementation status and date, and the test status and date, for each improvement action. The **implementation** and **test status** fields can be edited by any user with editing permissions, not just by the assigned person.

To edit an improvement action's status, select **Edit status** on the upper-left section of the details page. Below are the available fields and status options:

- **Implementation status**
  - **Not implemented** - action not yet implemented
  - **Implemented** - action implemented
  - **Alternative implementation** - select this option if you used other third-party tools or took other actions not included in Microsoft recommendations
  - **Planned** - action is planned for implementation
  - **Out of scope** – action isn't relevant to your organization and doesn't contribute to your score
- **Implementation date**: available to select when implementation status is "implemented" or "alternative implementation"
- **Test status**: available to select when implementation status is "implemented" or "alternative implementation":
  - **Not assessed** – action hasn't been tested
  - **Passed** - implementation has been verified by an assessor
  - **Failed low risk** - testing failed, low risk
  - **Failed medium risk** - testing failed, medium risk
  - **Failed high risk** – testing failed, high risk
  - **Out of scope** – the action is out of scope for the assessment and doesn't contribute to your score
- **Test date**: toggle through the calendar pop-up to select the date

Common actions synch across groups. When two different assessments in the same group share improvement actions that are managed by you, any updates you make to an action's implementation details or status will automatically synchronize to the same action in any other assessment in the group. This synchronization allows you to implement one improvement action and meet several requirements across multiple regulations.

Assign improvement action to assessor for completion

After you complete the work, conduct testing, and upload evidence, the next step is to assign the improvement action to an assessor for validation. The assessor validates the work and examines the documentation, and selects the appropriate test status.

If test status is set to "Passed": the action is complete and the points achieved shows the maximum points achieved. The points are then counted toward your overall compliance score.

If test status is set to "Failed": the action doesn't meet the requirements, and the assessor can assign it back to the appropriate user for additional work.

Accepting updates to improvement actions

When an update is available for an improvement action, you'll see a notification next to its name. You can either accept the update or defer it for a later time.

What causes an update

An update occurs when there are changes related to scoring, automation, or scope. Changes may involve new guidance for improvement actions based on regulatory changes, or could be because of product changes. Only the improvement actions managed by your organizations receive update notifications.

Where you'll see assessment update notifications

When an improvement action is updated, you'll see a **Pending update** label next to its name on the improvement actions page, and on the details page of its related assessments.

Go to the improvement action's details page, and select the **Review update** button in the top banner to review

details about the changes and accept or defer the update.

Review update to accept or defer

After selecting **Review update** from the improvement action details page, a flyout pane appears on the right side of your screen. The flyout pane provides key details about the update, such as the assessments impacted and changes in score and scope.

Select **Accept update** to accept all the changes to the improvement action. **Accepted changes are permanent**.

> **NOTE**
> When you accept an update to an action, you're also accepting updates to any other versions or instances of this action. Updates will propagate tenant-wide for technical actions, and will propagate group-wide for non-technical actions.

If you select **Cancel**, the update won't be applied to the improvement action. However, you'll continue to see the **Pending update** notification until you accept the update.

Why we recommend accepting updates

Accepting updates helps ensure you have the most updated guidance on using solutions and taking appropriate improvement actions to help you meet the requirements of the certification at hand.

Why you might want to defer an update

If you're in the middle of completing an assessment that includes the improvement action, you may want to ensure you've finished work on it before you accept the update. You can defer the update for a later time by selecting **Cancel** on the review update flyout pane.

Accept all updates at once

If you have multiple updates and want to accept them all at one time, select the **Accept all updates** link at the top of your improvement actions table. A flyout pane will appear which lists the number of actions to be updated. Select the **Accept updates** button to apply all updates.

Note that when you return to your improvement actions page, you may see a message across the top of the page asking you to refresh the page for the updates to be completed.

Export a report

Select **Export** in the upper-left corner of your screen to download an Excel worksheet containing all your improvement actions and the filter categories shown on the improvement actions page.

Compliance score calculation

11/2/2020 • 6 minutes to read • Edit Online

In this article: Learn how Compliance Manager calculates a compliance score for your organization. This article explains how to **interpret your score**, what the **Data Protection Baseline assessment** includes, **continuous monitoring**, and **how different types of actions are managed and scored**.

> **IMPORTANT**
>
> Recommendations from Compliance Manager should not be interpreted as a guarantee of compliance. It is up to you to evaluate and validate the effectiveness of customer controls per your regulatory environment. These services are subject to the terms and conditions in the Online Services Terms. See also Microsoft 365 licensing guidance for security and compliance.

How to read your compliance score

The Compliance Manager dashboard displays your overall compliance score. This score measures your progress in completing recommended improvement actions within controls. Your score can help you understand your current compliance posture. It can also help you prioritize actions based on their potential to reduce risk.

A score value is assigned at three levels:

1. **Improvement action score**: each action has a different impact on your score depending on the potential risk involved
2. **Control score**: this score is the sum of points earned by completing improvement actions within the control. This sum is applied in its entirety to your overall compliance score when the control meets both of the following conditions:
   - **Implementation Status** equals **Implemented** or **Alternative Implementation**, and
   - **Test Result** equals **Passed**.
3. **Assessment score**: this score is the sum of your control scores. It is calculated using action scores. Each Microsoft action and each improvement action managed by your organization is counted once, regardless of how often it is referenced in a control.

The overall compliance score is calculated using action scores, where each Microsoft action is counted once, each technical action you manage is counted once, and each non-technical action you manage is counted once per group. This logic is designed to provide the most accurate accounting of how actions are implemented and tested in your organization. You may notice that this can cause your overall compliance score to differ from the average of your assessment scores. Read more below about how actions are scored.

Initial score based on Microsoft 365 data protection baseline

Compliance Manager gives you an initial score based on the Microsoft 365 data protection baseline. This baseline is a set of controls that includes key regulations and standards for data protection and general data governance. This baseline draws elements primarily from NIST CSF (National Institute of Standards and Technology Cybersecurity Framework) and ISO (International Organization for Standardization), as well as from FedRAMP (Federal Risk and Authorization Management Program) and GDPR (General Data Protection Regulation of the European Union).

Your initial score is calculated according to the default Data Protection Baseline assessment provided to all

organizations. Upon your first visit, Compliance Manager is already collecting signals from your Microsoft 365 solutions. You'll see at a glance how your organization is performing relative to key data protection standards and regulations, and see suggested improvement actions to take.

Because every organization has specific needs, Compliance Manager relies on you to set up and manage assessments to help minimize and mitigate risk as comprehensively as possible.

How Compliance Manager continuously assesses controls

Compliance Manager automatically scans through your Microsoft 365 environment and detects your system settings, continuously and automatically updating your technical action status. Microsoft Secure Score is the underlying engine that performs the monitoring.

Your action status is updated on your dashboard every 24 hours. Once you follow a recommendation to implement a control, you'll typically see the control status updated the next day.

For example, if you turn on multi-factor authentication (MFA) in the Azure AD portal, Compliance Manager detects the setting and reflects it in the control access solution details. Conversely, if you didn't turn on MFA, Compliance Manager flags that as a recommended action for you to take.

Learn more about Secure Score and how it works.

Action types and points

Compliance Manager tracks two types of actions:

1. **Your improvement actions**: actions that your organization manages.
2. **Microsoft actions**: actions that Microsoft manages.

Both types of actions have points that count toward your overall score when completed.

Technical and non-technical actions

Actions are grouped by whether they are technical or non-technical in nature. The scoring impact of each action differs by type.

- **Technical actions** are implemented by interacting with the technology of a solution (for example, changing a configuration). The points for technical actions are granted once per action, regardless of how many groups it belongs to.
- **Non-technical actions** are managed by your organization and implemented in ways other than working with the technology of a solution. There are two types of non-technical actions: **documentation** and **operational**. The points for these actions are applied to your compliance score at a group level. This means that if an action exists in multiple groups, you will receive the action's point value each time you implement it within a group.

Example of how technical and non-technical actions are scored:

Let's say you have a technical action worth 3 points that exists in 5 groups, and you have a non-technical action worth 3 points that exists in the same 5 groups.

If you successfully implement the technical action, the total number of points you receive is 3. This is because you only need to implement the action once for your tenant. The implementation and test status for the technical action will show the same in all instances of that action, in every group it belongs to.

If you successfully implement the non-technical action in each of the 5 groups, the total number of points you receive is 15. This is because you need to implement the action in each group. The implementation and test status for the non-technical action will differ across groups because the action is implemented separately within each of its groups.

This scoring logic is designed to provide the most accurate accounting of how actions are implemented and tested in your organization.

How score values are determined

Actions are assigned a score value based on whether they're mandatory or discretionary, and whether they're preventative, detective, or corrective.

Mandatory and discretionary actions

- **Mandatory actions** can't be bypassed, either intentionally or accidentally. An example of a mandatory action is a centrally managed password policy that sets requirements for password length, complexity, and expiration. Users must follow these requirements to access the system.
- **Discretionary actions** rely upon users to understand and adhere to a policy. For example, a policy requiring users to lock their computer when they leave it is a discretionary action because it relies on the user.

Preventative, detective, and corrective actions

- **Preventative actions** address specific risks. For example, protecting information at rest using encryption is a preventative action against attacks and breaches. Separation of duties is a preventative action to manage conflict of interest and guard against fraud.
- **Detective actions** actively monitor systems to identify irregular conditions or behaviors that represent risk, or that can be used to detect intrusions or breaches. Examples include system access auditing and privileged administrative actions. Regulatory compliance audits are a type of detective action used to find process issues.
- **Corrective actions** try to keep the adverse effects of a security incident to a minimum, take corrective action to reduce the immediate effect, and reverse the damage if possible. Privacy incident response is a corrective action to limit damage and restore systems to an operational state after a breach.

Each action has an assigned value in Compliance Manager based on the risk it represents:

TYPE	ASSIGNED SCORE
Preventative mandatory	27
Preventative discretionary	9
Detective mandatory	3
Detective discretionary	1
Corrective mandatory	3
Corrective discretionary	1

How is your compliance score calculated?

	Preventative	Detective	Corrective
Mandatory	+27 points	+3 points	+3 points
Discretionary	+9 points	+1 points	+1 points

Microsoft Compliance Configuration Analyzer for Compliance Manager (preview)

2/18/2021 • 6 minutes to read • Edit Online

In this article: Learn how to install and run the Microsoft Compliance Configure Analyzer tool to get quickly started with Microsoft Compliance Manger.

Microsoft Compliance Configuration Analyzer (MCCA) (preview) overview

The Microsoft Compliance Configuration Analyzer (MCCA) is a preview tool that can help you get started with Microsoft Compliance Manager. MCCA is a PowerShell-based utility that will fetch your organization's current configurations and validate them against Microsoft 365 recommended best practices. These best practices are based on a set of controls that include key regulations and standards for data protection and data governance.

MCCA can help you quickly see which improvement actions in Compliance Manger apply to your current Microsoft 365 environment. Each action identified by MCCA will give you recommendations for implementation, with direct links to Compliance Manager and the applicable solution to start taking corrective action.

An additional resource for understanding MCCA is by visiting the README instructions on GitHub. This page provides detailed information about prerequisites and gives full installation instructions. You don't need a GitHub account to access this page.

Availability: MCCA is available to all organizations with Office 365 and Microsoft 365 licenses and US Government Community (GCC) Moderate customers, with plans underway to expand service to to GCC High customers.

Install MCCA and run a report

You can install the MCCA tool using Windows PowerShell. Once you download and install the tool, you don't need to repeat those steps in order to run reports. Each time you open MCCA, it will ask you for your login credentials, and it will generate a new, updated report.

Step 1: Install Windows PowerShell

To begin, you'll need the Exchange Online PowerShell module (v2.0.3 or higher) that's available in the PowerShell gallery. Get installation instructions.

Step 2: Install MCCA

To install MCCA, start by using PowerShell in administrator mode. Follow the steps below:

1. Select the Windows **Start** button.
2. Type **PowerShell**, right-click on **Windows PowerShell**, then select **Run as administrator**.
3. At the command prompt, type:

```
Install-Module -Name MCCAPreview
```

Step 3: Run a report

After you install MCCA, you can run MCCA and generate a report. To run a report:

1. Open PowerShell
2. Run the cmdlet:

```
Get-MCCAReport
```

3. Once MCCA runs, it does an initial version check and ask for credentials. At the Input the user name prompt, sign in with your Microsoft 365 account email address (view the roles eligible to create reports). Then enter your password at the password prompt.

Your report will then take approximately 2-5 minutes to generate. When it's done, a browser window opens and displays your HTML report. Every time you run the tool, it will ask for your credentials and generate a new report. This report is stored locally in the following directory:

C:\Users<username>\AppData\Local\Microsoft\MCCA.

You can access previously generated reports from this directory.

Understanding your report

Your report reflects data based on the date and time at which it was generated. The top section provides details on when it was generated, your organization name, and tenant ID.

Geolocation-based reporting

The **Note** section shows that your report is customized based on the geographic location of your tenant. Recommendations listed in the tool will be specific to your country or region.

Your geolocation selection is used to assess sensitive information types (SITs) which are relevant to that geolocation and generate a report that aligns to your country or region. Choose geolocations based on data you have in your tenant.

To change your report's location information, you need provide a geolocation (-Geo) input parameter. You can choose either one or multiple geolocations applicable for your tenant.

Follow these instructions to run a report based on a specific location:

1. Open PowerShell
2. To specify a certain region, you'll run a cmdlet using the numbers from the table below that correspond to the country or region. Enter multiple numbers by separating them with a comma. For example, the cmdlet below will run a customized report for Asia-Pacific and Japan:

```
Get-MCCAReport -Geo @(1,7)
```

INPUT	COUNTRY OR REGION
1	Asia-Pacific
2	Australia
3	Canada
4	Europe (excluding France) / Middle East / Africa
5	France

INPUT	COUNTRY OR REGION
6	India
7	Japan
8	Korea
9	North America (excluding Canada)
10	South America
11	South Africa
12	Switzerland
13	United Arab Emirates
14	United Kingdom

> **NOTE**
> The report will always include MCCA supported international sensitive information types such as SWIFT code, credit card number, etc.

Role-based reporting

Your report will also be customized based on your role.

The table below shows which roles have access to which sections of the report. Other roles within your organization (not listed in the table below) may not be able to run the tool, or they may run the tool and have limited access to information in the final report.

User Role	MIP		MIG		Insider Risk		Discovery & Response	
	DLP	IP	IG	RM	IRM	CC	Audit	eDiscovery
Azure Information Protection admin	No	No[1]	No	No	No	No	No[4]	No
Compliance admin	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Compliance Data Admin	Yes	Yes[2]	Yes	Yes	Yes	Yes[3]	Yes[5]	No
Customer Lockbox access approver	No	No	No	No	No	No	No	No
Exchange Admin	No	No[1]	No	No	No	No	No[4]	No
Global reader	Yes	Yes	Yes	Yes	No	No	Yes	No
Helpdesk admin	No	No[1]	No	No	No	No	No[4]	No
Non-Admin User	No	No	No	No	No	No	No	No
Reports reader	No	No	No	No	No	No	No	No
Security admin	Yes	Yes[2]	No	No	No	No	Yes[5]	No
Security operator	Yes	No	No	No	No	No	Yes[5]	No
Security reader	Yes	Yes[2]	No	No	No	No	Yes[5]	No
Service support admin	No	No	No	No	No	No	No	No
SharePoint admin	No	No	No	No	No	No	No	No

Exceptions:

1. User won't be able to generate report for IP apart from "Use IRM for Exchange Online" section.
2. User will be able to generate report for IP apart from "Use IRM for Exchange Online" section.
3. User will be able to generate report for IP apart from "Enable Communication Compliance in O365" section.
4. User won't be able to generate report for IP apart from "Enable Auditing in Office 365" section.
5. User will be able generate report for IP apart from "Enable Auditing in Office 365" section.

Solutions Summary section

The **Solutions Summary** section of the report gives an overview of improvement actions that your organization can take in Compliance Manager to help improve your compliance posture.

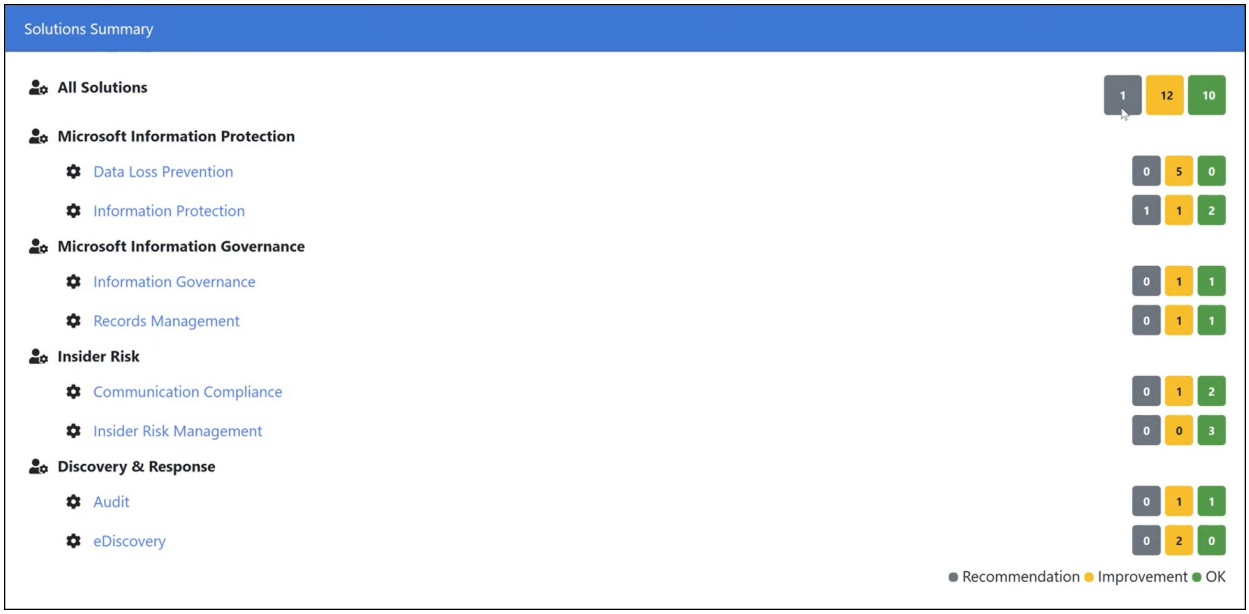

MCCA evaluates your current configurations against the recommended improvement actions in Compliance Manager. Any improvement action identified by the MCCA tool as needing attention will be listed in this section.

Next to each Microsoft solution are color-coded boxes indicating the number of items that correspond to improvement actions in Compliance Manager. The actions are broken down into three status states:

- **OK**: the actions that meet recommended conditions and need no attention at this time
- **Improvement**: actions that need attention
- **Recommendation**: actions that don't need attention, but for which we recommend best practices

Select a box to view improvements and recommendations.

Items with the Improvement status

Select the dropdown next to the **Improvement** label to the right of the improvement action. You'll see a quick summary and details about your current settings and the recommended improvement actions. The summary includes direct links into Compliance Manager, the applicable solution in the Microsoft 365 compliance center, and relevant documentation.

Clicking on the Compliance Manager link takes you to a filtered view of all the improvement actions within that solution that you have not yet implemented. From there, you can see the number of points you can achieve to increase your compliance score, and the assessments they apply to, and the applicable regulations and certifications.

For DLP, there's a **Remediation Script** button that gives you a pre-generated PowerShell script based on what's recommended. You can copy and paste it directly in your PowerShell console. It will create a DLP policy in test mode

Items with Recommendation status

Select the dropdown next to the **Recommendation** label to the right of the improvement action. You'll see a summary of your organization's current Microsoft 365 environment related to the improvement action, along with recommended best practices.

Resources

For more detailed information on installing, setting up, and using MCCA, see the README instructions on GitHub (no GitHub account required).

For more information on Windows PowerShell, start at How to use the PowerShell documentation. See also Starting Windows PowerShell.

Compliance Manager frequently asked questions

2/18/2021 • 4 minutes to read • Edit Online

Is Compliance Manager and Compliance Score the same thing, or are they different?

There is now just one solution: Compliance Manager. This section walks you through the transition, starting with a basic overview below. You may also find it helpful to jump directly to one of the following sections:

- Your organization primarily used Compliance Manager (either the classic or public preview versions), located in the Microsoft Service Trust Portal
- Your organization primarily used Compliance Score (public preview), located in the Microsoft 365 compliance center
- Your organization is new to Compliance Manager

The basics

Microsoft Compliance Manager began as a compliance management solution inside the Microsoft Service Trust Portal. As compliance solutions came into in the Microsoft 365 compliance center, we developed a new experience with a more user-friendly design for this location. Compliance Score public preview was released in the Microsoft 365 compliance center in November 2019. Compliance Score shared the same backend as Compliance Manager, allowing customers to work in both places. Since November 2019, we've released several updates as we built new functionality and responded to customer feedback.

The general availability of Compliance Manager in the Microsoft 365 compliance center in September 2020 completes this evolution. Compliance Manager is the unified, end-to-end compliance solution. Your compliance score remains a key component of Compliance Manager.

Read this blog post to learn more about what's new with the GA release of Compliance Manager.

Your organization regularly used Compliance Manager in the Service Trust Portal

If you used Compliance Manager in the Service Trust Portal, all of your organization's data now exists in Compliance Manager in the Microsoft 365 compliance center at https://compliance.microsoft.com/compliancemanager. There's nothing you need to do to resume your Compliance Manager work in its new location, other than to update any bookmarks you have to its previous location. All of your assessments and other data have been brought over for you.

Note that Compliance Manager (preview) is no longer accessible in the Service Trust Portal, and all links to it will redirect you to its new location in the Microsoft 365 compliance center. Compliance Manager (classic) remains in the Service Trust Portal, though its use is discouraged.

Everything you used to do in previous versions of Compliance Manager, such as completing actions (now called "improvement actions") and creating assessments, can be done in the new Compliance Manager. We've added over 150 new assessment templates and improved the template creation process. We'll add more enhancements in future releases.

Below are some helpful resources:

- Get familiar with your new Compliance Manager experience
- Find permissions and other setup information for Compliance Manager in its new home
- Learn more about the Microsoft 365 compliance center

Your organization used Compliance Score (public preview) in the Microsoft 365 compliance center

If you used Compliance Score in public preview, you'll notice Compliance Manager looks largely the same, with your score featured prominently on your dashboard. With the GA release, you no longer need to leave the Microsoft 365 compliance center in order to perform certain assessment management functions, such as creating and modifying templates for assessments. All functionality now resides in one place. Any data you had in the preview version of Compliance Score remains in the GA version of Compliance Manager.

Note that if you filtered your Compliance Score dashboard view, those filters were reset when we deployed the new Compliance Manager in September. You will need to reapply any filters you had.

Compliance Manager also has new licensing terms. See the question below on licensing.

You're new to Compliance Manager

Compliance Manager is an end-to-end solution in the Microsoft 365 compliance center for managing and tracking compliance activities. It's a great place to begin your compliance journey because it gives you an initial assessment of your compliance posture the first time you visit. Below are good places to start learning more:

- Get an overview of Compliance Manager
- Use our quickstart guide to help ramp up in stages
- Learn more about the Microsoft 365 compliance center

Are there licensing requirements for using Compliance Manager?

Yes. The GA release of Compliance Manager contains new licensing terms. All organizations with Office 365 and Microsoft 365 licenses, and US Government Community (GCC) Moderate and GCC High customers, have access to Compliance Manager. However, the assessments available to your organization and how you manage assessment templates depends on your licensing agreement. Visit the Microsoft 365 licensing guidance for security and compliance for details.

If I have a high score, does it mean I'm fully compliant?

No. Your compliance score measures your progress in completing recommended actions that help reduce risks around data protection and regulatory standards. It does not express an absolute measure of organizational compliance with regard to a particular standard or regulation. Compliance Manager, and your compliance score, should not be interpreted as a guarantee in any way.

Can I use Compliance Manager for non-Microsoft products?

While Compliance Manager provides continuous monitoring and recommended actions only for Microsoft cloud services, you can add custom assessments in Compliance Manager for your third-party services. In this way, you can use Microsoft Compliance Manager as a SaaS compliance management tool to help you manage all the controls across your digital assets.

What's happening to Compliance Manager (classic) in the Service Trust Portal?

The classic version of Compliance Manager, which resides in the Microsoft Service Trust Portal, will soon be retired. A Microsoft 365 Message Center notice will go out at least 60 days before the final retirement of Compliance Manager (classic). Customers who are managing their compliance activities in Compliance Manager (classic) will need to move their data, including assessments and controls, over to the new Compliance Manager solution in the Microsoft 365 compliance center. Customer data will not automatically transfer over to Compliance Manager in the Microsoft 365 compliance center when Compliance Manager (classic) is retired.

To learn how you can quickly set up the new Compliance Manager, read our Compliance Manager quickstart guide.

Microsoft Compliance Manager (classic)

2/18/2021 • 46 minutes to read • Edit Online

IMPORTANT

Compliance Manager (classic) will soon be removed from the Microsoft Service Trust Portal. We recommend that you transition to the new Compliance Manager in the Microsoft 365 compliance center, which provides an enhanced user experience and updated control mapping. Customers who have assessments in the classic version will need to create new assessments in the new Compliance Manager. Any existing data, including your assessments, controls, and other data, will not be transferred over to the new Compliance Manager. Learn more about the transition.

Compliance Manager isn't available in Office 365 operated by 21Vianet, Office 365 Germany, Office 365 U.S. Government Community High (GCC High), or Office 365 Department of Defense.

Compliance Manager, a workflow-based risk assessment tool in the Microsoft Service Trust Portal, enables you to track, assign, and verify your organization's regulatory compliance activities related to Microsoft Professional Services and Microsoft cloud services, such as Microsoft Office 365, Microsoft Dynamics 365, and Microsoft Azure.

Compliance Manager:

- Combines the detailed information provided by Microsoft to auditors and regulators as part of various third-party audits of Microsoft 's cloud services against various standards (for example, ISO 27001, ISO 27018, and NIST) and information that Microsoft compiles internally for its compliance with regulations (such as HIPAA and the EU General Data Protection Regulation, or GDPR) with your own self-assessment of your organization's compliance with these standards and regulations.
- Enables you to assign, track, and record compliance and assessment-related activities, which can help your organization cross team barriers to achieve your organization's compliance goals.
- Provides a Compliance Score to help you track your progress and prioritize the auditing controls that will help reduce your organization's exposure to risk.
- Provides a secure repository for you to upload and manage evidence and other artifacts related to your compliance activities.
- Produces richly detailed reports in Microsoft Excel that document the compliance activities performed by Microsoft and your organization, which can be provided to auditors, regulators, and other compliance stakeholders.

IMPORTANT

Compliance Manager is a dashboard that provides a summary of your data protection and compliance stature and recommendations to improve data protection and compliance. The Customer Actions provided in Compliance Manager are recommendations; it is up to each organization to evaluate the effectiveness of these recommendations in their respective regulatory environment prior to implementation. Recommendations found in Compliance Manager should not be interpreted as a guarantee of compliance.

What is Compliance Manager?

Compliance Manager is a workflow-based risk assessment tool designed to help you manage regulatory

compliance within the shared responsibility model of the cloud. Compliance Manager provides you with a dashboard view of standards and regulations and assessments that contain Microsoft's control implementation details and test results and customer control implementation guidance and tracking for your organization to enter. Compliance Manager provides certification assessment control definitions, guidance on implementation and testing of controls, risk-weighted scoring of controls, role-based access management, and an in-place control action assignment workflow to track control implementation, testing status and evidence management. Compliance Manager optimizes compliance workload by enabling customers to logically group assessments together and apply assessment control testing to identical or related controls, reducing the duplication of effort that might otherwise be required to satisfy identical control requirements across different certifications.

Assessments in Compliance Manager

The core component of Compliance Manager is called an *Assessment*. An Assessment is an assessment of a Microsoft service against a certification standard or data protection regulation (such as ISO 27001:2013, and the GDPR). Assessments help you to discern your organization's data protection and compliance posture against the selected industry standard for the selected Microsoft cloud service. Assessments are completed by the implementation of the controls that map to the certification standard being assessed.

The structure of an Assessment is based on the responsibility that is shared between Microsoft and your organization for assessing security and compliance risks in the cloud and for implementing the data protection safeguards specified by a compliance standard, a data protection standard, a regulation, or a law.

An Assessment is made of several components, which are:

- **In-Scope Services** - Each assessment applies to a specific set of Microsoft services, which are listed in the In-Scope Cloud Services section.
- **Microsoft-Managed Controls** - For each cloud service, Microsoft implements and manages a set of *controls* as part of Microsoft's compliance with various standards and regulations. These controls are organized into *control families* that align with the structure from the corresponding certification or regulation that the Assessment is aligned to. For each Microsoft-managed control, Compliance Manager provides details about how Microsoft implemented the control, along with how and when that implementation was tested and validated by an independent third-party auditor.

  Here's an example of three Microsoft-managed controls in the **Security** control family from an Assessment of Office 365 and the GDPR.

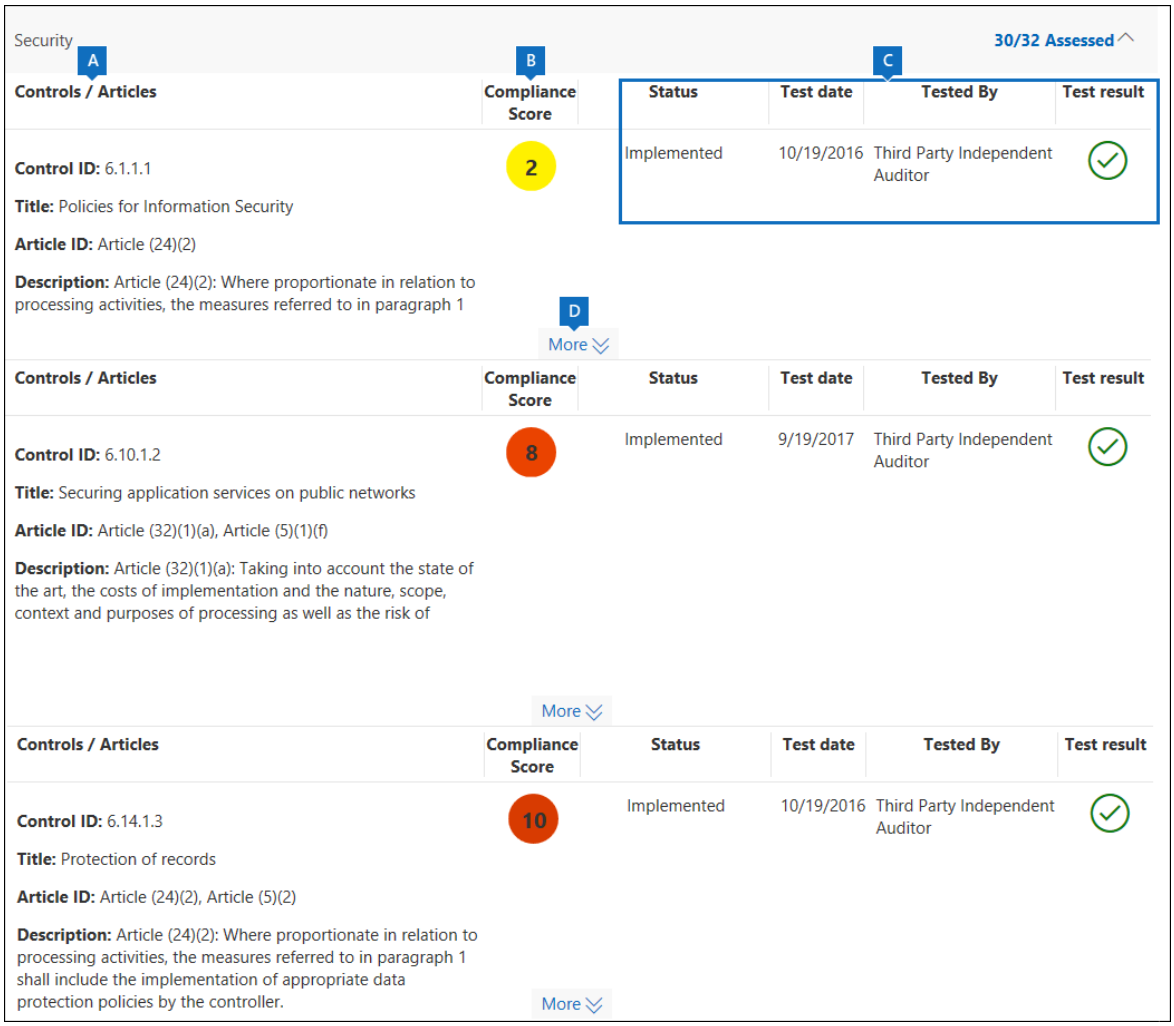

a. Specifies the following information from the certification or regulation that maps to the Microsoft-managed control.

- **Control ID** - The section or article number from the certification or regulation that the control maps to.
- **Title** - The title from the corresponding certification or regulation.
- **Article ID** - This field is included only for GDPR assessments, as it specifies the corresponding GDPR article number.
- **Description** - Text of the standard or regulation that maps to the selected Microsoft-managed control.

b. The Compliance Score for the control, which indicates the level of risk (due to non-compliance or control failure) associated with each Microsoft-managed control. See Understanding the Compliance Score for more information. Note that Compliance Scores are rated from 1 to 10 and are color-coded. Yellow indicates low risk controls, orange indicates medium-risk controls, and red indicated high-risk controls.

c. Information about the implementation status of a control, the date the control was tested, who performed the test, and the test result.

d. For each control, you can click **More** to see additional information, including details about Microsoft's implementation of the control and details about how the control was tested and validated by an independent third-party auditor.

- **Customer-Managed Controls** - This is the collection of controls that are managed by your organization. Your organization is responsible for implementing these controls as part of your

compliance process for a given standard or regulation. Customer-managed controls are also organized into control families for the corresponding certification or regulation. Use the customer-managed controls to implement the recommended actions suggested by Microsoft as part of your compliance activities. Your organization can use the prescriptive guidance and recommended Customer Actions in each customer-managed control to manage the implementation and assessment process for that control.

Customer-managed controls in Assessments also have built-in workflow management functionality that you can use to manage and track your organization's progress towards completing the Assessment. For example, a Compliance Officer in your organization can assign an Action Item to an IT admin who has the responsibility and necessary permissions to perform the actions that are recommended for the control. When that work is complete, the IT admin can upload evidence of their implementation tasks (for example, screenshots of configuration or policy settings) and then assign the Action Item back to the Compliance Officer to evaluate the collected evidence, test the implementation of the control, and record the implementation date and test results in Compliance Manager. For more information, see the Managing the assessment process section in the article.

Permissions and role-based access control

Compliance Manager uses a role-based access control permission model. Only users who are assigned a user role may access Compliance Manager, and the actions allowed by each user are restricted by role type.

Note that there is no longer a default **Guest access** role. Each user must be assigned a role in order to access and work within Compliance Manager.

The following table describes each Compliance Manager permission and what it allows the user do. The table also indicates the role that each permission is assigned to.

PERMISSION	COMPLIANCE MANAGER READER	COMPLIANCE MANAGER CONTRIBUTOR	COMPLIANCE MANAGER ASSESSOR	COMPLIANCE MANAGER ADMINISTRATOR	PORTAL ADMIN
Read data - Users can read but not edit data.	✓	✓	✓	✓	✓
Edit data - Users can edit all fields, except the Test Result and Test Date fields.		✓	✓	✓	✓
Edit test results - Users can edit the Test Result and Test Date fields.			✓	✓	✓
Manage assessments - Users can create, archive, and delete Assessments.				✓	✓

PERMISSION	COMPLIANCE MANAGER READER	COMPLIANCE MANAGER CONTRIBUTOR	COMPLIANCE MANAGER ASSESSOR	COMPLIANCE MANAGER ADMINISTRATOR	PORTAL ADMIN
Manage users - Users can add other users in their organization to the Reader, Contributor, Assessor, and Administrator roles. Only those users with the Global Administrator role in your organization can add or remove users from the Portal Admin role.					✓

Understanding the Compliance Score

On the Dashboard, Compliance Manager displays a total score for Office 365 assessments in the upper right-hand corner of the tile. This is the overall total Compliance Score for the Assessment, and is the accumulation of points received for each control assessment that has been marked as Implemented and Tested in the Assessment. When adding an Assessment, you will see that the Compliance Score is already on the way towards completion because the points for the Microsoft-managed controls that have been implemented by Microsoft and tested by independent third parties are already applied.

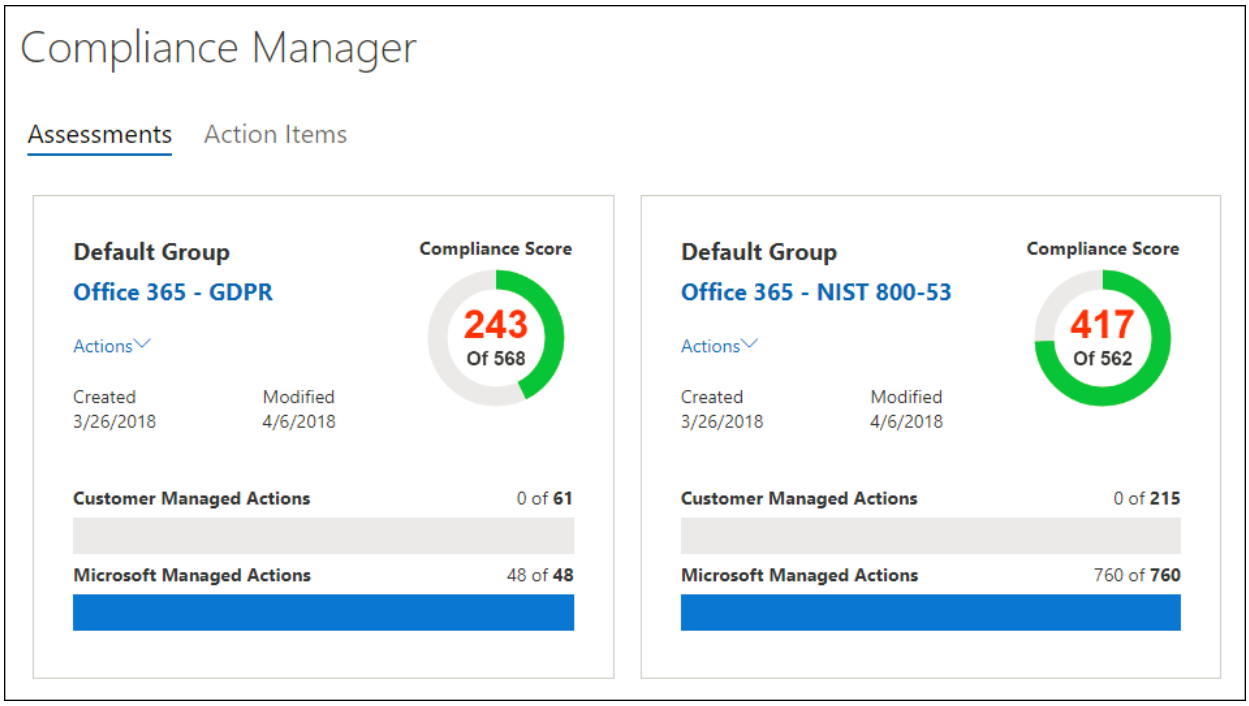

The remaining points come from the successful customer control assessment, from the implementation and testing of the customer-managed controls, each of which has a specific value that contributes to the overall compliance score.

Each Assessment displays a risk-based Compliance Score to help you assess the level of risk (due to non-compliance or control failure) associated with each control (including both Microsoft managed and customer-managed controls) in an Assessment. Each customer-managed control is assigned a possible number of points (called a *severity ranking) on a scale from 1 to 10, where more points are awarded for controls associated with a higher risk factor if the control fails, and fewer points are awarded for lower-risk controls.

For example, the User Access Management assessment control shown below has a very high severity risk ranking, and displays an assigned value of 10.

By comparison, the Information Backup assessment control shown below has a lower severity risk ranking, and displays an assigned value of 3.

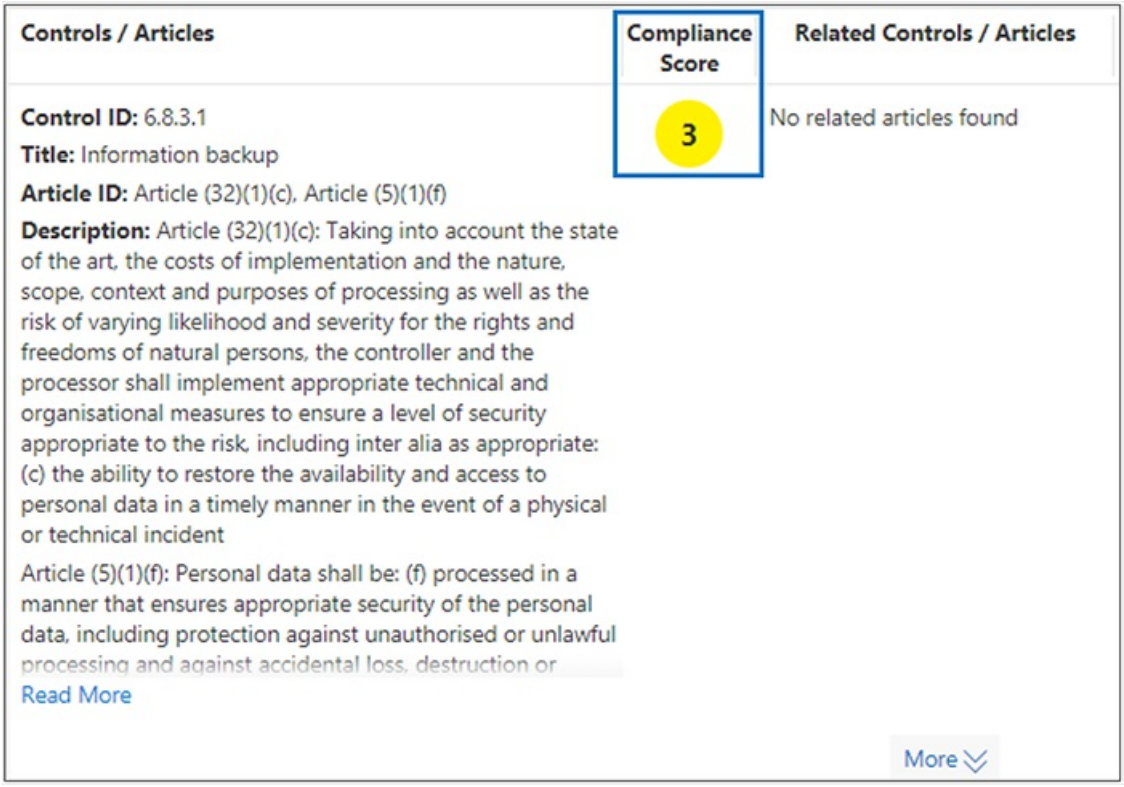

The Compliance Manager assigns a default severity ranking to each control. Risk rankings are calculated based on the following criteria:

- Whether a control prevents incidents from happening (highest ranking), detects incidents that have happened, or corrects the impact of an incident (lowest ranking). In terms of severity ranking, a mandatory control that prevents a threat is assigned the highest number of points; controls that are detective or corrective (regardless of whether they're mandatory or discretionary) are assigned the lowest number of points.

- Whether a control (after it's been implemented) is mandatory and therefore can't be by-passed by users (for example, users having to reset their password and meet password length and character requirements) or discretionary and can be by-passed by users (for example, business rules that require users to lock their screens when their computers are unattended).
- Controls related to risks to data confidentiality, integrity, and availability, whether these risks come from internal or external threats, and whether the threat is malicious or accidental. For example, controls that would help prevent an external attacker from breaching that network and gaining access to personally identifiable information would be assigned more points than a control related to preventing an employee from accidentally mis-configuring a network router setting that results in a network outage).
- Risks related to legal and external drivers, such as contracts, regulations, and public commitments, for each control.

The displayed Compliance Score values for the control are applied *in their entirety* to the Total Compliance Score on a pass/fail basis--either the control is implemented and passes the subsequent assessment test or it does not; there is no partial credit for a partial implementation. Only when the control has its **Implementation Status** set to **Implemented** or **Alternative Implementation** and the **Test Result** is set to **Passed** are the assigned points added to the Total Compliance Score.

Most importantly, the Compliance Score can help you prioritize which controls to focus on for implementation by indicating which controls that have a higher potential risk if there is a failure related to a control. In addition to risk-based prioritization, when assessment controls are related to other controls (either within the same assessment or in another assessment in the same assessment grouping), completing a single control successfully can result in a significant reduction of effort based on the synchronization of control test results.

For example, in the image below we see that the Office 365 - GDPR Assessment is currently 46% assessed, with 51 of 111 control assessments completed for a Total Compliance score of 289 out of a possible 600.

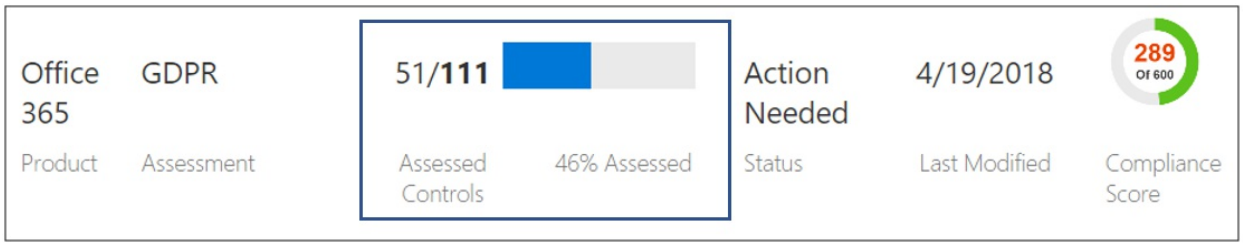

Within the assessment GDPR control 7.5.5 is related to 5 other controls (7.4.1, 7.4.3, 7.4.4, 7.4.8, and 7.4.9) each with a moderate to high severity risk rating score of 6 or 8). Using the assessment filter, we have selected all of these controls, making them visible in the assessment view, and can see below that none of them have been assessed.

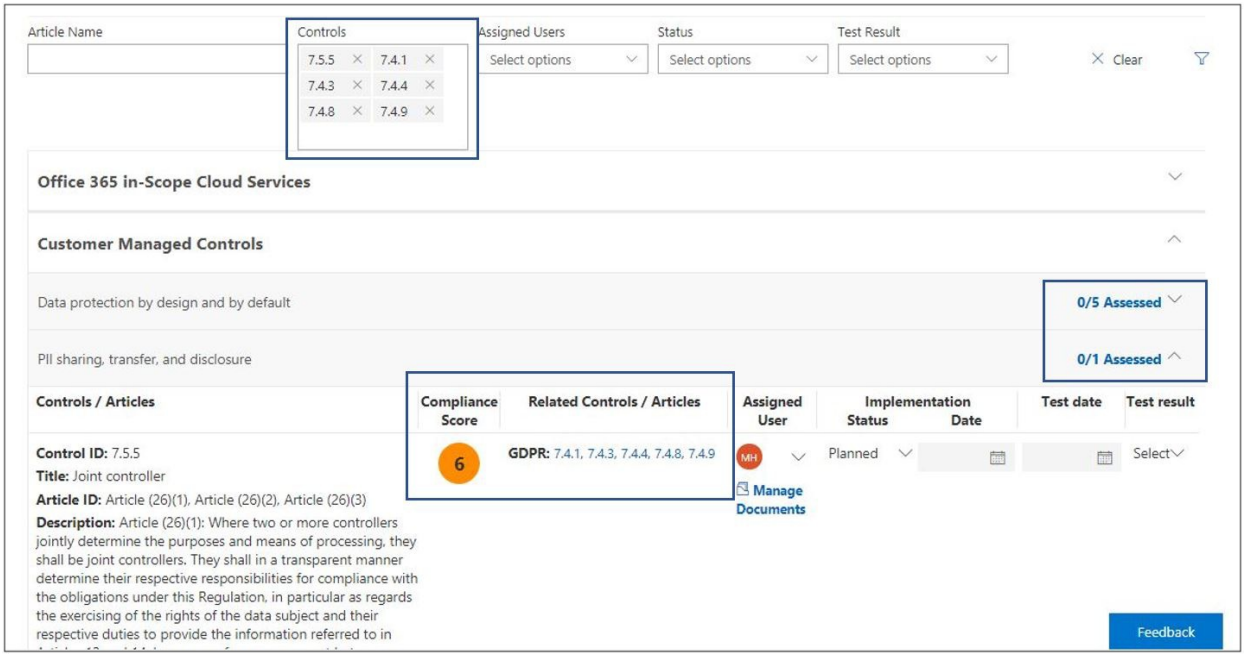

As those 6 controls are related, the completion of any one them will result in a synchronization of those test results across the related controls within this assessment (just as it will for any related controls in an assessment that is in the same assessment grouping). Upon completion of the implementation and testing of GDPR control 7.5.5, the control detail area refreshes to show that all 6 controls have been assessed, with a corresponding increase in the number of assessed controls to 57 and 51% assessed, and a change in total Compliance Score of +40.

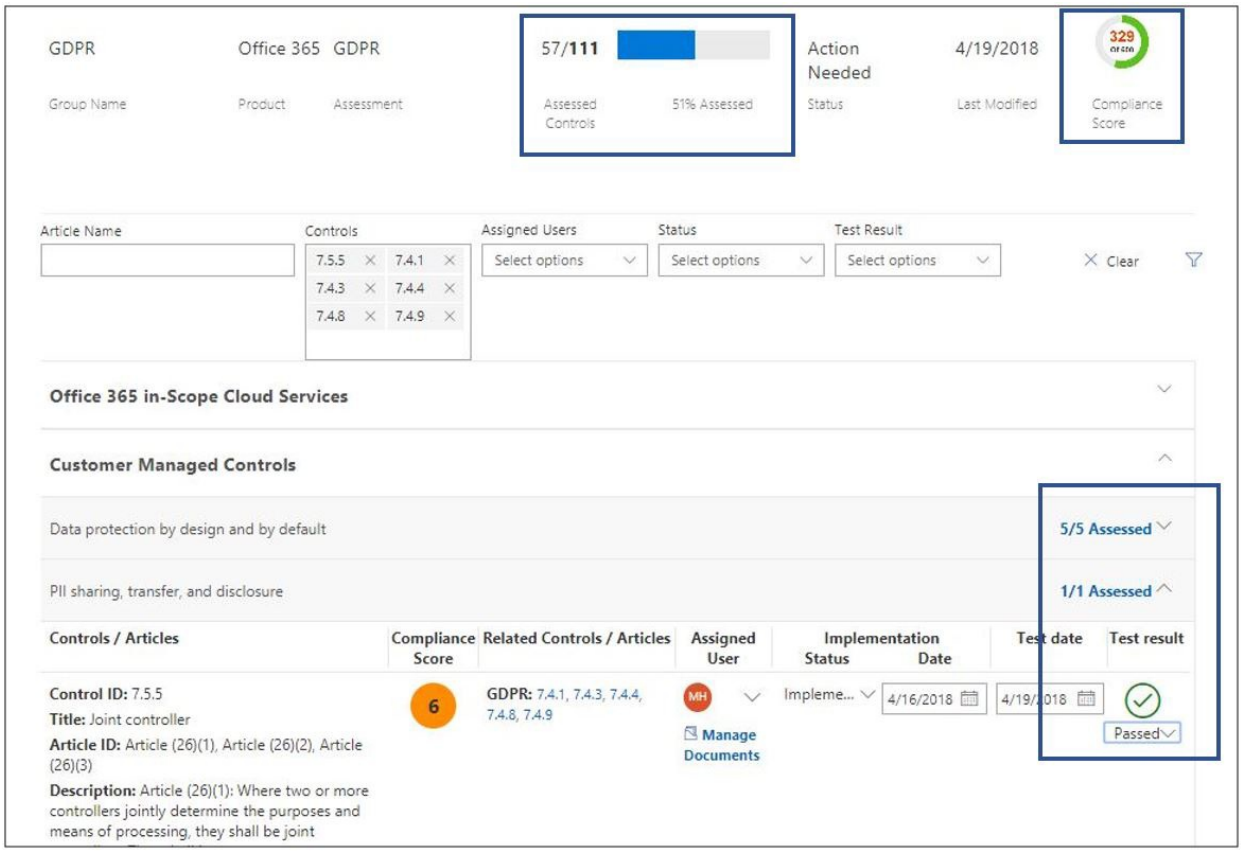

This confirmation update dialog box will appear if you are about to change the Implementation Status of a related control in a way that will impact the other related controls.

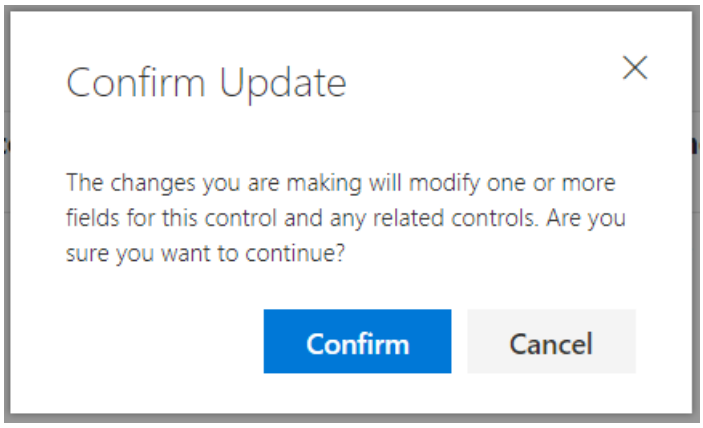

NOTE

Currently, only Assessments for Office 365 cloud services include a Compliance Score. Assessments for Azure and Dynamics show an assessment status.

Compliance Score methodology

The Compliance Score, like the Microsoft Secure Score, is similar to other behavior-based scoring systems; your organization's activity can increase its Compliance Score by performing activities related to data protection, privacy, and security.

NOTE

The Compliance Score does not express an absolute measure of organizational compliance with any particular standard or regulation. It expresses the extent to which you have adopted controls which can reduce the risks to personal data and individual privacy. No service can guarantee that you are compliant with a standard or regulation, and the Compliance Score should not be interpreted as a guarantee in any way.

Assessments in Compliance Manager are based on the shared responsibility model for cloud computing. In the shared responsibility model, Microsoft and each customer share responsibility for the protection of the customer's data when that data is stored in our cloud.

As shown in the Office 365 GDPR Assessment below, Microsoft and customers are each responsible for performing a variety of Actions that are designed to satisfy the requirements of the standard or regulation being assessed. To rationalize and understand the required. Actions across a variety of standards and regulations, Compliance Manager treats all standards and regulations as if they were control frameworks. Thus, the Actions performed by Microsoft and by customers for each Assessment involve the implementation and validation of various controls.

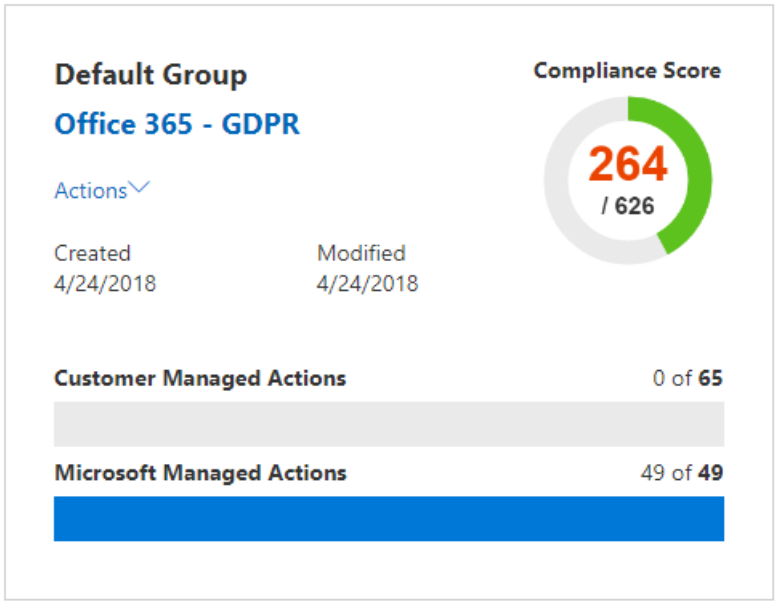

Here's the basic workflow for a typical Action:

1. The Compliance, Risk, Privacy, and/or Data Protection Officer of an organization assigns the task to someone in the organization to implement a control. That person could be:
   - A business policy owner
   - An IT implementer
   - Another individual in the organization who has responsibility for performing the task
2. That individual performs the tasks necessary to implement the control, uploads evidence of implementation into Compliance Manager, and marks the control(s) tied to the Action as implemented. Once these tasks are completed, they assign the Action to an Assessor for validation. Assessors can be:
   - Internal assessors that perform validation of controls within an organization
   - External assessors that examine, verify, and certify compliance, such as the third-party independent organizations that audit Microsoft's cloud services
3. The Assessor validates the control and examines the evidence and marks the control(s) as assessed and the results of the assessment (e.g., passed).

Once all the controls associated with an Assessment have been assessed, the Assessment is considered completed.

Every Assessment in Compliance Manager comes pre-loaded with information that provides details about the Actions taken by Microsoft to satisfy the requirements of the controls for which Microsoft is responsible. This information includes details about how Microsoft has implemented each control and how and when Microsoft's implementation was assessed and verified by a third-party auditor. For this reason, the Microsoft Managed Controls for each Assessment are marked as Assessed, and the Compliance Score for the Assessment reflects this.

Each Assessment includes a total Compliance Score based on the shared responsibility model. Microsoft's implementation and testing of controls for Office 365 contributes a portion of the total possible points associated with a GDPR assessment. As the customer implements and tests each of the customer Actions, the Compliance Score for the Assessment will increase by the value assigned to the control.

Risk-based scoring methodology

Compliance Manager uses a risk-based scoring methodology with a scale from 1-10 that assigns a higher value to controls that represent a higher risk in the event the control fails or is non-compliant. The scoring system

used by Compliance Score is based on several key factors, such as:

- The essence of the control
- The level of risk of the control based on the kinds of threats
- The external drivers for the control

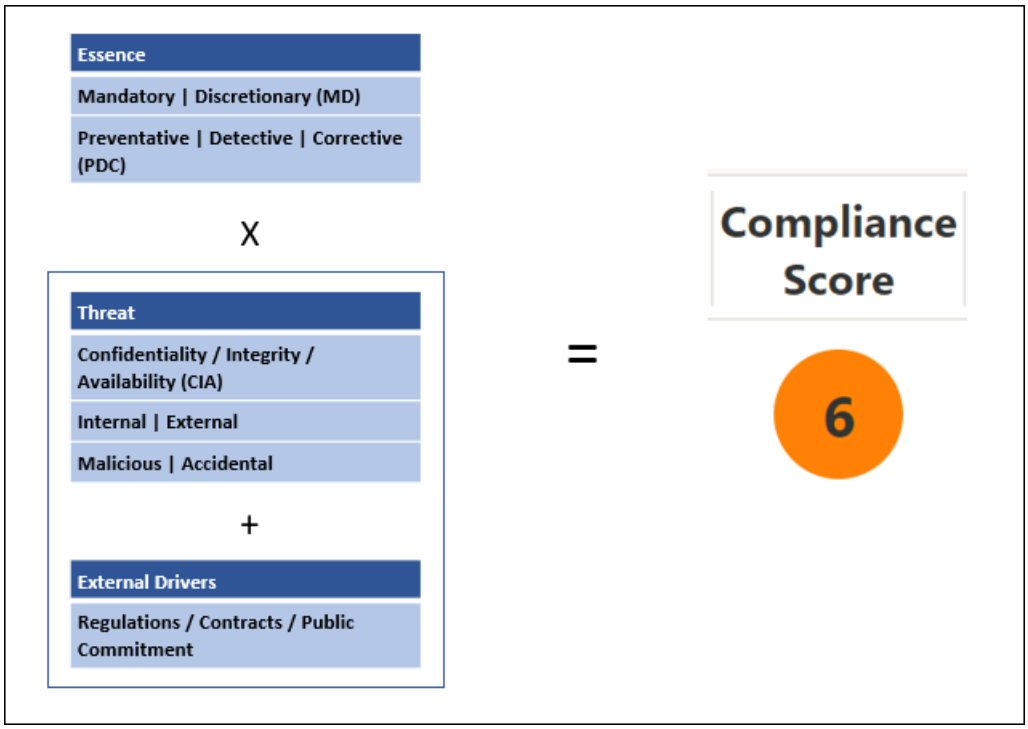

Essence of the control

The essence of the control is based on whether the control is Mandatory or Discretionary, and whether it is Preventative, Detective, or Corrective.

Mandatory or discretionary

Mandatory controls are controls that cannot be bypassed either intentionally or accidentally. An example of a common mandatory control is a centrally-managed password policy that sets requirements for password length, complexity, and expiration. Users must comply with these requirements in order to access the system.

Discretionary controls rely upon users to understand policy and act accordingly. For example, a policy requiring users to lock their computer when they leave it is a discretionary control because it relies on the user.

Preventative, detective, or corrective

Preventative controls are those that prevent specific risks. For example, protecting information at rest using encryption is a preventative control against attacks, breaches, etc. Separation of duties is a preventative control to manage conflict of interest and to guard against fraud.

Detective controls are those that actively monitor systems to identify irregular conditions or behaviors that represent risk or that can be used to detect intrusions or determine if a breach has occurred. System access auditing and privileged administrative actions auditing are types of detective monitoring controls; regulatory compliance audits are a type of detective control used to find process issues.

Corrective controls are those that try to keep the adverse effects of a security incident to a minimum, take corrective action to reduce the immediate effect, and reverse the damage, if possible. Privacy incident response is a corrective control to limit damage and restore systems to an operational state after a breach.

By evaluating each control using these factors, we determine the essence of the control and assign it a value relative to the risk that it represents.

Threat:

CONTROL	MANDATORY	DISCRETIONARY
Preventative	High risk	Medium risk
Detective	Medium risk	Low risk
Corrective	Medium risk	Low risk

Threat refers to anything that poses a risk to the fundamental, universally-accepted security standard known as the CIA triad for data: Confidentiality, Integrity, and Availability:

- Confidentiality means that information can be read and understood only by trusted, authorized parties.
- Integrity means that information has not been modified or destroyed by unauthorized parties.
- Availability means that information can be accessed readily with a high level of quality of service.

A failure of any of these characteristics is considered a compromise of the system as a whole. Threats can come from both internal and external sources, and an actor's intent can be accidental or malicious. These factors are estimated in a threat matrix that assigns threat levels of either High, Moderate, or Low to each combination of scenarios.

FACTOR	INTERNAL	INTERNAL	EXTERNAL	EXTERNAL
	Malicious	*Accidental*	*Malicious*	*Accidental*
Confidentiality	(H, M, or L)	(H, M, or L)	(H, M, or L)	(H, M, or L)
Integrity	(H, M, or L)	(H, M, or L)	(H, M, or L)	(H, M, or L)
Availability	(H, M, or L)	(H, M, or L)	(H, M, or L)	(H, M, or L)

External drivers:

CONTRACTS	REGULATIONS	PUBLIC COMMITMENTS
(H, M, or L)	(H, M, or L)	(H, M, or L)

External factors such as applicable regulations, contracts, and public commitments can influence controls designed to protect data and prevent data breaches, and each of these factors are assigned risk values or High, Moderate or Low.

The estimated number of occurrences of these risk values of High, Moderate, or Low across the 15 possible risk scenarios represented in the CIA/Threat and Legal/External Drivers are combined to provide a risk weighting, which considers the likelihood and number of occurrences of risks at a given value as significant and is taken into consideration when calculating the severity ranking of the control.

Based on the control's severity ranking, the control is assigned its compliance score value, a number between 1 (low) and 10 (high), grouped into the following categories of risk:

RISK LEVEL	CONTROL VALUE
Low	1-3
Moderate	6
High	8
Severe	10

By prioritizing assessment controls with the highest compliance score values, the organization will be concentrating on the highest risk items and receive proportionally higher positive feedback in the form of more points added to the total compliance score for the assessment for each control assessment completed.

Summary of scoring methodology

The Compliance Score is a core component of the way that Compliance Manager helps organizations understand and manage their compliance. The Compliance Score for an assessment is an expression of the company's compliance with a given standard or regulation as a number, where the higher the score (up to the maximum number of points allocated for the Assessment), the better the company's compliance posture. Understanding the compliance scoring methodology in which assessment controls are assigned risk severity values between 1- 10 (low to high), and how completed control assessments add to the total compliance score is crucial to organizations for prioritizing their actions.

Grouping Assessments

When you create a new Assessment, you're prompted to create a group to assign the Assessment to or assign the Assessment to an existing group. Groups allow you to logically organize Assessments and share common information and workflow tasks between Assessments that have the same or related customer-managed controls.

For example, you could group Assessments by year or teams, departments, or agencies within your organization or group them by year. Here are some examples of groups and the Assessments they might contain.

- GDPR Assessments — 2018
  - Office 365 + GDPR
  - Azure + GDPR
  - Dynamics + GDPR
- Azure Assessments — 2018
  - Azure + GDPR
  - Azure + ISO 27001:2013
  - Azure + ISO 27018:2014
- Data Security and Privacy Assessments
  - Office 365 + ISO 27001:2013
  - Office 365 + ISO 27018:2014
  - Azure + ISO 27001:2013
  - Azure + ISO 27018:2014

TIP

We recommend that your determine a grouping strategy for your organization before adding new assessments.

These are the requirements for grouping Assessments:

- Group names (also called *Group IDs) must be unique within your organization.
- Groups can contain Assessments for the same certification/regulation, but each group can only contain one Assessment for a specific cloud service/certification pair. For example, a group can't contain two Assessments for Office 365 and GDPR. Similarly, a group can contain multiple Assessments for the same cloud service as long as the corresponding certification/regulation for each one is different.

Once an assessment has been added to an assessment grouping, the grouping cannot be changed. You can rename the assessment group, which changes the name of the assessment grouping for all of the assessments associated with that group. You can create an assessment and a new assessment group and copy information from an existing assessment, which effectively creates a duplicate of that assessment in a different assessment group. Archiving an assessment breaks the relationship between that assessment and the assessment group. Any further updates to other related assessments are no longer reflected in the archived assessment.

As previous explained, one key advantage of using groups is that when two different Assessments in the same group share the same customer-managed control (and therefore the customer actions would be the same for each control), then the completion of implementation details, testing information, and status for the control in one Assessment would be synchronized to the same control in any other Assessment in the group. In other words, if Assessments share the same control and those Assessments are in the same group, you'd only have to manage the assessment process for the control in one Assessment. The results for that control will be automatically synchronized to other Assessments. For example, ISO 27001 and ISO 27018 both have a control related to password policies. If the Test Status for the control is set to "Passed" in one Assessment, the control is updated (and marked as "Passed") in the other Assessment, as long as both assessments are part of the same Assessment Group.

As an example of this, consider these two related assessment controls, each having to do with encryption of data on public networks, control 6.10.1.2 in the Office 365 — GDPR assessment, and control SC-13 in the Office 365 — NIST 800-53 assessment. These are related assessment controls, in two different assessments, both in the Default Group. Initially, neither assessment has completed any customer control assessments, as is displayed on the Compliance Manager Dashboard that shows these two Assessments.

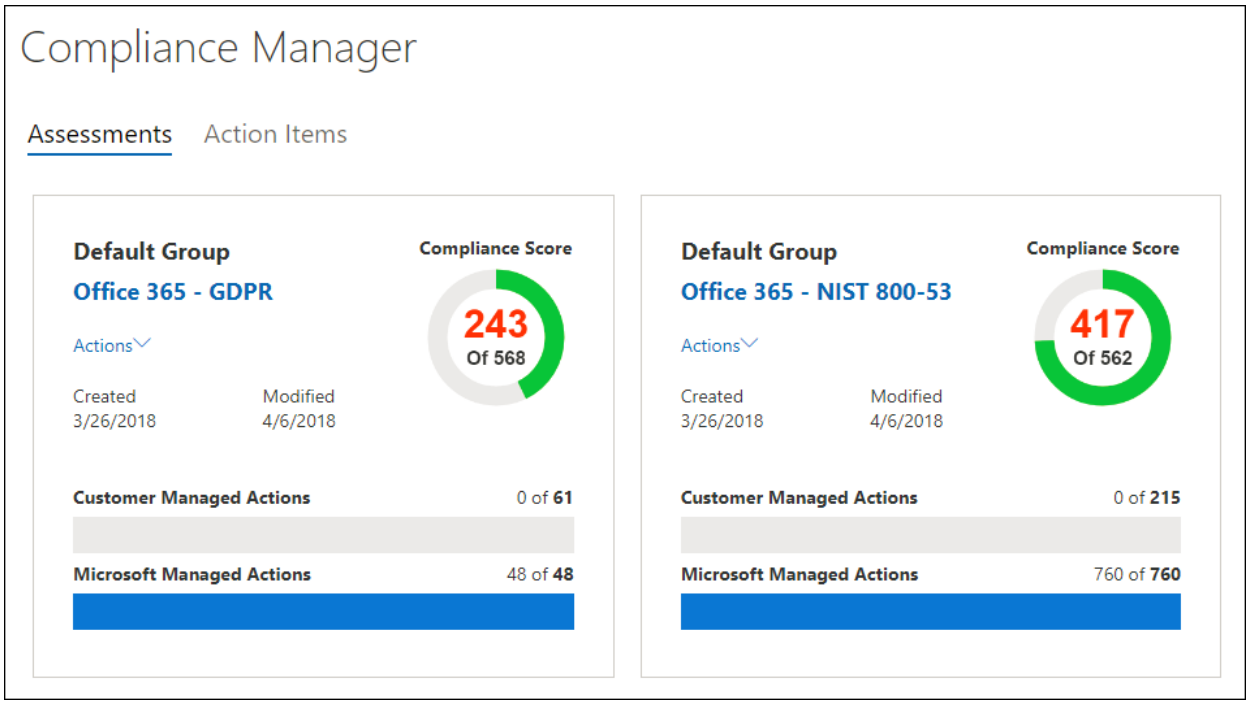

By clicking the **Office 365 — GDPR** assessment, and using the filter controls to view GDPR control 6.10.1.2, we see that NIST 800-53 control SC-13 is listed as a related control.

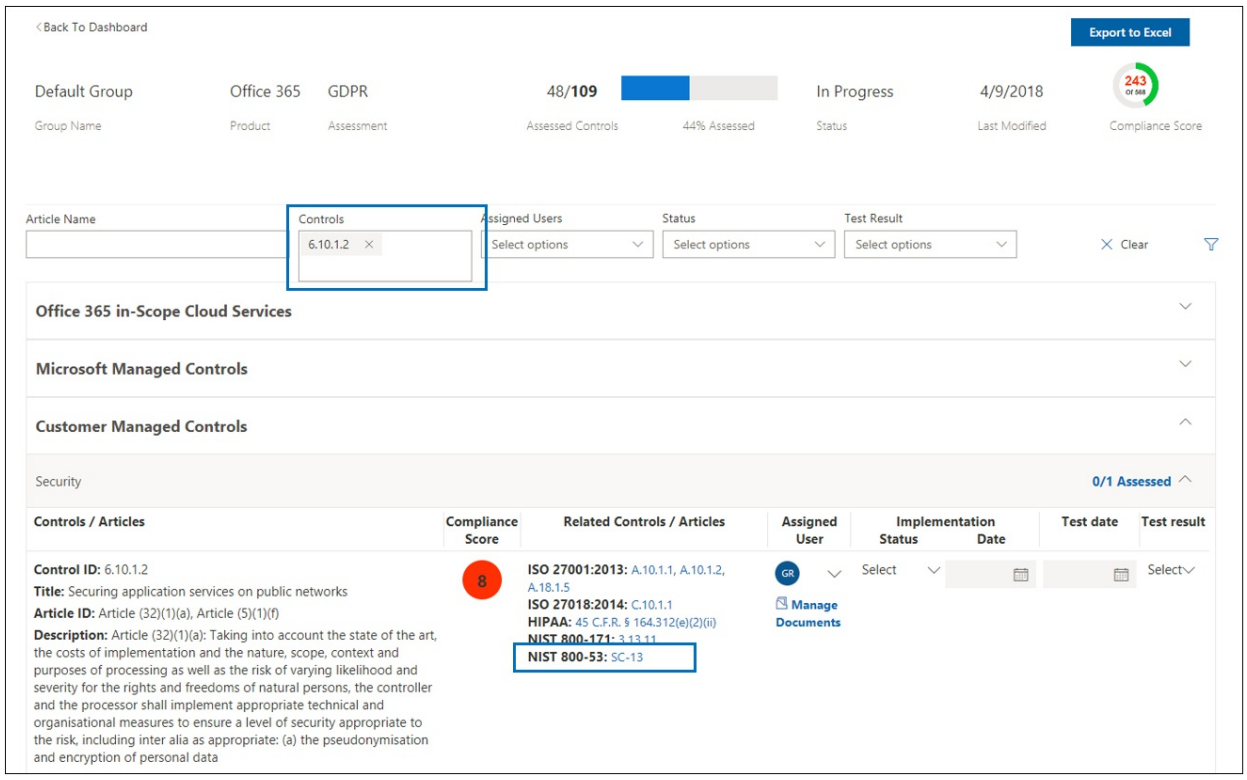

Here we show the completion of the implementation and testing of GDPR control 6.10.1.2.

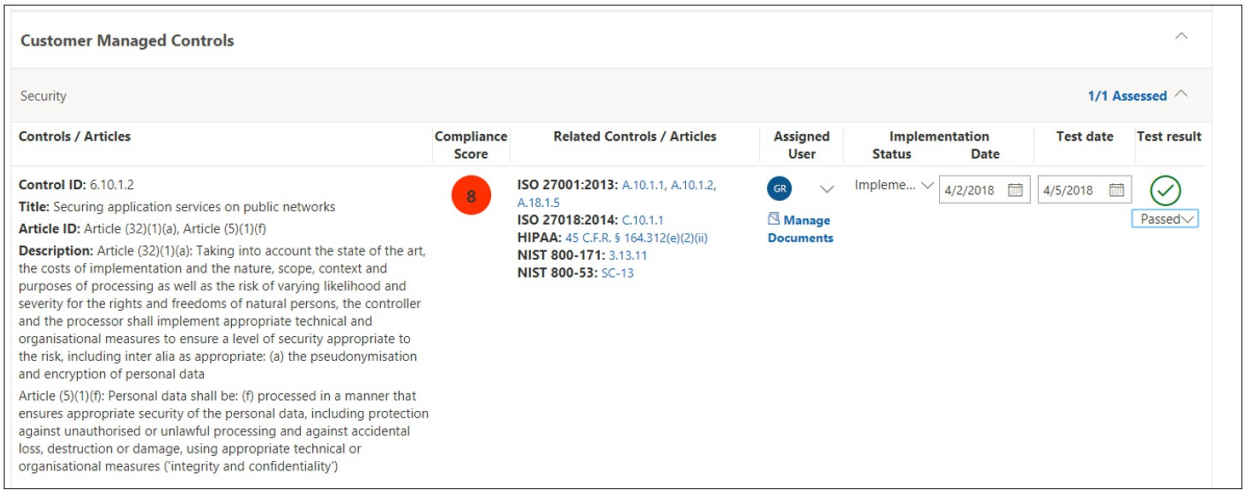

By navigating to the related control in the grouped assessment, we see that NIST 800-53 SC-13 has also been marked as completed with the same date and time, with no additional implementation or testing effort.

Back at the Dashboard, we can see that each assessment has one control assessment completed and that the total Compliance Score for each assessment has increased by 8 (the compliance score value of that shared control).

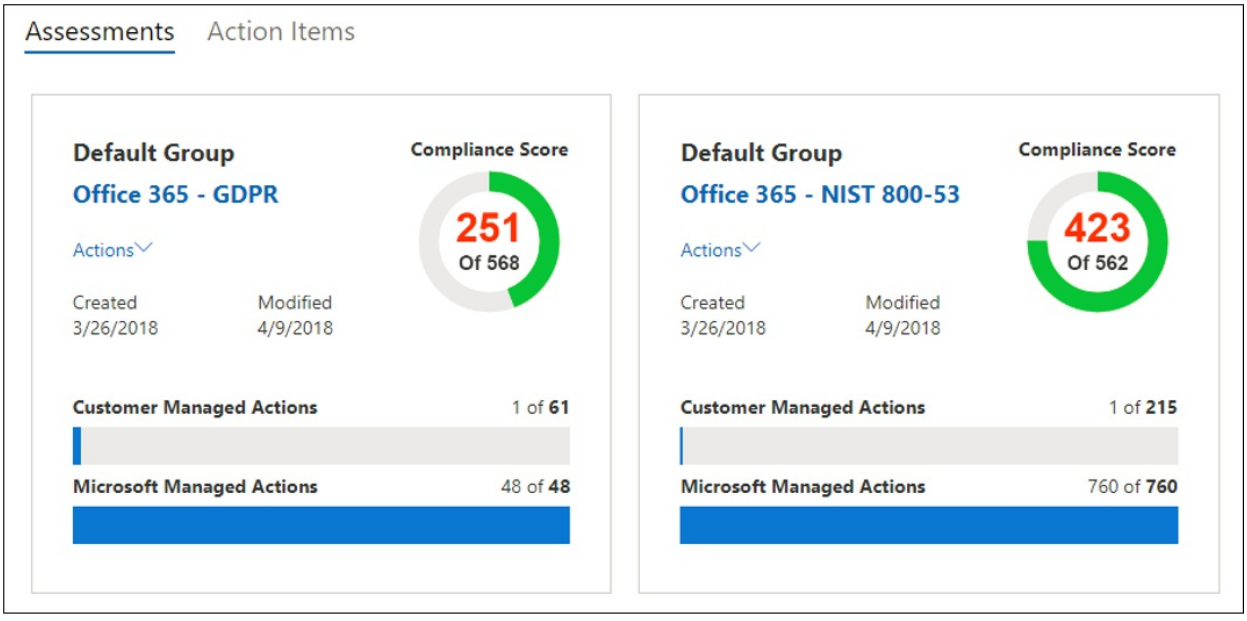

Administrative functions

There are specific administrative functions that are only available to the tenant administrator account, and will only be visible when logged in as a global administrator.

> **NOTE**
>
> The Access to Restricted Documents permission in the drop-down list will allow administrators to give users access to restricted documents that Microsoft shares on the Service Trust Portal. The Restricted Documents feature isn't available, but is coming soon.

Assigning Compliance Manager roles to users

Each Compliance Manager role has slightly different permissions. You can view the permissions assigned to each role, see which users are in which roles, and add or remove users from that role through the Service Trust Portal by selecting the **Admin** menu item, and then choosing **Settings**.

To add or remove users from Compliance Manager roles.

1. Go to https://servicetrust.microsoft.com.
2. Sign in with your Azure Active Directory global administrator account.
3. On the Service Trust Portal top menu bar, click **Admin** and then choose **Settings**.
4. In the **Select Role** drop-down list, click the role that you want to manage.
5. Users added to each role are listed on the **Select Role** page.
6. To add users to this role, click **Add**. In the **Add Users** dialog, click the user field. You can scroll through the list of available users or begin typing the user name to filter the list based on your search term. Click the user to add that account to the **Add Users** list to be provisioned with that role. If you would like to add multiple users concurrently, begin typing a user name to filter the list, and then click the user to add to the list. Click **Save** to provision the selected role to these users.

   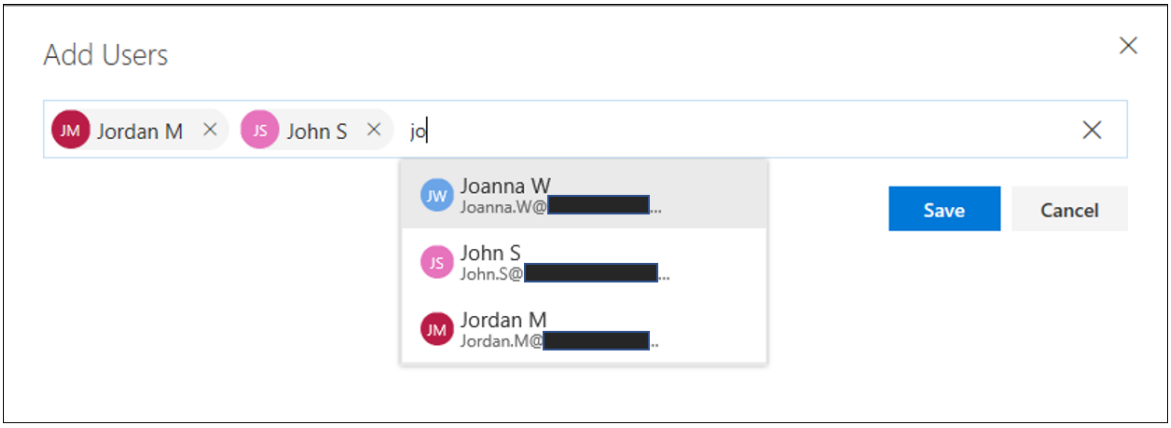

7. To remove users from this role, select the user(s) and click **Delete**.

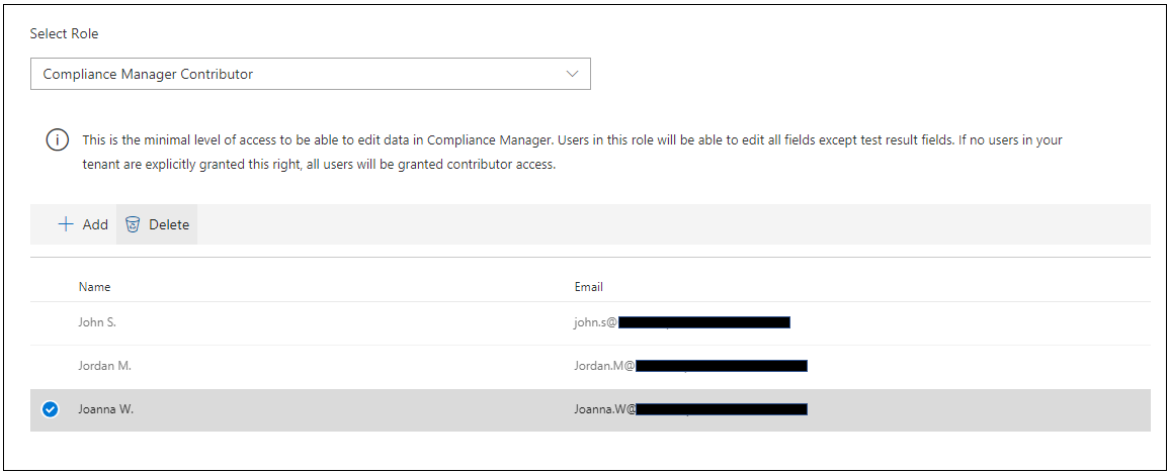

User Privacy settings

Certain regulations require that an organization must be able to delete user history data. To enable this, Compliance Manager provides the **User Privacy Settings** functions, that allow administrators to:

- Search for a user
- Export a report of account data history
- Reassign action items
- Delete user data history

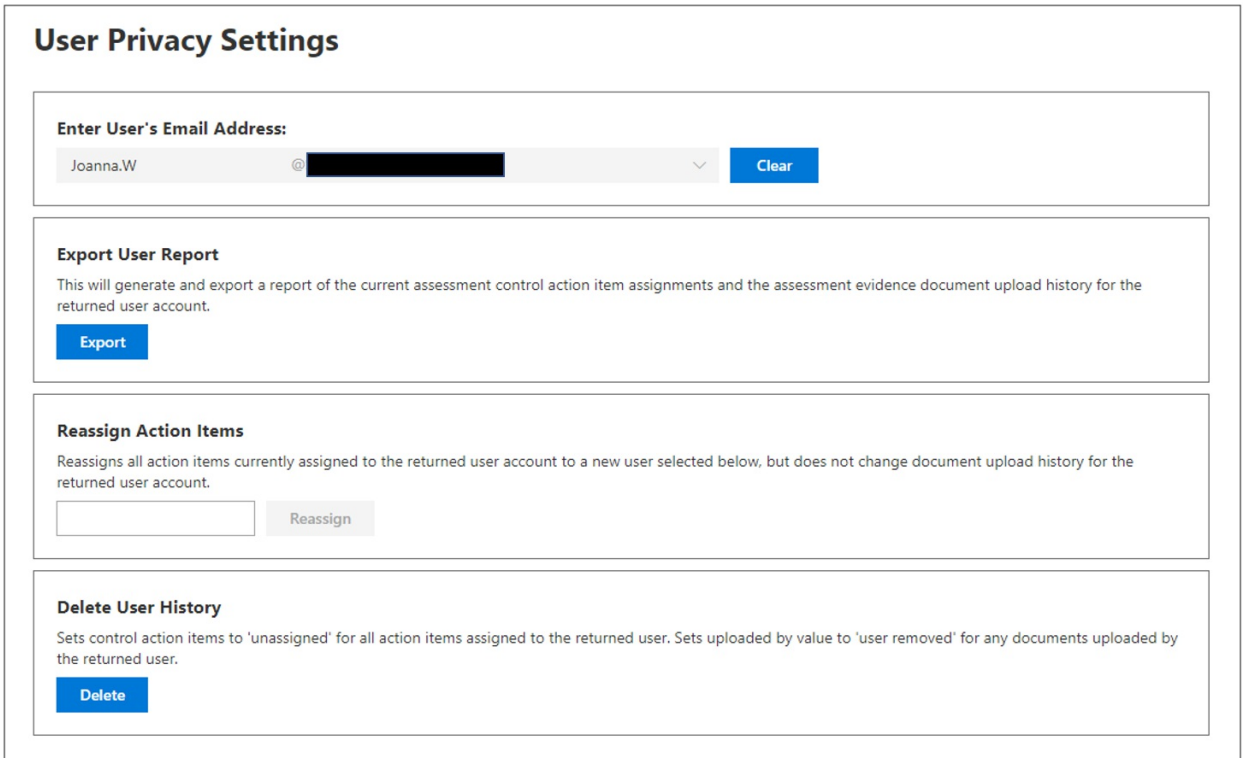

Search for a user

To search for a user account:

1. Enter the user email address by typing in the alias (the information to the left of the @ symbol) and choosing the domain name by clicking the domain suffix list on the right. If this is tenant with multiple registered domains, you can double check the email address domain name suffix to ensure that it is correct.
2. When you have the username correctly entered, click **Search**.

3. If the user account is not found, the error message 'User not found' will be displayed on the page. Check the user's email address information, make corrections as necessary and click **Search** to try again.
4. If user account is found, the text of the button changes from **Search** to **Clear**, which indicates that the returned user account is the operating context for the additional functions that will be displayed below, that running those functions will apply to this user account.
5. To clear search results and search for a different user, click **Clear**.

Export a report of account data history

Once the user account has been identified, you may wish to generate a report of dependencies that exist linked to this account. This information allows you to reassign open action items or ensure access to previously uploaded evidence.

To generate and export a report:

1. Click **Export** to generate and download a report of the Compliance Manager control action items currently assigned to the returned user account and the list of documents uploaded by that user. If there are no assigned actions or uploaded documents, an error message will state "No data for this user".
2. The report downloads in the background of the active browser window — if you don't see a download popup you want to check your browser download history.
3. Open the document to review the report data.

> **NOTE**
> This is not a historical report that retains and displays state changes to action item assignment history. The generated report is a snapshot of the control action items assigned at the time that the report is run (date and time stamp written into the report). For instance, any subsequent reassignment of action items will result in different snapshot report data if this report is generated again for the same user.

Reassign action items

This function enables an organization to remove any active or outstanding dependencies on the user account by reassigning all action item ownership (which includes both active and completed action items) from the returned user account to a new user selected below. This action does not change document upload history for the returned user account.

To reassign action items to another user:

1. Click the input box to browse for and select another user within the organization to whom the returned user's action items should be assigned.
2. Select **Replace** to reassign all control action items from the returned user to the newly selected user.
3. A confirmation dialog box appears stating "This will reassign all control action items from the current user to the selected user. This action cannot be undone. Are you sure you want to continue?"
4. To continue click **OK**, otherwise click **Cancel**.

> **NOTE**
> All action items (both active and completed) will be assigned to the newly selected user. However, this action does not affect the document upload history; any documents uploaded by the previously assigned user will still show the date/time and name of the previously assigned user.

Changing the document upload history to remove the previously assigned user will have to be done as a

manual process. In that case, the administrator will need to:

1. Open the previously downloaded Export report.
2. Identify and navigate to the desired control action item.
3. Click **Manage Documents** to navigate to the evidence repository for that control.
4. Download the document.
5. Delete the document in the evidence repository.
6. Re-upload the document. The document will now have a new upload date, time and Uploaded By username.

Delete user data history

This sets control action items to 'unassigned' for all action items assigned to the returned user. This also sets uploaded by value to 'user removed' for any documents uploaded by the returned user

To delete the user account action item and document upload history:

1. Click **Delete**.

   A confirmation dialog will be displayed, stating "This will remove all control action item assignments and the document upload history for the selected user. This action cannot be undone. Are you sure you want to continue?"

2. To continue click **OK**, otherwise click **Cancel**.

Using Compliance Manager

Compliance Manager provides you with tools to assign, track, and record compliance and assessment-related activities, and to help your organization cross team barriers to achieve your organization's compliance goals.

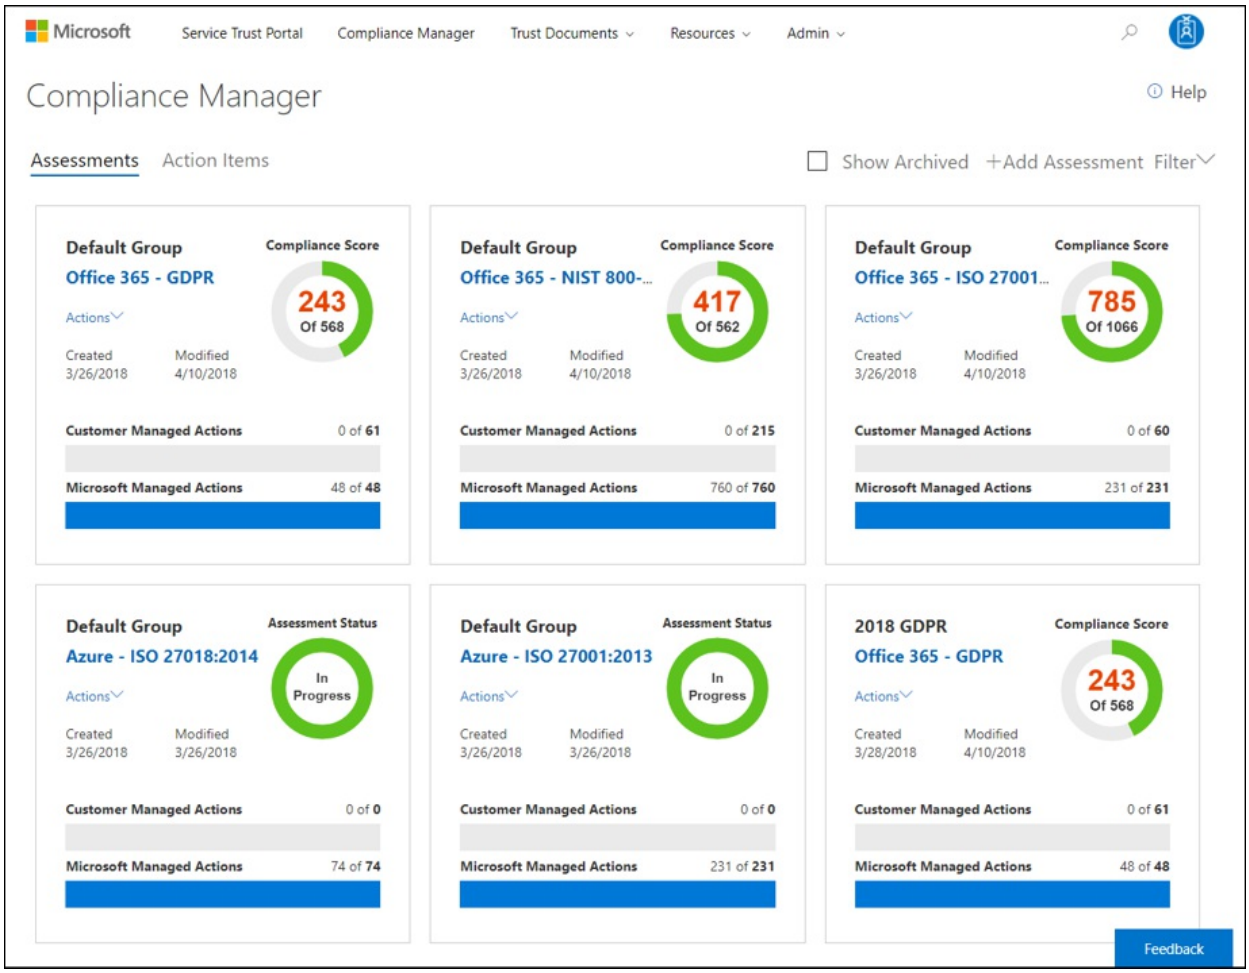

Accessing Compliance Manager

You access Compliance Manager from the Service Trust Portal. Anyone with a Microsoft account or Azure Active Directory organizational account can access Compliance Manager.

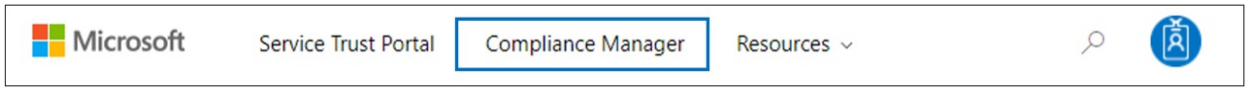

1. Go to https://servicetrust.microsoft.com.
2. Sign in with your Azure Active Directory (Azure AD) user account.
3. In the Service Trust Portal, click **Compliance Manager**.
4. When the Non-Disclosure Agreement is displayed, read it, and then click **Agree** to continue. You'll only have to do this once, and then the Compliance Manager dashboard is displayed.

   To get you started, we've added the following Assessments by default:

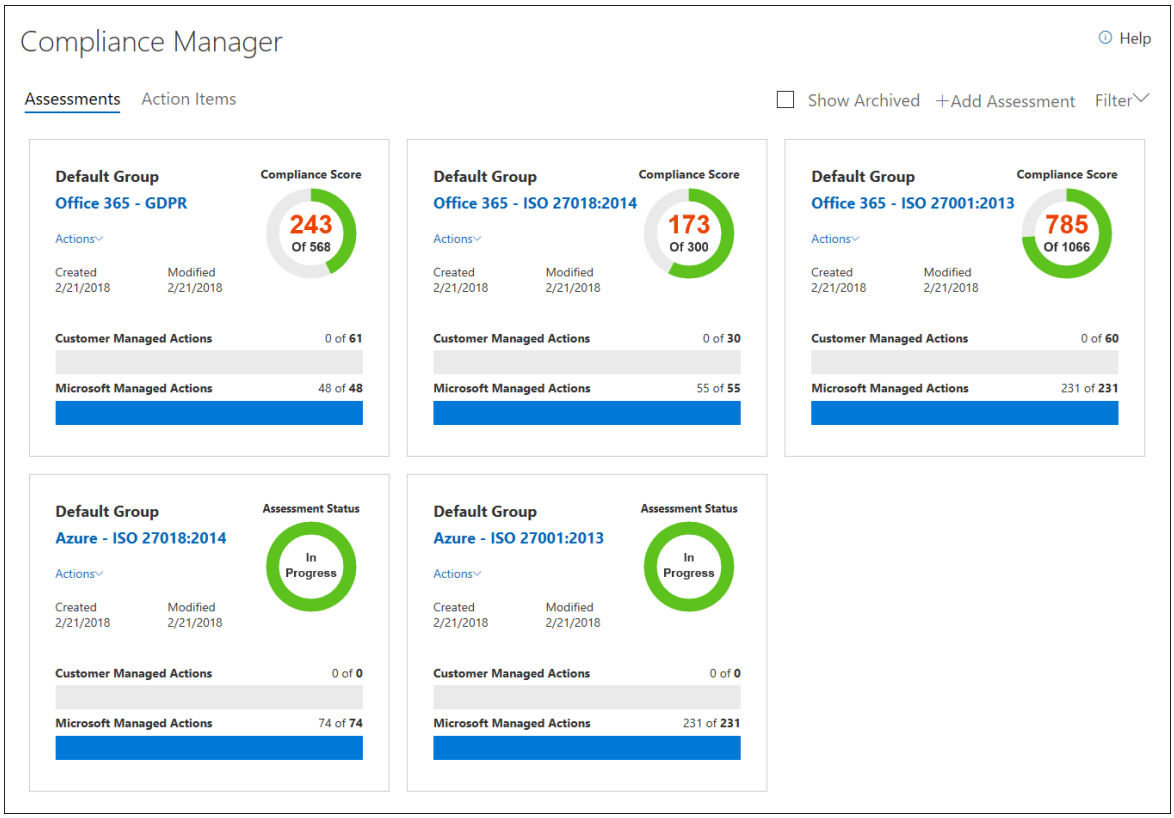

5. Click ⓘ **Help** to take a short tour of Compliance Manager.

Viewing action items

Compliance Manager provides a convenient view of all your assigned control assessment action items, enabling you to quickly and easily take action on them. You can view all action items or select the action items that correspond with a specific certification by clicking the tab associated with that assessment. For instance, in the image below, the GDPR tab has been selected, showing controls that related to the GDPR assessment.

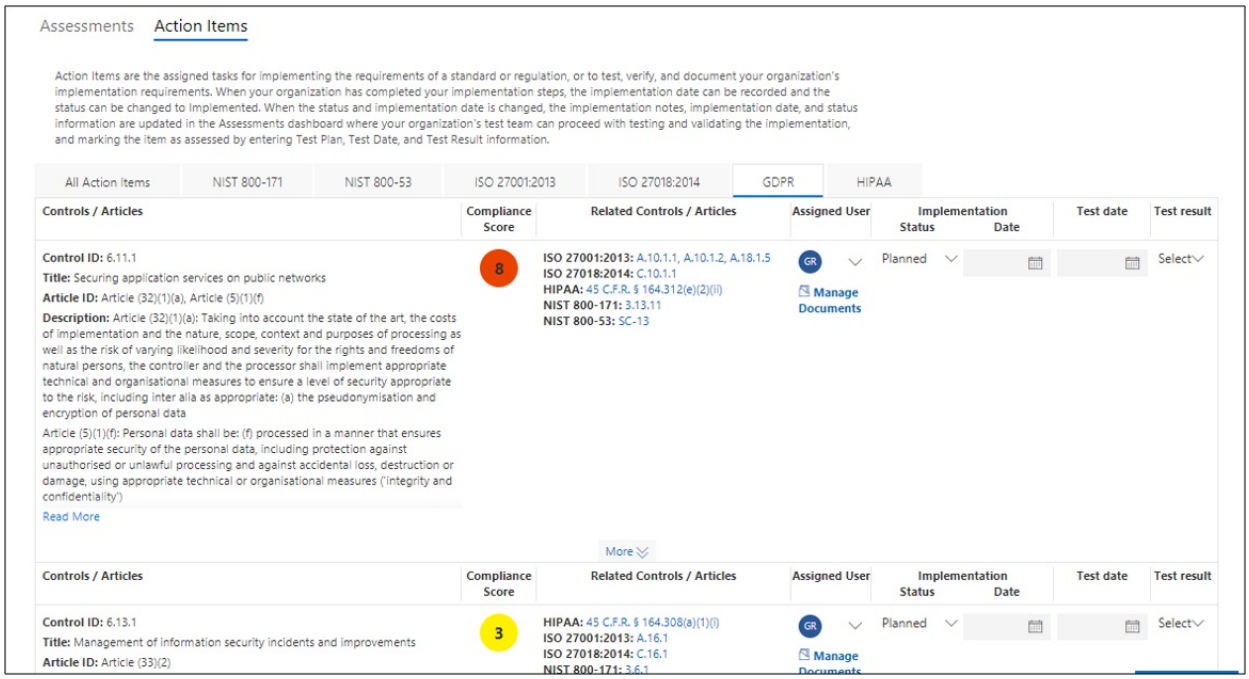

To view your action items:

1. Go to the Compliance Manager dashboard
2. Click the **Action Items** link, and the page will refresh to show the action items that have been assigned to you.

By default, all action items are shown. If you have action items across multiple certifications, the names of the certifications will be listed in tabs across the top of the assessment control. To see the action items for a specific certification, click that tab.

Adding an Assessment

To add an Assessment to Compliance Manager:

1. In the Compliance Manager dashboard, click **+ Add Assessment**.
2. In the **Add an Assessment** window, you can create a new group to add the Assessment to or you can add it to an existing group (the built-in group is named "Initial Group".) Depending on the option you choose, either type the name of a new group or select an existing group from the drop-down list. For more information, see Grouping Assessments.

   If you create a group, you also have the option to copy information from an existing group to the new Assessment. That means any information that was added to the Implementation Details and Test Plan and Management Response fields of customer-managed controls from Assessments in the group that you're copying from are copied to the same (or related) customer-managed controls in the new Assessment. If you're adding a new Assessment to an existing group, common information from Assessments in that group will be copied to the new Assessment. For more information, see Copying information from existing Assessments.
3. Click **Next**, and do the following:

   a. Choose a Microsoft cloud service to assess for compliance from the **Select a product** drop-down list.

   b. Choose a certification to assess the selected cloud service against from the **Select a certification** drop-down list.
4. Click **Add to Dashboard** to create the Assessment; the assessment will be added to the Compliance Manager dashboard as a new tile at the end of the list of existing tiles.

   The **Assessment Tile** on the Compliance Manager dashboard, displays the assessment grouping, the name of the assessment (automatically created as a combination of the Service name and the certification selected), the date it was created and when it was last modified, the Total Compliance Score (which is the sum of all of the assigned control risk values that have been implemented, tested, and passed), and progress indicators along the bottom that show the number of controls that have been assessed.
5. Click the Assessment name to open it, and view the details of the Assessment.
6. Click the **Actions** menu to view your assigned action items, rename the assessment group, export the assessment report, or archive the assessment.

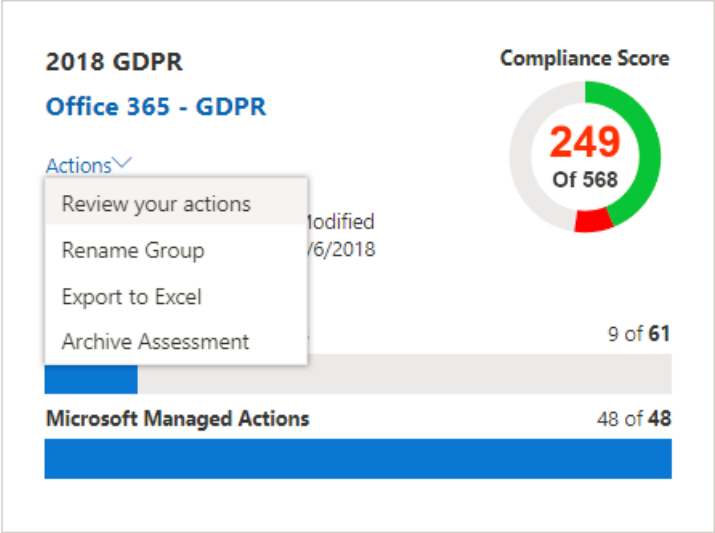

Copying information from existing Assessments

As previously explained, when you create an assessment group, you have the option to copy information from Assessments in an existing group to the new Assessment in the new group. This allows you to apply the assessment and testing work that's been completed to the same customer-managed controls in the new Assessment. For example, if you have a group for all GDPR-related Assessments in your organization, you can copy common information from existing assessment work when add a new Assessment to the group.

You can copy the following information from customer to a new Assessment:

- Assessment Users. An Assessment user is a user who the control is assigned to.
- Status, Test Date, and Test Results.
- Implementation details and test plan information.

Similarly, information from shared customer-managed controls within the same Assessment group is synchronized. And information in related customer-managed controls within the same Assessment is also synchronized.

Viewing Assessments

1. Locate the Assessment Tile corresponding to the assessment you wish to view, then click the assessment name to open it and view the Microsoft and customer-managed controls associated with the Assessment, along with a list of the cloud services that are in-scope for the Assessment. Here's an example of the Assessment for Office 365 and GDPR.

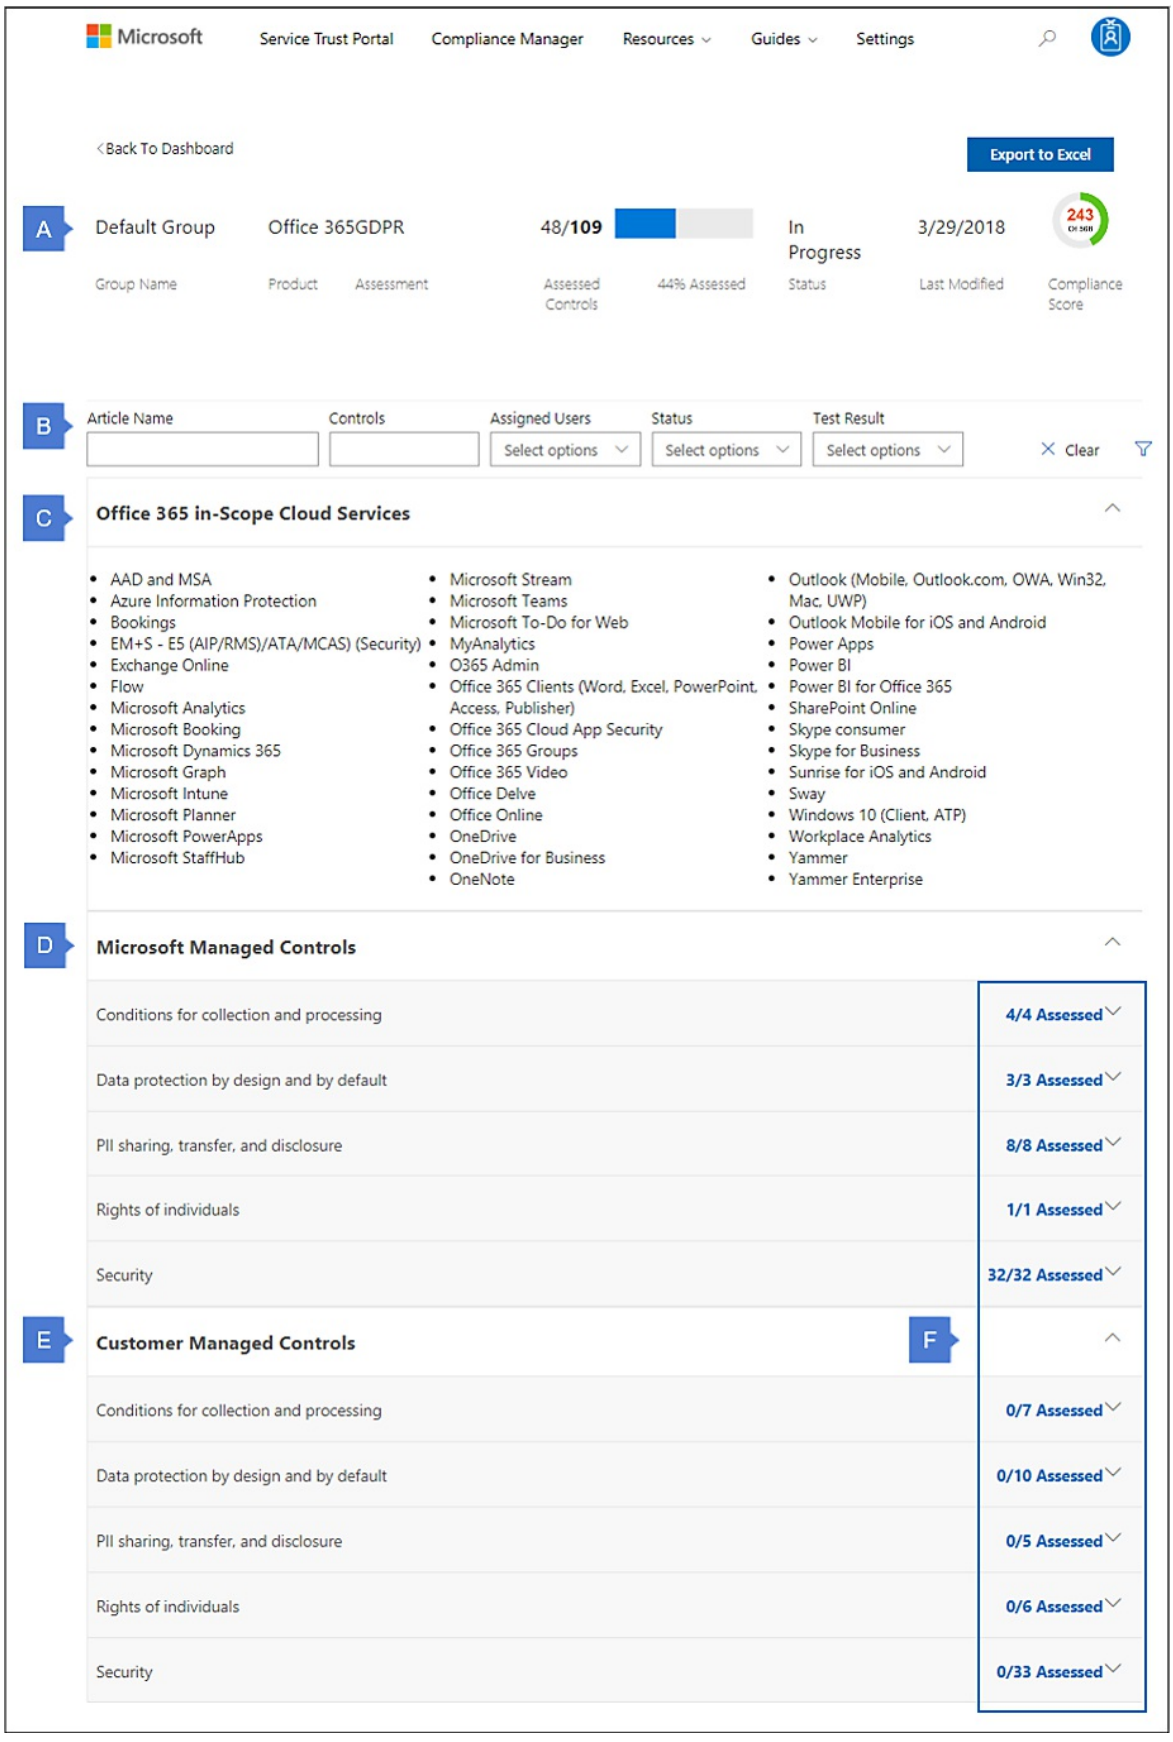

2. This section shows the Assessment summary information, including the name of the Assessment Grouping, Product, Assessment name, number of Assess controls

3. This section shows the Assessment Filter controls. For a more detailed explanation of how to use the Assessment Filter controls see the Managing the assessment process section.

4. This section shows the individual cloud services that are in-scope for the assessment.

5. This section contains Microsoft-managed controls. Related controls are organized by control family. Click

a control family to expand it and display individual controls.

6. This section contains customer-managed controls, which are also organized by control family. Click a control family to expand it and display individual controls.

7. Displays the total number of controls in the control family, and how many of those controls have been assessed. A key capability of Compliance Manager is tracking your organization's progress on assessing the customer-managed controls. For more information, see the Understanding the Compliance Score section.

Managing the assessment process

The creator of an Assessment is initially the only Assessment User. For each customer-managed control, you can assign an Action Item to a person in your organization so that person becomes an Assessment User who can perform the recommended Customer Actions, and gather and upload evidence. When you assign an Action Item, you can choose to send an email to the person that contains details including the recommended Customer Actions and the Action Item priority. The email notification includes a link to the **Action Items** dashboard, which lists all Action Items assigned to that person.

Here's a list of tasks that you can perform using the workflow features of Compliance Manager.

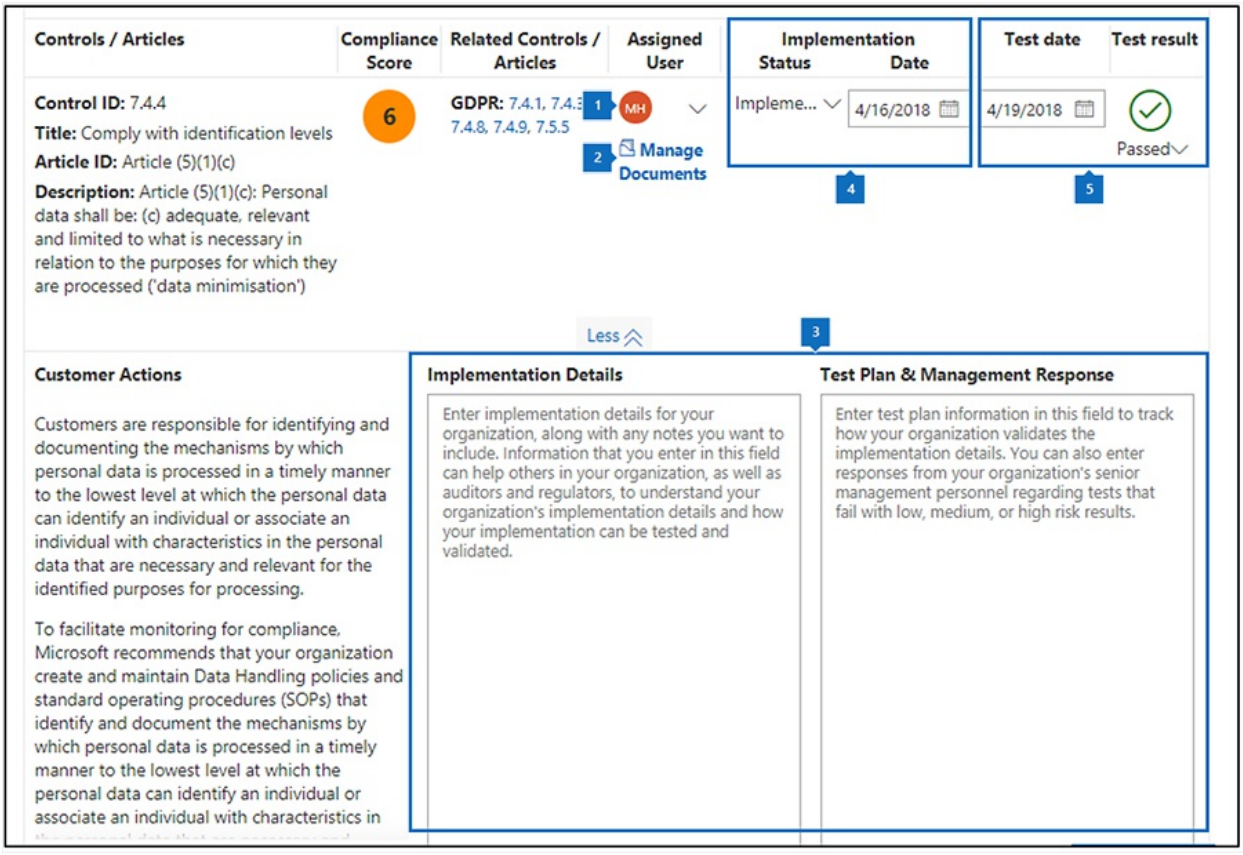

1. **Use the Filter Options to find specific assessment controls** - Compliance Manager provides **Filter Options**, giving you highly granular selection criteria for displaying assessment controls, helping you to precisely target specific areas of your compliance efforts.

   Click the funnel icon on the right-hand side of the page to show or hide the **Filter Options** controls. These controls allow you to specify filter criteria, and only the assessment controls that fit those criteria will be displayed below.

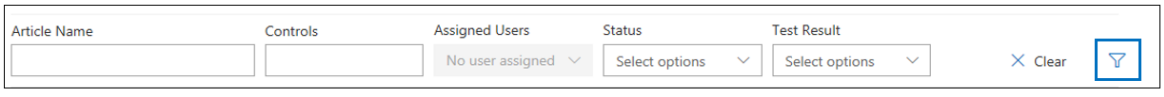

   - **Articles** - filters on the article name and returns the assessment controls associated to that article.

For instance, typing in "Article (5)" returns a selection list of articles whose name includes that string, i.e. Article (5)(1)(a), Article (5)(1)(b), Article (5)(1)(c), etc. Selecting Article (5)(1)(c) will return the controls associated with Article (5)(1)(c). This is multiselect field that uses an OR operator with multiple values — for instance, if you select Article (5)(1)(a) and then add Article (5)(1)(c), the filter will return controls associated with either Article (5)(1)(a) or Article (5)(1)(c).

- **Controls** - returns the list of controls whose names fit the filter, i.e. typing in 7.3 returns a selection list of items like 7.3.1, 7.3.4, 7.3.5, etc. This is multiselect field that uses an OR operator with multiple values — for instance, if you select 7.3.1 and then add 7.3.4, the filter returns controls associated with either 7.3.1 or 7.3.4.

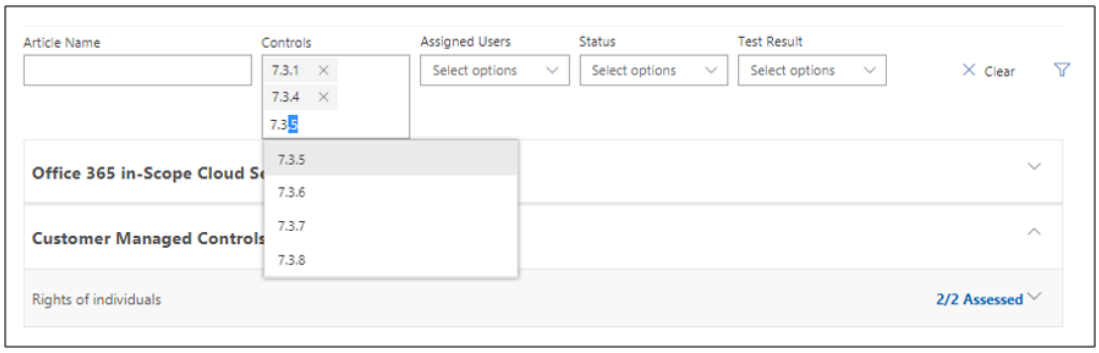

- **Assigned Users** - returns the list of controls who are assigned to the selected user.
- **Status** - returns the list of controls with the selected status.
- **Test Result** - returns the list of controls with the selected test result.

As you apply filter conditions, the view of applicable controls will change to correspond to your filter conditions. Expand the control family sections to show the control details below.

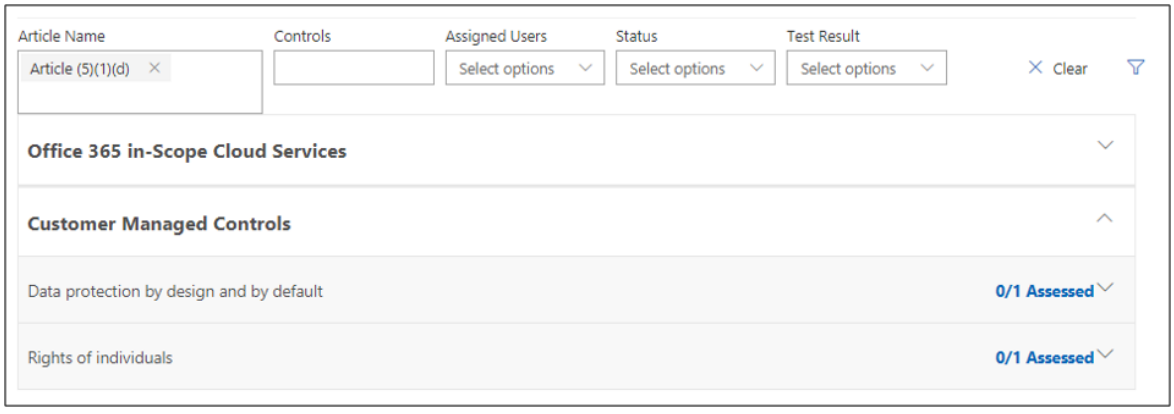

2. If after selecting the desired filters no results are shown, that means there are no controls that correspond to the specified filter conditions. For instance, if you select a particular **Assigned User** and then choose a

Control name that does correspond to the control assigned to that user, no assessments will be shown in the page below.

3. **Assign an Action Item to a user** - You can assign an Action Item to a person to implement the requirements of a certification/regulation, or to test, verify, and document your organization's implementation requirements. When you assign an Action Item, you can choose to send an email to the person that contains details including the recommended Customer Actions and the Action Item priority. You can also unassign or reassign an Action Item to a different person.

4. **Manage documents** - Customer-managed controls also have a place to manage documents that are related to performing implementation tasks and for performing testing and validation tasks. Anyone with permissions to edit data in Compliance Manager can upload documents by clicking **Manage Documents**. After a documented has been uploaded, you can click **Manage Documents** to view and download files.

5. **Provide implementation and testing details** - Every customer-managed control has an editable field where users can add implementation details that document the steps taken by your organization to meet the requirements of the certification/regulation, and to validate and document how your organization meets those requirements.

6. **Set Status** - Set the Status for each item as part of the assessment process. Available status values are **Implemented**, **Alternative Implementation**, **Planned**, and **Not in Scope**.

7. **Enter test date and test result** - The person with the Compliance Manager Assessor role can verify that proper testing performed, review the implementation details, test plan, test results, and any uploaded evidence, and then set the Test Date and Test Result. Available test result values are **Passed**, **Failed-Low Risk**, **Failed-Medium Risk**, and **Failed-High Risk**.

Managing action items

The people involved in the assessment process in your organization can use Compliance Manager to review the customer-managed controls from all Assessments for which they are users. When a user signs in to Compliance Manager and opens the **Action Items** dashboard, a list of Action Items assigned to them is displayed. Depending on the Compliance Manager role assigned to the user, they can provide implementation or test details, update the Status, or assign Action Items.

As certification controls are generally implemented by one person and tested by another, the control action item can be initially assigned to one person for implementation, and once that is complete, that person can reassign the control action item to the next person for control testing and uploading of evidence. This assignment/reassignment of control actions can be performed by any users who have a Compliance Manager role with sufficient permissions, allowing for central management of control assignments, or decentralized routing of control action items, from implementer to tester as appropriate.

To assign an action item:

1. On the Compliance Manager dashboard, locate the assessment tile of the assessment you wish to work with and click on the name of the assessment to go to the assessment details page.

2. You can click **Filter** and use the filter controls to find the specific assessment control you wish to assign, or

3. Scroll down to the Customer-Managed Controls section, expand the control family, and scroll through the list of control until you have located the assessment control to be assigned

4. Under the **Assigned User** column, click **Assign**.

5. In the Assign Action Item dialog box, click the **Assign To** field to populate the list of users to whom the

action can be assigned. You can scroll through the list to find the target user or start typing in the field to search for the username.

6. Click the user to assign them this action item.
7. If you wish to send an email notification to the user notifying them, ensure that the **Send Email Notification** checkbox is checked.
8. Type any notes you wish to be displayed to that user and click **Assign**.

   The user will receive notification of their action item assignment and any notes you have provided.

The notes that are associated with the action item are persisted in the notes section, available for the next time the action item is assigned. These notes are not read-only, can be edited, replaced or removed by the person assigning the action item.

Exporting information from an Assessment

You can export an Assessment to an Excel file, which can be reviewed by compliance stakeholders in your organization, and provided to auditors and regulators. This assessment report is a snapshot of the assessment as of the date and time that the report is created, and it contains the details of both the Microsoft-managed controls and the customer-managed controls for that assessment, including control implementation status, control test date and test results, and provides links to the uploaded evidence documents. It is recommended that you export the assessment report prior to archiving an assessment, as archived assessments do not retain their links to uploaded documents.

To export an Assessment report:

- On the Compliance Manager dashboard, click **Actions** on the tile of the assessment you wish to export, and then choose **Export to Excel**

  Or

- If you are viewing the Assessment details page, click on the **Export to Excel** button, which is located in the upper right-hand corner of the page above the assessment's Compliance Score.

The assessment report will be downloaded in your browser session. If you don't see a popup informing you of this, you may wish to check your browser's downloads folder.

Archiving an Assessment

When you have completed an Assessment and no longer need it for compliance purposes, you can archive it. When an Assessment is archived, it is removed from Assessments dashboard.

> **NOTE**
> When an Assessment is Archived, it cannot be 'unarchived' or restored to a read-write in progress state. Please note that Archived Assessments do not retain their links to uploaded evidence documents, so it is highly recommended that you perform an Export of the Assessment before archiving it, as the exported assessment report will contain links to the evidence documents, enabling you to continue to access them.

To archive an assessment:

1. On the dashboard tile of the desired assessment, click **Actions**.
2. Select **Archive Assessment**.

   The **Archive Assessments** dialog is displayed, asking you to confirm that you want to archive the

assessment.

3. To continue with archiving, click **Archive**, or else click **Cancel**.

To view archived Assessments:

1. On the Compliance Manager dashboard, check the **Show Archived** checkbox.

   The archived assessments will appear in a newly visible section below the rest of the active assessments under a bar titled **Archived Assessments**.

2. Click the name of the assessment you wish to view.

When viewing an archived assessment, none of the normally editable controls (i.e. Implementation, Test Results) will be active, and the **Managed Documents** button will be absent.

Using search

Click the magnifying glass in the upper right-hand corner of the page by to expand the Search input field, enter your search terms and press Enter. The Search control will appear, with the search term in the search pane input field, and search results will appear beneath.

By default, Search returns Document results, and you can use the Filter By dropdown lists to refine the list of documents displayed, to add or remove search results from view. You can use multiple filter attributes at the same time to narrow the returned documents to specific cloud services, categories of compliance or security practices, regions of the world, or industries. Click the document name link to download the document.

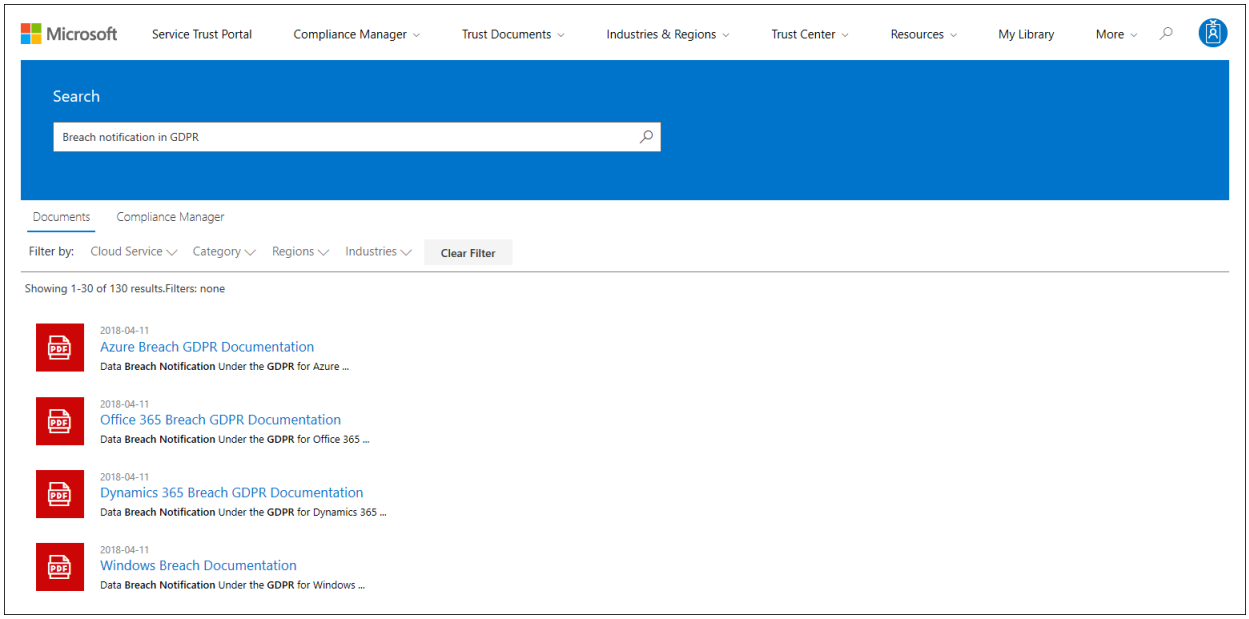

Click on the Compliance Manager link to display Search results for Compliance Manager assessment controls. The listed search results show the date the assessment was created, the name of the assessment grouping, the applicable cloud service, and whether the controls are Microsoft or Customer Managed.

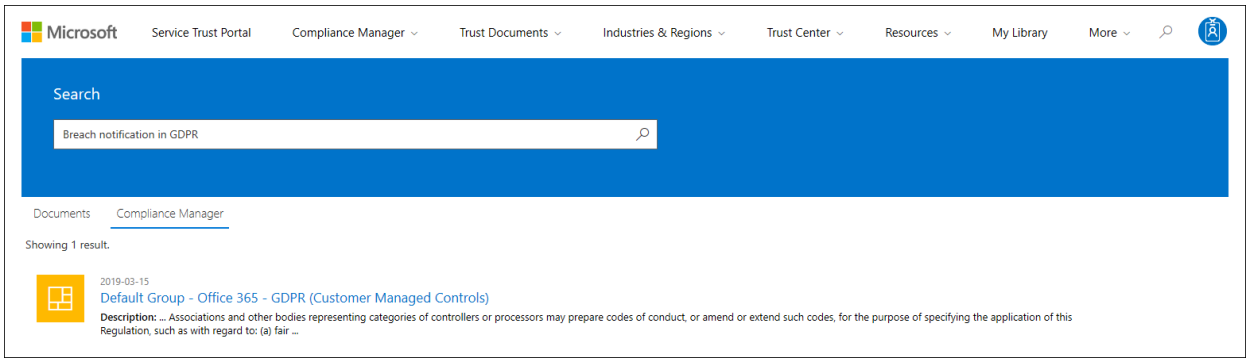

NOTE

Service Trust Portal reports and documents are available to download for at least twelve months after publishing or until a new version of document becomes available.

Localization support

Service Trust Portal enables you to view the page content in different languages. To change the page language, simply click on the globe icon in the lower left corner of the page and select the language of your choice.

Change log for Customer-Managed Controls

Compliance Manager is designed to be regularly updated to keep pace with changes in regulatory requirements, as well as changes in our cloud services. These updates include changes to the Customer-Managed Controls. A Change Log is provided to help you understand the impact of these changes, including the details of the content being added or changed, and guidance as to what effect the changes have on existing Assessments. Generally, there are two types of changes:

- A **Major** change is a significant change to a Customer Action, such as the addition or removal of a control or specific numbered steps, or a change in the guidance around responsibilities, recommendations, or evidence. For Major changes, we recommend that you re-evaluate your implementation and/or assessment of the affected control.
- A **Minor** change is an insignificant change to a Customer Actions, such as fixing a typo or formatting issues, or updating or correcting hyperlinks. Minor changes generally do not require the control to be re-evaluated; however, we do recommend that you review the updated Customer Action.

Customer-managed controls - Change Log for July 2018

CONTROL ID	ASSESSMENT	TYPE OF CHANGE	DESCRIPTION OF CHANGE	RECOMMENDED ACTIONS FOR CUSTOMERS
45 C.F.R. § 164.308(a)(7)(ii)(A)	Office 365: HIPAA	Major	Added HITECH control to HIPAA Assessment for Office 365	Review the added control and recommended Customer Actions
45 C.F.R. 164.312(a)(6)(ii)	Office 365: HIPAA	Major	Added HITECH control to HIPAA Assessment for Office 365	Review the added control and recommended Customer Actions
45 C.F.R. § 164.312(c)(1)	Office 365: HIPAA	Major	Added HITECH control to HIPAA Assessment for Office 365	Review the added control and recommended Customer Actions
45 C.F.R. § 164.316(b)(2)(iii)	Office 365: HIPAA	Major	Added HITECH control to HIPAA Assessment for Office 365	Review the added control and recommended Customer Actions

Customer-managed controls - Change Log for April 2018

GDPR	HIPAA	ISO 27001	ISO 27018	NIST 800-53	NIST 800-171	TYPE OF CHANGE	DESCRIPTION OF CHANGE	RECOMMENDED ACTIONS FOR CUSTOMERS
6.13.2			C.16.1.1			Major	Previously numbered as 6.12.1.1. Added details to recommendations.	Re-assess the control: Review the updated guidance in the Customer Actions and follow the recommended steps for implementing and assessing the control.

GDPR	HIPAA	ISO 27001	ISO 27018	NIST 800-53	NIST 800-171	TYPE OF CHANGE	DESCRIPTION OF CHANGE	RECOMMENDED ACTIONS FOR CUSTOMERS
					3.1.6	Major	Added steps to guidance that include enabling auditing and searching audit logs.	Review the updated recommendations in the Customer Actions.
6.8.2			A.10.2			Major	Previously numbered as 6.7.2.9. Updated guidance with additional recommendations and action items.	Re-assess the control: Review the updated guidance in the Customer Actions and follow the recommended steps for implementing and assessing the control.
6.6.4	45 C.F.R. § 164.312(a)(2)(i) 45 C.F.R. § 164.312(d)	A.9.4.2		IA-2	3.5.1	Major	Previously numbered as 6.5.2.3. Updated guidance with additional recommendations and action items.	Re-assess the control: Review the updated guidance in the Customer Actions and follow the recommended steps for implementing and assessing the control.

GDPR	HIPAA	ISO 27001	ISO 27018	NIST 800-53	NIST 800-171	TYPE OF CHANGE	DESCRIPTION OF CHANGE	RECOMMENDED ACTIONS FOR CUSTOMERS
6.13.1	45 C.F.R. § 164.308(a)(1)(i)	A.16.1	C.16.1	IR-4(a)	3.6.1	Major	Previously numbered as 6.12.1. Updated guidance with additional recommendations and action items.	Re-assess the control: Review the updated guidance in the Customer Actions and follow the recommended steps for implementing and assessing the control.
6.7						Major	Previously numbered as 6.6.1.1. Updated guidance with additional recommendations and action items.	Re-assess the control: Review the updated guidance in the Customer Actions and follow the recommended steps for implementing and assessing the control.

GDPR	HIPAA	ISO 27001	ISO 27018	NIST 800-53	NIST 800-171	TYPE OF CHANGE	DESCRIPTION OF CHANGE	RECOMMENDED ACTIONS FOR CUSTOMERS
6.6.5			A.10.8	IA-3	3.5.2	Major	Previously numbered as 6.5.4.2. Updated guidance with additional recommendations and action items.	Re-assess the control: Review the updated guidance in the Customer Actions and follow the recommended steps for implementing and assessing the control.
6.15.1						Major	Previously numbered as 6.14.1.3. Updated guidance with additional recommendations and action items.	Re-assess the control: Review the updated guidance in the Customer Actions and follow the recommended steps for implementing and assessing the control.
				AC-2(h)(2)		Minor	Added link to Enable Auditing blade.	No action necessary.
				AC-2(7)(b)		Minor	Added link to Enable Auditing blade.	No action necessary.

GDPR	HIPAA	ISO 27001	ISO 27018	NIST 800-53	NIST 800-171	TYPE OF CHANGE	DESCRIPTION OF CHANGE	RECOMMENDED ACTIONS FOR CUSTOMERS
				AC-2(h)(1)		Minor	Added link to Enable Auditing blade.	No action necessary.
	45 C.F.R. § 164.308(a)(5)(ii)(C)			AC-2(g)		Minor	Added link to Enable Auditing blade.	No action necessary.
				AC-2(12)		Minor	Added link to Enable Auditing blade.	No action necessary.
	45 C.F.R. § 164.312(b)	A.12.4.3		AU-2(d)		Minor	Added link to Enable Auditing blade.	No action necessary.
				AC-2(4)		Minor	Added link to Enable Auditing blade.	No action necessary.
					3.1.7	Minor	Added link to Enable Auditing blade.	No action necessary.
		A.16.1.7	C.12.4.2, Part 2			Minor	Added link to Enable Auditing blade.	No action necessary.
				AC-2(h)(3)		Minor	Added link to Enable Auditing blade.	No action necessary.
		A.12.4.2				Minor	Added link to Enable Auditing blade.	No action necessary.

GDPR	HIPAA	ISO 27001	ISO 27018	NIST 800-53	NIST 800-171	TYPE OF CHANGE	DESCRIPTION OF CHANGE	RECOMMENDED ACTIONS FOR CUSTOMERS
		A.7.2.8				Minor	Added links to Content Search blade and to DSR portal.	No action necessary.
	45 C.F.R. § 164.308(a)(3)(ii)(C)					Minor	Added links to Enable Auditing blade and to Office 365 admin role support topics.	No action necessary.
5.2.1						Minor	Previously numbered as 5.2.2. Clarified customer responsibilities within guidance.	Review the updated recommendations in the Customer Actions.
6.11.1	45 C.F.R. § 164.312(e)(2)(ii)	A.10.1.1 A.10.1.2 A.18.1.5	C.10.1.1	SC-13	3.13.11	Minor	Previously numbered as 6.10.1.2. Fixed typo.	No action necessary.
7.5.1						Minor	Previously numbered as A.7.4.1. Fixed typo.	No action necessary.

GDPR	HIPAA	ISO 27001	ISO 27018	NIST 800-53	NIST 800-171	TYPE OF CHANGE	DESCRIPTION OF CHANGE	RECOMMENDED ACTIONS FOR CUSTOMERS
		A.8.2.3			3.1.3	Minor	Removed extra unnecessary sentence.	No action necessary.
	45 C.F.R. § 164.308(a)(4)(i)	A.6.1.2		AC-5(a)	3.1.2 3.1.4	Minor	Updated guidance with additional recommendations and action items.	Review the updated recommendations in the Customer Actions.
	45 C.F.R. § 164.308(a)(7)(ii)(E)			RA-2(a)		Minor	Updated import service help topic link to use FWLink.	No action necessary.

GDPR Assessment Control ID Change Reference - Change Log for February 2018

PREVIOUS CONTROL ID (NOVEMBER 2017 PREVIEW)	NEW CONTROL ID (FEBRUARY 2018 GA RELEASE)
5.2.2	5.2.1
5.2.3	5.2.2
5.2.4	5.2.3
6.1.1.1	6.2
6.10.1.2	6.11.1
6.10.2.5	6.11.2
6.11.1.2	6.12
6.12.1	6.13.1
6.12.1.1	6.13.2
6.12.1.5	6.13.3
6.14.1.3	6.15.1

PREVIOUS CONTROL ID (NOVEMBER 2017 PREVIEW)	NEW CONTROL ID (FEBRUARY 2018 GA RELEASE)
6.14.2.1	6.15.2
6.14.2.3	6.15.3
6.2.1.1	6.3
6.3.2.2	6.4
6.4.3.1	6.5.2
6.4.3.2	6.8.1
6.4.3.3	6.5.3
6.5.2	6.6.1
6.5.2.1	6.6.2
6.5.2.2	6.6.3
6.5.2.3	6.6.4
6.5.4.2	6.6.5
6.6.1.1	6.7
6.7.2.7	6.8.1
6.7.2.9	6.8.2
6.8.1.4	6.9.1
6.8.4.1	6.9.3
6.8.4.2	6.9.4
6.9.2.1	6.10.1
6.9.2.3	6.10.2
A.7.1.1	7.2.1
A.7.1.2	7.2.2
A.7.1.3	7.2.3
A.7.1.4	7.2.4
A.7.1.5	7.2.5

PREVIOUS CONTROL ID (NOVEMBER 2017 PREVIEW)	NEW CONTROL ID (FEBRUARY 2018 GA RELEASE)
A.7.1.6	7.2.6
A.7.1.7	7.2.7
A.7.2.1	7.3.1
A.7.2.10	7.3.9
A.7.2.11	7.3.10
A.7.2.2	7.3.2
A.7.2.3	7.3.3
A.7.2.4	7.3.4
A.7.2.5	7.3.5
A.7.2.6	7.3.6
A.7.2.7	7.3.7
A.7.2.8	7.3.8
A.7.3.1	7.4.1
A.7.3.10	7.4.10
A.7.3.2	7.4.2
A.7.3.3	7.4.3
A.7.3.4	7.4.4
A.7.3.5	7.4.5
A.7.3.6	7.4.6
A.7.3.7	7.4.7
A.7.3.8	7.4.8
A.7.3.9	7.4.9
A.7.4.1	7.5.1
A.7.4.2	7.5.2
A.7.4.3	7.5.3

PREVIOUS CONTROL ID (NOVEMBER 2017 PREVIEW)	NEW CONTROL ID (FEBRUARY 2018 GA RELEASE)
A.7.4.4	7.5.4
A.7.4.5	7.5.5
B.8.1.1	8.2.1
B.8.1.2	8.2.2
B.8.1.3	8.2.3
B.8.1.4	8.2.4
B.8.1.5	8.2.5
B.8.1.6	8.2.6
B.8.2.1	8.3.1
B.8.3.1	8.4.1
B.8.3.2	8.4.2
B.8.3.3	8.4.3
B.8.4.1	8.5.1
B.8.4.2	8.5.2
B.8.4.3	8.5.4
B.8.4.4	8.5.5
B.8.4.5	8.5.3
B.8.4.6	8.5.6
B.8.4.7	8.5.7
B.8.4.8	8.5.8

Get started with the Microsoft Service Trust Portal

11/2/2020 • 7 minutes to read • Edit Online

The Microsoft Service Trust Portal provides a variety of content, tools, and other resources about Microsoft security, privacy, and compliance practices.

Accessing the Service Trust Portal

The Service Trust Portal contains details about Microsoft's implementation of controls and processes that protect our cloud services and the customer data therein. To access some of the resources on the Service Trust Portal, you must log in as an authenticated user with your Microsoft cloud services account (either an Azure Active Directory organization account or a Microsoft Account) and review and accept the Microsoft Non-Disclosure Agreement for Compliance Materials.

Existing customers

Existing customers can access the Service Trust Portal at https://aka.ms/STP with one of the following online subscriptions (trial or paid):

- Microsoft 365
- Dynamics 365
- Azure

> **NOTE**
> Azure Active Directory accounts associated with organizations have access to the full range of documents and features like Compliance Manager. Microsoft accounts created for personal use have limited access to Service Trust Portal content.

New customers and customers evaluating Microsoft online services

To create a new account or to create a trial account, use one of the following sign-up forms (also used for trial accounts) to get access to the STP.

- Sign up for a new Microsoft 365 Apps for business trial account or a new Office 365 Enterprise trial account
- Sign up for a new Dynamics 365 trial account
- Sign up for a new Azure trial account.

When you sign up for either a free trial, or a subscription, you must enable Azure Active Directory to support your access to the STP.

Using the Service Trust Portal

The Service Trust Portal features and content are accessible from the main menu.

The following sections describe each item in the main menu.

Service Trust Portal

The **Service Trust Portal** link displays the home page. It provides a quick way to get back to the home page.

Compliance Manager

> **IMPORTANT**
> Compliance Manager has moved from the Service Trust Portal to its new location in the Microsoft 365 compliance center. All customer data has been moved over to the new location, so you can continue using Compliance Manager without interruption. Refer to the Compliance Manager documentation for setup information and to learn about new features. Although the classic version of Compliance Manager remains in the Service Trust Portal, all users are encouraged to use Compliance Manager in the Microsoft 365 compliance center.

Trust Documents

Provides a wealth of security implementation and design information with the goal of making it easier for you to meet regulatory compliance objectives by understanding how Microsoft Cloud services keep your data secure. To review content, select one of the following options on the **Trust Documents** pull-down menu.

- **Audit Reports:** A list of independent audit and assessment reports on Microsoft's Cloud services is displayed. These reports provide information about Microsoft Cloud services compliance with data protection standards and regulatory requirements, such as:
  - International Organization for Standardization (ISO)
  - Service Organization Controls (SOC)
  - National Institute of Standards and Technology (NIST)
  - Federal Risk and Authorization Management Program (FedRAMP)
  - General Data Protection Regulation (GDPR)
- **Data Protection:** Contains a wealth of resources such as audited controls, white papers, FAQs, penetration tests, risk assessment tools, and compliance guides.
- **Azure Security and Compliance Blueprints:** Resources that help you build secure and compliant cloud-based applications. This area contain blueprint-guidance for government, finance, healthcare, and retail verticals.

Industries & Regions

Provides industry- and region-specific compliance information about Microsoft Cloud services.

- **Industries:** At this time, this page provides an industry-specific landing page the for the Financial Services industry. This contains information such as compliance offerings, FAQs, and success stories. Resources for more industries will be released in the future, however you can find resources for more industries by going to the **Trust Documents > Data Protection** page in the STP.
- **Regions:** Provides legal opinions on Microsoft Cloud services compliance with various the laws of various countries. Specific countries include Australia, Canada, Czech Republic, Denmark, Germany, Poland, Romania, Spain, and the United Kingdom.

Trust Center

Links to the Microsoft Trust Center, which provides more information about security, compliance, and privacy in the Microsoft Cloud. This includes information about the capabilities in Microsoft Cloud services that you can use to address specific requirements of the GDPR, documentation helpful to your GDPR accountability and to your understanding of the technical and organizational measures Microsoft has taken to support the GDPR.

My Library

This new feature lets you save (or *pin*) documents so that you can quickly access them on your My Library page. You can also set up notifications so that Microsoft sends you an email message when documents in your My Library are updated. For more information, see the My Library section in this article.

More

Go to **More** > **Admin** to access administrative functions that are only available to the global administrator account. This option is visible only when you are signed in as a global administrator. There are two options in the **Admin** pull-down menu:

- **Settings:** This page lets you assign user roles for Compliance Manager (classic).
- **User Privacy Settings:** This page lets you export a report that contains action item assignments in Compliance Manager (classic) for a specific user. You can also reassign all action items to a different user and remove any assigned action item from the specified user.

Search

Click the magnifying glass in the upper right-hand corner of the Service Trust Portal page to expand the box, enter your search terms, and press **Enter**. The **Search** page is displayed, with the search term displayed in the search box and the search results listed below.

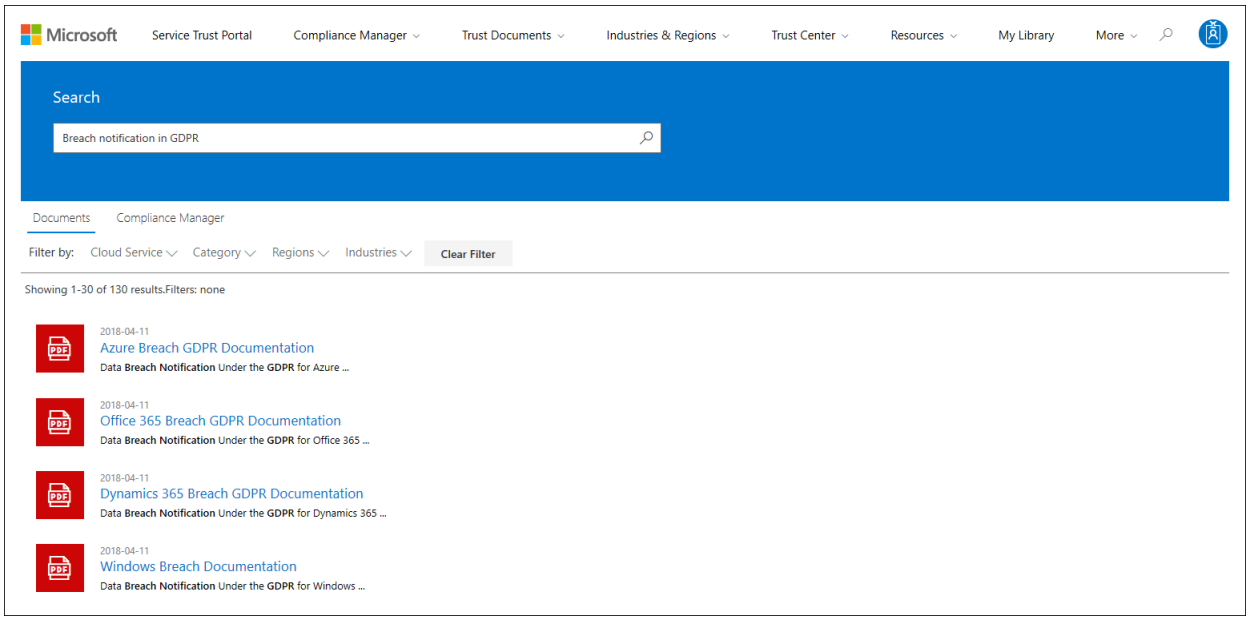

By default, the search returns document results. You can filter the results by using the dropdown lists to refine the list of documents displayed. You can use multiple filters to narrow the list of documents. Filters include the specific cloud services, categories of compliance or security practices, regions, and industries. Click the document name link to download the document.

To list controls from Assessments in Compliance Manager (classic) related your search terms, click **Compliance Manager**. The search results show the date the assessment was created, the name of the assessment grouping, the applicable Microsoft Cloud service, and whether the control is Microsoft or Customer Managed. Click the name of the control to view the control in the Assessment in Compliance Manager (classic).

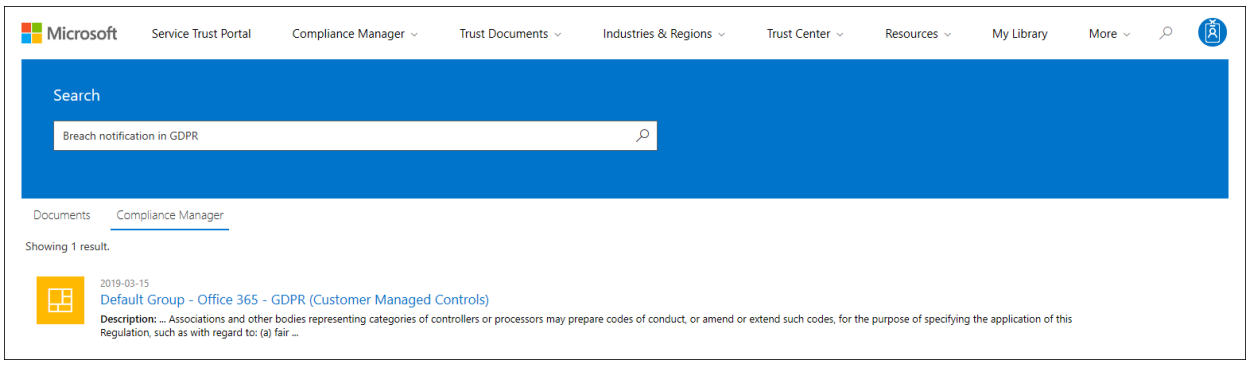

> **NOTE**
>
> Service Trust Portal reports and documents are available to download for at least 12 months after publishing or until a new version of document becomes available.

My Library

Use the My Library feature to add documents and resources on the Service Trust Portal to your My Library page. This lets you access documents that are relevant to you in a single place. To add a document to your My Library, click the ... menu to the right of a document and then select **Save to library**. You can add multiple documents to your My Library by clicking the checkbox next to one or more documents, and then clicking **Save to library** at the top of the page.

Additionally, the notifications feature lets you configure your My Library so that an email message is sent to you whenever Microsoft updates a document that you've added to your My Library. To set up notifications, go to your My Library and click **Notification Settings**. You can choose the frequency of notifications and specify an email address in your organization to send notifications to. Email notifications include links to the documents that have been updated and a brief description of the update.

Also note that we identify any documents in your My Library that have been updated within the last 30 days, regardless of whether or not you turn on notifications. A brief description of the update is also displayed in a tool tip.

Starter packs

Starter packs are a Microsoft-curated set of documentation about Microsoft Cloud services for specific industries. Currently, the Service Trust Portal offers the following three starter packs for financial services organizations. These starter packs help organizations evaluate and assess security, compliance, and privacy in the Microsoft Cloud and provide guidance to help implement Microsoft Cloud services in the highly regulated financial services industry.

- **Evaluation Starter Pack:** Use for early evaluation of the Microsoft cloud for financial services organizations.
- **Assessment Starter Pack:** After evaluation, use the checklists and other guidance in this starter pack to help your organization assess risks related to security, compliance, and privacy.
- **Audit Starter Pack:** User this starter pack for guidance on using auditing controls and other tool to help guide your implementation of Microsoft Cloud services in a way that helps reduce your organization's exposure to risk.

To access these starter packs, go to **Service Trust Portal > Industries & Regions > Industry Solutions > Financial Services**. You can open or a download documents from a starter pack or save them to your My Library.

Localization support

The Service Trust Portal enables you to view the page content in different languages. To change the page language, simply click on the globe icon in the lower left corner of the page and select the language of your choice.

Give feedback

We can help with questions about the Service Trust Portal, or errors you experience when you use the portal. You can also contact us with questions and feedback about Service Trust Portal compliance reports and trust resources by using the Feedback link on the bottom of the STP pages.

Your feedback is important to us. Click on the Feedback button at the bottom of the page to send us comments about what you did or did not like, or suggestions you may have for improving our products or product features.

Key compliance and security considerations for US banking and capital markets

2/18/2021 • 37 minutes to read • Edit Online

Introduction

Financial services institutions surpass nearly all commercial businesses in their demand for stringent security, compliance, and governance controls. The protection of data, identities, devices, and applications is not only critical to their business, it's subject to compliance requirements and guidelines from regulatory bodies such as the U.S. Securities and Exchange Commission (SEC), the Financial Industry Regulatory Authority (FINRA), the Federal Financial Institutions Examination Council (FFIEC), and the Commodity Futures Trading Commission (CFTC). In addition, financial institutions are subject to laws such as Dodd-Frank and the Sarbanes-Oxley Act of 2002.

In today's climate of increased security vigilance, insider risk concerns and public data breaches, customers also demand high levels of security from their financial institutions in order to trust them with their personal data and banking assets.

Historically, the need for comprehensive controls directly impacted and constrained the IT systems and platforms that financial institutions use to enable collaboration internally and externally. Today, financial services employees need a modern collaboration platform that's easy to adopt and easy to use. But financial services can't trade the flexibility to collaborate between users, teams, and departments with security and compliance controls that enforce policies to protect users and IT systems from threats.

In the financial services sector, careful consideration is required for the configuration and deployment of collaboration tools and security controls, including:

- Risk assessment of common organizational collaboration and business process scenarios
- Information protection and data governance requirements
- Cybersecurity and insider threats
- Regulatory compliance requirements
- Other operational risks

Microsoft 365 is a modern workplace cloud environment that can address the contemporary challenges financial services organizations face. Secure and flexible collaboration across the enterprise is combined with controls and policy enforcement to adhere to stringent regulatory compliance frameworks. This article describes how the Microsoft 365 platform helps financial services move to a modern collaboration platform, while helping keep data and systems secure and compliant with regulations:

- Enable organizational and employee productivity by using Microsoft 365 and Microsoft Teams
- Protect modern collaboration by using Microsoft 365
- Identify sensitive data and prevent data loss
- Defend the fortress
- Govern data and comply with regulations by effectively managing records
- Establish ethical walls with information barriers
- Protect against data exfiltration and insider risk

As a Microsoft partner, Protiviti contributed to and provided material feedback to this article.

The following downloadable illustrations supplement this article. Woodgrove Bank and Contoso are used to demonstrate how capabilities described in this article can be applied to address common regulatory requirements of financial services. Feel free to adapt these illustrations for your own use.

Microsoft 365 information protection and compliance illustrations

ITEM	DESCRIPTION
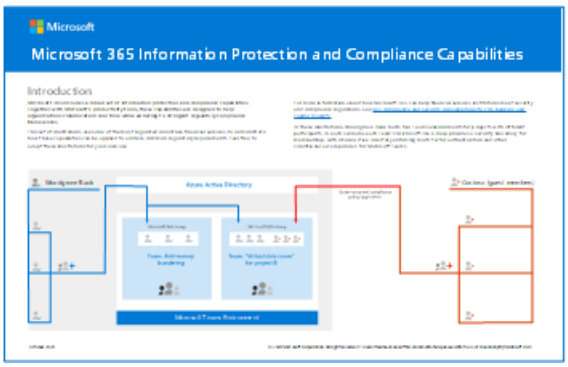 English: Download as a PDF | Download as a Visio Japanese: Download as a PDF | Download as a Visio Updated November 2020	Includes: • Microsoft information protection and data loss prevention • Retention policies and retention labels • Information barriers • Communication compliance • Insider risk • Third-party data ingestion

Empower organizational and employee productivity by using Microsoft 365 and Teams

Collaboration typically requires various forms of communication, the ability to store and access documents/data, and the ability to integrate other applications as needed. Employees in financial services typically need to collaborate and communicate with members of other departments or teams and sometimes with external entities. Therefore, using systems that create silos or make information sharing difficult is undesirable. Instead, it's preferable to leverage platforms and applications that enable employees to communicate, collaborate, and share information securely and according to corporate policy.

Providing employees with a modern, cloud-based collaboration platform allows them to choose and integrate tools that make them more productive and empower them to find agile ways to work. Using Teams in conjunction with security controls and information governance policies that protect the organization can help your workforce communicate and collaborate effectively.

Teams provides a collaboration hub for the organization. It helps bring people together to work productively on common initiatives and projects. Teams lets team members conduct 1:1 and multi-party chat conversations, collaborate and coauthor documents, and store and share files. Teams also facilitates online meetings through integrated enterprise voice and video. Teams can also be customized with Microsoft apps such as Microsoft Planner, Microsoft Dynamics 365, PowerApps, Power BI, and third-party line-of-business applications. Teams is designed for use by both internal team members and permitted external users who can join team channels, participate in chat conversations, access stored files, and leverage other applications

Every Microsoft Team is backed by a Microsoft 365 group. That group is considered the membership service for numerous Office 365 services, including Teams. Microsoft 365 groups are used to securely distinguish between "owners" and "members" and to control access to various capabilities within Teams. When coupled with appropriate governance controls and regularly administered access reviews, Teams allows only members and owners to utilize authorized channels and capabilities.

A common scenario where Teams benefits financial services is when running internal projects or programs. For example, many financial institutions, including banks, wealth management firms, credit unions, and insurance providers, are required to have anti-money laundering and other compliance programs in place. A cross-functional team comprised of IT, lines of business such as retail and wealth management, and a financial crime

unit may be required to share data with each other and communicate about the program or specific investigations. Traditionally, these programs have used shared network drives, but this approach can present numerous challenges, including:

- Only one person can edit a document at a time.
- Managing security is time-consuming because adding/removing individuals typically involves IT.
- Data remains resident on shared network drives much longer than required or wanted.

Teams can provide a collaboration space to securely store sensitive client data and conduct conversations between team members where sensitive topics may be discussed. Multiple members of the team can edit or collaborate on a single document at the same time. The program owner or coordinator can be configured as the team owner and can then add and remove members as needed.

Another common scenario is to use Teams as a "virtual data room" to securely collaborate, including storing and managing documents. Team members and syndicates within investment banking, asset management, or private equity firms can securely collaborate on a deal or investment. Cross-functional teams are often involved in planning and fulfilling such deals, and the ability to securely share data and conduct conversations is a core requirement. Securely sharing related documents with external investors is also a key requirement. Teams provides a secure and fully auditable location from which to centrally store, protect, and share investment data.

Teams: Improve collaboration and reduce compliance risk

Microsoft 365 provides other common policy capabilities for Teams through its use of Microsoft 365 groups as an underlying membership service. These policies can help improve collaboration and meet compliance needs.

Microsoft 365 group naming policies help ensure that Microsoft 365 groups, and therefore teams, are named according to corporate policy. Names can be problematic if they aren't appropriate. For example, employees may not know which teams to work with or share information with if names aren't appropriately applied. Group-naming policies (including support for prefix/suffix-based policies and custom blocked words) can enforce good "hygiene" and prevent use of specific words, such as reserved words or inappropriate terminology.

Microsoft 365 group expiration policies help ensure that Microsoft 365 groups and therefore teams, aren't retained for longer periods of time than the organization wants or needs. This capability helps prevent two key

information-management issues:

- Proliferation of teams that aren't necessary or used.
- Over-retention of data that's no longer required or used by the organization (except in cases of legal hold/preservation).

Administrators can specify an expiration period for Microsoft 365 groups, such as 90, 180, or 365 days. If a service that's backed by a Microsoft 365 group is inactive for the expiration period, group owners are notified. If no action is taken, the Microsoft 365 group and all its related services, including Teams, is deleted.

The over-retention of data that's stored in Teams and other group-based services can pose risks to financial services organizations. Microsoft 365 group expiration policies are a recommended way to help prevent retention of data that's no longer needed. Combined with built-in retention labels and policies, Microsoft 365 helps ensure that organizations are only retaining the data that's required to meet corporate policies and regulatory compliance obligations.

Teams: Integrate custom requirements with ease

Teams enables self-service creation of teams by default. However, many regulated organizations want to control and understand which collaboration channels are currently in use by their employees, which channels may contain sensitive data, and ownership of organizational channels. To facilitate these governance controls, Microsoft 365 lets organization disable self-service teams creation. By using business process automation tools such as Microsoft Power Apps and Power Automate, organizations can build and deploy simple forms and approval processes for employees to request creation of a new team. When approved, the team can be automatically provisioned and a link sent to the requestor. In this way, organizations can design and integrate their compliance controls and custom requirements into the team-creation process.

Acceptable digital communication channels

FINRA emphasizes that the digital communications of regulated firms meet the record-keeping requirements of Exchange Act rules 17a-3 and 17a-4, as well as FINRA Rule Series 4510. FINRA releases an annual report that contains key findings, observations, and effective practices to help organizations improve compliance and risk management. In its 2019 Report on Examination Findings and Observations, FINRA identified digital communications as a key area where firms encounter challenges complying with supervision and record-keeping requirements.

If an organization permits its employees to use a specific application, such as an app-based messaging service or collaboration platform, the firm must archive business records and supervise the activities and communications of those employees in that application. Organizations are responsible for conducting due diligence to comply with FINRA rules and securities laws, and for following up on potential violations of those rules related to employee use of such apps.

Effective practices recommended by FINRA include the following:

- Establish a comprehensive governance program for digital communication channels. Manage the organization's decisions about which digital communication channels are permitted and define compliance processes for each digital channel. Closely monitor the rapidly changing landscape of digital communication channels and keep compliance processes up to date.
- Clearly define and control permissible digital channels. Define both approved and prohibited digital channels. Block or restrict the use of prohibited digital channels, or prohibited features within digital channels, that limit the organization's ability to comply with records management and supervisory requirements.
- Provide training for digital communications. Implement mandatory training programs before giving registered representatives access to approved digital channels. Training helps clarify an organization's expectations for business and personal digital communications, and it guides staff through using permitted features of each channel in a compliant manner.

FINRA's findings and observations for Digital Communications relate directly to an organization's ability to

comply with SEC Rule 17a-4 for retaining all business-related communications, FINRA rules 3110 and 3120 for supervision and review of communications, and Rule Series 4510 for record keeping. The Commodity Futures Trading Commission (CFTC) promulgates similar requirements under 17 CFR 131. These regulations are discussed in depth later in this article.

Teams, along with the comprehensive suite of Microsoft 365 security and compliance offerings, provides a corporate digital communication channel for financial services institutions to effectively conduct business and comply with regulations. The remainder of this article describes how Microsoft 365 built-in capabilities for records management, information protection, information barriers, and supervisory control gives Teams a robust toolset to help meet these regulatory obligations.

Protect modern collaboration with Microsoft 365

Secure user identities and control access

Protecting access to customer information, financial documents, and applications begins with strongly securing user identities. This requires a secure platform for the enterprise to store and manage identities, providing a trusted means of authentication, and dynamically controlling access to those applications.

As employees work, they may move from application to application or between multiple locations and devices. Access to data must be authenticated at each step along the way. The authentication process has to support a strong protocol and multiple factors of authentication (such as one-time SMS pass code, authenticator app, and certificate) to ensure that identities aren't compromised. Enforcing risk-based access policies is critical to protecting financial data and applications from insider threats, inadvertent data leaks, and data exfiltration.

Microsoft 365 provides a secure identity platform in Azure Active Directory (Azure AD), where identities are centrally stored and securely managed. Azure AD, along with a host of related Microsoft 365 security services, forms the basis for providing employees with the access they need to work securely while also protecting the organization from threats.

Azure AD Multi-Factor Authentication (MFA) is built into the platform and provides an additional proof of authentication to help confirm user identity when they access sensitive financial data and applications. Azure MFA requires at least two forms of authentication, such as a password plus a known mobile device. It supports several second-factor authentication options, including:

- The Microsoft Authenticator app
- A one-time passcode delivered via SMS
- A phone call where a user must enter a PIN

If password is somehow compromised, a potential hacker would still need the user's phone to gain access to organizational data. In addition, Microsoft 365 uses Modern Authentication as a key protocol, which brings the same strong and rich authentication experience from web browsers to the collaboration tools that employees use day to day, including Microsoft Outlook and the other Microsoft Office applications.

Passwordless

Passwords are the weakest link in a security chain. They can be a single point of failure if there's no additional verification. Microsoft supports a broad range of authentication options to fit the needs of financial institutions.

Passwordless methods help make MFA more convenient for users. While not all MFA is passwordless, passwordless technologies employ multi-factor authentication. Microsoft, Google, and other industry leaders have developed standards to enable a simpler, stronger authentication experience across the web and mobile devices in a group called Fast IDentity Online (FIDO). The recently developed FIDO2 standard enables users to authenticate easily and securely without requiring a password to eliminate phishing.

Microsoft MFA methods that are passwordless include:

- Microsoft Authenticator: For flexibility, convenience, and cost, we recommend using the Microsoft

Authenticator mobile app. Microsoft Authenticator supports biometrics, push notifications, and one-time passcodes for any Azure AD-connected app. It's available from the Apple and Android app stores.

- Windows Hello: For a built-in experience on the PC, we recommend using Windows Hello. It uses biometric information (such as face or fingerprint) to sign in automatically.
- FIDO2 Security keys are now available from several Microsoft partners: Yubico, Feitian Technologies, and HID Global in a USB, NFC-enabled badge or biometric key.

Azure AD Conditional Access provides a robust solution for automating access control decisions and enforcing organizational policies to protect company assets. A classic example is when a financial planner wants to access an application that has sensitive customer data. They are automatically required to perform a multi-factor authentication to specifically access that application, and access must be from a corporate-managed device. Azure Conditional Access brings together signals about a user's access request, such as properties about the user, the device, location, and network, and the application that the user is trying to access. It dynamically evaluates attempts to access the application against configured policies. If user or device risk is elevated, or other conditions are not met, Azure AD can automatically enforce policies such as requiring MFA, requiring a secure password reset, or restricting or blocking access. This helps ensure that sensitive organizational assets are protected in dynamically changing environments.

Azure AD, and the related Microsoft 365 security services, provide the foundation on which a modern cloud collaboration platform can be rolled out to financial institutions so that access to data and applications can be secured, and regulator compliance obligations can be met. These tools provide the following key capabilities:

- Centrally store and securely manage user identities.
- Use a strong authentication protocol, including multi-factor authentication, to authenticate users on access requests and provide a consistent and robust authentication experience across all applications.
- Dynamically validate policies on all access requests, incorporating multiple signals into the policy decision-making process, including identity, user/group membership, application, device, network, location, and real-time risk score.
- Validate granular policies based on user behavior and file properties and dynamically enforce additional security measures when required.
- Identify "shadow IT" in the organization, and allow InfoSec teams to sanction or block cloud applications.
- Monitor and control access to Microsoft and non-Microsoft cloud applications.
- Proactively protect against email phishing and ransomware attacks.

Azure AD Identity Protection

While Conditional Access protects resources from suspicious requests, Identity Protection goes further by providing ongoing risk detection and remediation of suspicious user accounts. Identity Protection keeps you informed of suspicious user and sign-in behavior in your environment around the clock. Its automatic response proactively prevents compromised identities from being abused.

Identity Protection is a tool that allows organizations to accomplish three key tasks:

- Automate the detection and remediation of identity-based risks.
- Investigate risks by using data in the portal.
- Export risk detection data to third-party utilities for further analysis.

Identity Protection uses knowledge that Microsoft has acquired from its position in organizations with Azure AD, in the consumer space with Microsoft Accounts, and in gaming with Xbox to protect your users. Microsoft analyzes 65 trillion signals per day to identify and protect customers from threats. The signals generated by and fed to Identity Protection can be further fed into tools like Conditional Access to make access decisions. They can also be fed back to a security information and event management (SIEM) tool for further investigation based on your organization's enforced policies.

Identity Protection helps organizations automatically protect against identity compromise by taking advantage

of cloud intelligence powered by advanced detection based on heuristics, user and entity behavior analytics (UEBA), and machine learning (ML) across the Microsoft ecosystem.

Identify sensitive data and prevent data loss

Microsoft 365 allows all organizations to identify sensitive data within the organization through a combination of powerful capabilities, including:

- **Microsoft Information Protection (MIP)** for both user-based classification and automated classification of sensitive data.
- **Office 365 Data Loss Prevention (DLP)** for automated identification of sensitive data using sensitive data types (in other words, regular expressions) and keywords and policy enforcement.

Microsoft Information Protection (MIP) enables organizations to classify documents and emails intelligently by using sensitivity labels. Sensitivity labels can be applied manually by users to documents in Microsoft Office applications and to emails in Outlook. The labels can automatically apply document markings, protection through encryption, and rights-management enforcement. Sensitivity labels can also be applied automatically by configuring policies that use keywords and sensitive data types (such as credit card numbers, social insurance numbers, and identity numbers) to automatically find and classify sensitive data.

In addition, Microsoft provides "trainable classifiers" that use machine-learning models to identify sensitive data based on the content, as opposed to simply through pattern matching or by the elements within the content. A classifier learns how to identify a type of content by looking at numerous examples of the content to be classified. Training a classifier begins by giving it examples of content in a particular category. After it learns from those examples, the model is tested by giving it a mix of matching and non-matching examples. The classifier predicts whether a given example falls into the category or not. A person then confirms the results,

sorting the positives, negatives, false positives, and false negatives to help increase the accuracy of the classifier's predictions. When the trained classifier is published, it processes content in Microsoft SharePoint Online, Exchange Online, and OneDrive for Business and automatically classifies the content.

Applying sensitivity labels to documents and emails embeds metadata that identifies the chosen sensitivity within the object. The sensitivity then travels with the data. So even if a labeled document is stored on a user's desktop or within an on-premises system, it's still protected. This functionality enables other Microsoft 365 solutions, such as Microsoft Cloud App Security or network edge devices, to identify sensitive data and automatically enforce security controls. Sensitivity labels have the added benefit of educating employees about which data within an organization is considered sensitive and how to handle that data when they receive it.

Office 365 Data Loss Prevention (DLP) automatically identifies documents, emails, and conversations that contain sensitive data by scanning them for sensitive data and then enforcing policy on those objects. Policies are enforced on documents in SharePoint and OneDrive for Business. They're also enforced when users send email, and in Teams chats and channel conversations. Policies can be configured to look for keywords, sensitive data types, retention labels, and whether data is shared within the organization or externally. Controls are provided to help organizations fine-tune DLP policies to reduce false positives. When sensitive data is found, customizable policy tips can be displayed to users within Microsoft 365 applications to inform them that their content contains sensitive data and then propose corrective actions. Policies can also prevent users from accessing documents, sharing documents, or sending emails that contain certain types of sensitive data. Microsoft 365 supports more than 100 built-in sensitive data types. Organizations can configure custom sensitive data types to meet their policies.

Rolling out MIP and DLP policies to organizations requires careful planning and a user education program so that employees understand the organization's data classification schema and which types of data are considered sensitive. Providing employees with tools and educational programs that help them identify sensitive data and understand how to handle it makes them part of the solution for mitigating information security risks.

The signals generated by and fed to Identity Protection can also be fed into tools like Conditional Access to make access decisions or to a security information and event management (SIEM) tool for investigation based on an organization's enforced policies.

Identity Protection helps organizations automatically protect against identity compromise by taking advantage of cloud intelligence powered by advanced detections based on heuristics, user and entity behavior analytics, and machine learning across the Microsoft ecosystem.

Defend the fortress

Microsoft recently launched the Microsoft 365 Defender solution, which is designed to secure the modern organization from the evolving threat landscape. By leveraging the Intelligent Security Graph, the Threat Protection solution offers comprehensive, integrated security against multiple attack vectors.

The Intelligent Security Graph

Security services from Microsoft 365 are powered by the Intelligent Security Graph. To combat cyberthreats, the Intelligent Security Graph uses advanced analytics to link threat intelligence and security signals from Microsoft and its partners. Microsoft operates global services at a massive scale, gathering trillions of security signals that power protection layers across the stack. Machine-learning models assess this intelligence, and the signal and threat insights are widely shared across our products and services. This enables us to detect and respond to threats quickly and bring actionable alerts and information to customers for remediation. Our machine learning models are continuously trained and updated with new insights, helping us build more-secure products and provide more proactive security.

Microsoft Defender for Office 365 provides an integrated Microsoft 365 service that protects organizations from malicious links and malware delivered through email and Office documents. One of the most common attack vectors that affects users today is email phishing attacks. These attacks can be targeted at specific users and can be very convincing, with some call to action that prompts the user to click a malicious link or open an attachment that contains malware. Once a computer is infected, the attacker can either steal the user's credentials and move laterally across the organization or exfiltrate emails and data to look for sensitive information. Defender for Office 365 supports safe attachments and safe links by evaluating documents and links at click-time for potentially malicious intent and blocks access. Email attachments are opened in a protected sandbox before they're delivered to a user's mailbox. It also evaluates links in Office documents for malicious URLs. Defender for Office 365 also protects links and files in SharePoint Online, OneDrive for Business, and Teams. If a malicious file is detected, Defender for Office 365 automatically locks that file to reduce potential damage.

Microsoft Defender for Endpoint is a unified endpoint security platform for preventative protection, post-breach detection, and automated investigation and response. Defender for Endpoint provides built-in capabilities for

discovery and protection of sensitive data on enterprise endpoints.

Microsoft Cloud App Security (MCAS) enables organizations to enforce policies at a granular level and to detect behavioral anomalies based on individual user profiles that are automatically defined by using machine learning. MCAS policies can build on Azure Conditional Access policies to protect sensitive company assets by evaluating additional signals related to user behavior and properties of the documents that are accessed. Over time, MCAS learns what's typical behavior for each employee with regard to the data they access and the applications they use. Based on learned behavior patterns, policies can then automatically enforce security controls if an employee acts outside of that behavioral profile. For example, if an employee typically accesses an accounting application from 9 AM to 5 PM Monday through Friday but suddenly starts to access that application heavily on a Sunday evening, MCAS can dynamically enforce policies to require the user to reauthenticate. This helps ensure that the user's credentials haven't been compromised. MCAS can also help identify "shadow IT" in the organization, which helps information security teams ensure that employees are using sanctioned tools when they work with sensitive data. Finally, MCAS can protect sensitive data anywhere in the Cloud, even outside the Microsoft 365 platform. It allows organizations to sanction (or unsanction) specific external Cloud apps, controlling access and monitoring usage.

Microsoft Defender for Identity is a cloud-based security solution that leverages your on-premises Active Directory signals to identify, detect, and investigate advanced threats, compromised identities, and malicious insider actions directed at your organization. AATP enables SecOp analysts and security professionals detect advanced attacks in hybrid environments to:

- Monitor users, entity behavior, and activities by using learning-based analytics.
- Protect user identities and credentials stored in Active Directory.
- Identify and investigate suspicious user activities and advanced attacks throughout the kill chain.
- Provide clear incident information on a simple timeline for fast triage.

Govern data and manage records

Financial institutions must retain their records and information according to their regulatory, legal, and business obligations as represented within their corporate retention schedule. For example, the SEC mandates retention periods of three to six years, based on record type, with immediate accessibility for the first two years.

Organizations face legal and regulatory compliance risks if data is under-retained (discarded too early), and now also manage regulations that mandate disposal when information is no longer required. Effective records-management strategies emphasize a practical and consistent approach so that information is disposed of appropriately while minimizing cost and risk to the organization.

In addition, regulatory mandates from the New York State Department of Financial Services require covered entities to maintain policies and procedures for disposal of nonpublic information. 23 NYCRR 500, Section 500.13, Limitations on Data Retention requires that "As part of its cybersecurity program, each Covered Entity shall include policies and procedures for the secure disposal on a periodic basis of any Nonpublic Information identified in section 500.01(g)(2)-(3) of this Part that is no longer necessary for business operations or for other legitimate business purposes of the Covered Entity, except where such information is otherwise required to be retained by law or regulation."

Financial institutions manage vast amounts of data. And some retention periods are triggered by events, such as a contract expiring or an employee leaving the organization. In this atmosphere, it can be challenging to apply record retention policies. Approaches to assigning record retention periods accurately across organizational documents can vary. Some apply retention policies broadly or leverage autoclassification and machine-learning techniques. Others identify an approach that requires a more granular process that assigns retention periods uniquely to individual documents.

Microsoft 365 provides flexible capabilities to define retention labels and policies to intelligently implement records-management requirements. A record manager defines a retention label, which represents a "record type" in a traditional retention schedule. The retention label contains settings that define these details:

- How long a record is retained
- What occurs when the retention period expires (delete the document, start a disposition review, or take no action)
- What triggers the retention period to start (created date, last modified date, labeled date, or an event) and marks the document or email as a record (meaning it can't be edited or deleted)

The retention labels are then published to SharePoint or OneDrive sites, Exchange mailboxes, and Microsoft 365 groups. Users can apply the retention labels manually to documents and emails. Record managers can use intelligence to automatically apply the labels. Intelligent capabilities can be based on ninety-plus built-in sensitive information types (such as ABA outing number, US bank account number, or US Social Security Number). They're also customizable based on keywords or sensitive data found in documents or emails, such as credit card numbers or other personally identifiable information or based on SharePoint metadata. For data that's not easily identified through manual or automated pattern matching, trainable classifiers can be used to classify documents intelligently based on machine learning techniques.

The **Securities and Exchange Commission (SEC)** requires broker-dealers and other regulated financial institutions to retain all business-related communications. These requirements apply to many types of communications and data, including emails, documents, instant messages, faxes, and more. **SEC rule 17a-4** defines the criteria that these organizations must meet to store records in an electronic data storage system. In 2003, the SEC issued a release that clarified these requirements. It included the following criteria:

- Data preserved by an electronic storage system must be non-rewriteable and non-erasable. This is referred to as a WORM requirement (write once, read many).
- The storage system must be able to store data beyond the retention period required by the rule, in case of a subpoena or other legal order.
- An organization wouldn't violate the requirement in paragraph (f)(2)(ii)(A) of the rule if it used an electronic storage system that prevents the overwriting, erasing, or otherwise altering of a record during its required retention period through the use of integrated hardware and software control codes.
- Electronic storage systems that merely "mitigate" the risk that a record will be overwritten or erased, for

example by relying on access control, don't meet the requirements of the rule.

To help financial institutions meet the requirements of SEC rule 17a-4, Microsoft 365 provides a combination of capabilities related to how data is retained, policies are configured, and data is stored within the service. These include:

- **Preservation of data (Rule 17a-4(a), (b)(4))** – Retention labels and policies are flexible to meet organizational needs and may be automatically or manually applied to different types of data, documents, and information. A wide variety of data types and communications are supported, including documents in SharePoint and OneDrive for Business, data within Exchange Online mailboxes, and data in Teams.

- **Non-rewriteable, non-erasable format (Rule 17a-4(f)(2)(ii)(A))** – Preservation Lock capability for retention policies allows records managers and administrators to configure retention policies to be restrictive, such that they can no longer be modified. This prohibits anyone from removing, disabling, or modifying the retention policy in any way. This means that once Preservation Lock is enabled, it can't be disabled, and there is no method by which any data to which the retention policy has been applied can be overwritten, modified, or deleted during the retention period. In addition, the retention period can't be shortened. However, the retention period can be lengthened, when there's a legal requirement to continue retention of data.

  When a Preservation Lock is applied to a retention policy, the following actions are restricted:

  - The retention period of the policy can only be increased. It can't be shortened.
  - Users can be added to the policy, but existing users configured in the policy can't be removed.
  - The retention policy can't be deleted by any administrator within the organization.

  Preservation Lock helps ensure that no user, not even administrators with the highest levels of privileged access, can change the settings, modify, overwrite or delete the data that has been stored, bringing archiving in Office 365 in line with the guidance provided in the SEC 2003 Release.

- **Quality, accuracy, and verification of storage/serialization and indexing of data (Rule 17a-4(f)(2) (ii)(B) and (C))** – Office 365 workloads each contain capabilities for automatically verifying the quality and accuracy of the process for recording data on storage media. In addition, data is stored by utilizing metadata and timestamps to ensure sufficient indexing to allow for effective searching and retrieval of data.

- **Separate storage for duplicate copies (Rule 17a-4(f)(3(iii))** – The Office 365 cloud service stores duplicate copies of data as a core aspect of its high availability. This is accomplished by implementing redundancy at all levels of the service, including at the physical level on all servers, at the server level within the data center, and at the service level for geographically dispersed data centers.

- **Downloadable and accessible data (Rule 17a-4(f)(2)(ii)(D))** – Office 365 generally permits data that's been labeled for retention to be searched for, accessed, and downloaded in place. And it allows data in Exchange Online Archives to be searchable by using built-in eDiscovery features. Data can then be downloaded as needed in standard formats, including EDRML and PST.

- **Audit requirements (Rule 17a-4(f)(3)(v))** – Office 365 provides audit logging for every administrative and user action that modifies data objects, configures or modifies retention policies, performs eDiscovery searches, or modifies access permissions. Office 365 maintains a comprehensive audit trail, including data about who performed an action, when it was performed, details about the action, and the commands that were performed. The audit log can then be output and included as part of formal audit processes as required.

Finally, Rule 17a-4 requires organizations to retain records for many types of transactions so that they're immediately accessible for two years. Records must be further retained for three to six years with non-

immediate access. Duplicate records must also be kept for the same period at an off-site location. Office 365 records-management capabilities enable records to be retained such that they can't be modified or deleted but can be easily accessed for a time period that's controlled by the record manager. These periods can span days, months, or years, depending on the organization's regulatory-compliance obligations.

Upon request, Microsoft will provide an attestation letter of compliance with SEC 17a-4 if required by an organization.

In addition, these capabilities also help Microsoft 365 meet storage requirements for CFTC Rule 1.31(c)-(d) from the **U.S. Commodity Futures Trading Commission** and FINRA Rule Series 4510 from the **Financial Industry Regulatory Authority.** Collectively, these rules represent the most-prescriptive guidance globally for financial institutions to retain records.

Additional details about how Microsoft 365 complies with SEC rule 17a-4 and other regulations is available at Assessment of Office 365 Exchange Online SEC 17a-4(f) / CFTC 1.31(c)-(d) by Cohasset Associates.

Establish ethical walls with information barriers

Financial institutions can be subject to regulations that prevent employees in certain roles from exchanging information or collaborating with other roles. For example, FINRA has published rules 2241(b)(2)(G), 2242(b)(2)(D), (b)(2)(H)(ii) and (b)(2)(H)(iii) that require members to:

"(G) establish information barriers or other institutional safeguards reasonably designed to ensure that research analysts are insulated from the review, pressure, or oversight by persons engaged in investment banking services activities or other persons, including sales and trading personnel, who might be biased in their judgment or supervision;" and "(H) establish information barriers or other institutional safeguards reasonably designed to ensure that debt research analysts are insulated from the review, pressure, or oversight by persons engaged in: (i) investment banking services; (ii) principal trading or sales and trading activities; and (iii) other persons who might be biased in their judgment or supervision;"

Ultimately, these rules require organizations to establish policies and implement information barriers between roles involved in banking services, sales, or trading from exchanging information and communications with analysts.

Information barriers provides the ability to establish ethical walls within your Office 365 environment, allowing compliance administrators or other authorized administrators to define policies that allow or prevent communications between groups of users in Teams. Information barriers perform checks on specific actions to prevent unauthorized communication. Information barriers can also restrict communication in scenarios where internal teams are working on mergers/acquisitions or sensitive deals, or working with sensitive internal information that must be heavily restricted.

Information barriers in Microsoft 365 support conversations and files in Teams. They can prevent the following types of communications-related actions to help comply with FINRA regulations:

- Search for a user
- Add a member to a team, or continue to participate with another member in a team
- Start or continue a chat session
- Start or continue a group chat
- Invite someone to join a meeting
- Share a screen
- Place a call

Implement supervisory control

Financial institutions are typically required to establish and maintain a supervisory function within their

organizations to monitor the activities of employees and to help it achieve compliance with applicable securities laws. Specifically, FINRA has established these supervision requirements:

- FINRA Rule 3110 (Supervision) requires firms to have written supervisory procedures (WSPs) to supervise activities of its employees and the types of businesses in which it engages. In addition to other requirements, procedures must include:
  - Supervision of supervisory personnel
  - Review of a firm's investment banking, securities business, internal communications, and internal investigations
  - Review of transactions for insider trading
  - Review of correspondence and complaints

  Procedures must describe the individuals responsible for reviews, supervisory activity each person will perform, review frequency, and the types of documentation or communications under review.
- FINRA Rule 3120 (Supervisory Control System) requires firms to have a system of supervisory control policies and procedures (SCPs) that validates their written supervisory procedures as defined by Rule 3110. Firms are required not only to have WSPs but also to have policies that test these procedures annually to validate their ability to ensure compliance with applicable securities laws and regulations. Risk-based methodologies and sampling may be used to define the scope of testing. Among other requirements, this rule requires firms to provide an annual report to senior management that includes a summary of test results and any significant exceptions or amended procedures in response to test results.

Communication compliance

Communication compliance in Microsoft 365 enables organizations to pre-configure policies to capture employee communications for monitoring and review by authorized supervisors. Policies in communication compliance can capture internal/external email and attachments, Teams chat and channel communications, and Skype for Business Online chat communications and attachments. In addition, communication compliance can ingest communications and data from third-party services (such as Bloomberg, Thomson Reuters, LinkedIn, Twitter, Facebook, Box, and Dropbox). The comprehensive nature of communications that can be captured and reviewed within an organization, and the extensive conditions with which policies may be configured, allow

communication compliance policies to help financial institutions comply with FINRA Rule 3110. Policies may be configured to review communications for individuals or groups. Designated supervisors can be assigned at an individual or group level. Comprehensive conditions can be configured to capture communications based on inbound or outbound messages, domains, retention labels, keywords or phrases, keyword dictionaries, sensitive data types, attachments, message size, or attachment size. Reviewers get a dashboard in which they can review flagged communications, act on communications that potentially violate policies, and mark flagged items as resolved. They can also review the results of reviews and items that were previously resolved.

Communication compliance provides reports that enable policy review activities to be audited based on the policy and the reviewer. Reports are available to validate that policies are working as defined by an organization's written supervision policies. They can also be used to identify communications that require review and those that are not compliant with corporate policy. Finally, all activities related to configuring policies and reviewing communications are audited in the Office 365 unified audit log. As a result, communication compliance in Microsoft 365 also helps financial institutions to comply with FINRA Rule 3120.

In addition to complying with FINRA rules, communication compliance allows organizations to monitor communications for compliance with other legal requirements, corporate policies, and ethical standards. Communication compliance provides built-in threat, harassment, and profanity classifiers that help reduce false positives when reviewing communications, saving reviewers time during the investigation and remediation process. It also allows organizations to reduce risk by monitoring communications when they undergo sensitive changes, such as mergers and acquisitions or leadership changes.

Protect against data exfiltration and insider risk

A common threat to enterprises is data exfiltration, or the act of extracting data from an organization. This risk can be a significant concern for financial institutions due to the sensitive nature of the information that can be accessed day to day. With the increasing number of communications channels available and the proliferation of tools for moving data, advanced capabilities are typically required to mitigate the risks of data leaks, policy violations, and insider risk.

Insider risk management

Enabling employees with online collaboration tools that can be accessed anywhere inherently brings risk to the

organization. Employees may inadvertently or maliciously leak data to attackers or competitors. Alternatively, they may exfiltrate data for personal use or take data with them to a future employer. These scenarios present serious risks to financial services institutions from both security and compliance standpoints. Identifying these risks when they occur and quickly mitigating them requires both intelligent tools for data collection and collaboration across departments such as legal, human resources, and information security.

Microsoft 365 recently launched an insider risk management solution that correlates signals across Microsoft 365 services and uses machine-learning models to analyze user behavior for hidden patterns and signs of insider risk. This tool enables collaboration between security operations, internal investigators, and HR so that they can easily remediate cases based on predetermined workflows.

For example, insider risk management in Microsoft 365 can correlate signals from a user's Windows 10 desktop, such as copying files to a USB drive or emailing a personal email account, with activities from online services such as Office 365 email, SharePoint Online, Microsoft Teams, or OneDrive for Business, to identify data exfiltration patterns. It can also correlate these activities with employees leaving an organization, which is a common data exfiltration pattern. It can monitor multiple activities and behavior over time. When common patterns emerge, it can raise alerts and help investigators focus on key activities to verify a policy violation with a high degree of confidence. Insider risk management can pseudo-anonymize data from investigators to help meet data privacy regulations, while still surfacing key activities that help them perform investigations efficiently. It allows investigators to package and securely send key activity data to the HR and legal departments, following common escalation workflows for raising cases for remediation action.

Insider risk management in Microsoft 365 significantly increases capabilities of organizations to monitor and investigate insider risks while allowing organizations to still meet data privacy regulations and follow established escalation paths when cases require higher-level action. For more information about insider risk management in Microsoft 365, see Modern risk pain points and Workflow in Insider risk management in Microsoft 365.

Tenant restrictions

Organizations that deal with sensitive data and put a strict emphasis on security typically want to control the online resources that users can access. At the same time, they want to enable secure collaboration through online services such as Office 365. As a result, controlling the Office 365 environments that users can access

becomes a challenge because noncorporate Office 365 environments can be used to exfiltrate data from corporate devices either maliciously or inadvertently. Traditionally, organizations restrict the domains or IP addresses that users can access from corporate devices. But this doesn't work in a cloud-first world, where users need to legitimately access Office 365 services.

Microsoft 365 provides the tenant restrictions the capability to address this challenge. Tenant restrictions can be configured to restrict employee access to external Office 365 enterprise tenants using rogue identities (identities that aren't part of your corporate directory). Today, tenant restrictions apply across the tenant, allowing access to only those tenants that appear on the list that you configure. Microsoft is continuing to develop this solution to increase granularity of control and enhance the protections it provides.

Conclusion

Microsoft 365 and Teams provide an integrated and comprehensive solution for financial services companies, enabling simple yet powerful cloud-based collaboration and communications capabilities across the enterprise. By using security and compliance technologies from Microsoft 365, institutions can operate in a more secure and compliant manner with robust security controls to protect data, identities, devices, and applications from various operational risks, including cybersecurity and insider risks. Microsoft 365 provides a fundamentally secure platform on which financial services organizations can achieve more while protecting their company, employees, and customers.

Key Compliance and Security Considerations for the Energy Industry

2/18/2021 • 31 minutes to read • Edit Online

Introduction

The energy industry provides society with fuel and critical infrastructure that people rely on every day. In order to ensure the reliability of infrastructure related to bulk power systems, regulatory authorities impose strict standards on energy industry organizations. These regulatory standards relate not only to the generation and transmission of power, but also to the data and communications that are critical to the day to day operations of energy companies.

Organizations in the energy industry work with and exchange numerous types of information as part of their regular operations, including customer data, capital engineering design documentation, resource location maps, project management artifacts, performance metrics, field service reports, environmental data, and performance metrics. As these organizations look to transform their operations and collaboration systems into modern digital platforms, they are looking to Microsoft as a trusted Cloud Service Provider (CSP) and Microsoft 365 as their best-of-breed collaboration platform. Since Microsoft 365 is intrinsically built on the Microsoft Azure platform, organizations should examine both platforms as they consider their compliance and security controls when moving to the Cloud.

In North America, the North America Electric Reliability Corporation (NERC) enforces reliability standards that are referred to as NERC Critical Infrastructure Protection (CIP) standards. NERC is subject to oversight by the U.S. Federal Energy Regulatory Commission (FERC) and governmental authorities in Canada. All bulk power system owners, operators, and users must register with NERC and must comply with NERC CIP standards. Cloud Service Providers and third-party vendors such as Microsoft are not subject to NERC CIP standards; however, the CIP standards include objectives that should be considered when Registered Entities use vendors in the operation of the Bulk Electric System (BES). Microsoft customers operating Bulk Electric Systems are wholly responsible for ensuring their own compliance with NERC CIP standards.

For information about Microsoft cloud services and NERC, see the following resources:

- NERC CIP Standards and Cloud Computing
- Cloud Implementation Guide for NERC Audits

Regulatory standards that are recommended for consideration by energy organizations include FedRAMP (US Federal Risk and Authorization Management Program) which is based on and augments the NIST SP 800-53 Rev 4 standard (National Institute of Standards and Technology).

- Microsoft Office 365 and Office 365 U.S. Government have each been granted a FedRAMP ATO (Authorization to Operate) at the Moderate Impact Level.
- Azure and Azure Government have each been granted a FedRAMP High P-ATO (Provisional Authorization to Operate), which represents the highest level of FedRAMP authorization.

For information about Microsoft cloud services and FedRAMP, see the following resources:

- Microsoft FedRAMP overview
- Office 365 FedRAMP reports

These achievements are significant for the energy industry because a comparison between the FedRAMP Moderate control set and NERC CIP requirements shows that FedRAMP Moderate controls encompass all the NERC CIP requirements. For additional information, Microsoft has developed a Cloud Implementation Guide for NERC Audits that includes a control mapping between the current set of NERC CIP standards and FedRAMP Moderate control set as documented in NIST 800-53 Rev 4.

As the energy industry looks to modernize their collaboration platforms, careful consideration is required for the configuration and deployment of collaboration tools and security controls, including:

- Assessment of common collaboration scenarios
- Access to data required by employees to be productive
- Regulatory compliance requirements
- Associated risks to data, customers and the organization

Microsoft 365 is a modern workplace cloud environment that can provide secure and flexible collaboration across the enterprise, as well as controls and policy enforcement to adhere to the most stringent regulatory compliance frameworks. Through the following topics, this paper will explore how the Microsoft 365 platform helps the energy industry move to a modern collaboration platform, while helping keep data and systems both secure and compliant with regulations:

- Provide a Comprehensive Collaboration Platform with Microsoft Teams
- Provide Secure and Compliant Collaboration in the Energy Industry
- Identify Sensitive Data and Prevent Data Loss
- Govern Data by Effectively Managing Records
- Comply with FERC and FTC Regulations for Energy Markets
- Protect Against Data Exfiltration and Insider Risk

As a Microsoft partner, Protiviti contributed to and provided material feedback to this article.

Provide a Comprehensive Collaboration Platform with Microsoft Teams

Collaboration typically requires multiple forms of communication, the ability to store and access documents, and the ability to integrate other applications as needed. Whether they are global enterprises or local companies, employees in the energy sector typically need to collaborate and communicate with members of other departments or across teams. They also often need to communicate with external partners, vendors, or clients. As a result, utilizing systems that create silos or make it difficult to share information is typically not recommended. That said, we still want to ensure that employees are sharing information securely and according to policy.

Providing employees with a modern, cloud-based collaboration platform, which allows them to choose and

easily integrate the tools that make them most productive, empowers them to find the best ways to work and collaborate. Using Microsoft Teams, in conjunction with security controls and governance policies to protect the organization, can help your workforce to easily collaborate in the cloud.

Microsoft Teams provides a collaboration hub for your organization, which quickly brings people together to work effectively and collaborate with other team members on common initiatives or projects. It allows team members to conduct conversations, collaborate, and co-author documents. It enables people to store and share files with team members or those outside the team. It also allows them to hold live meetings with integrated enterprise voice and video. Microsoft Teams can be customized with easy access to Microsoft apps such as Planner, Dynamics 365, Power BI, and other third-party line-of-business applications. Teams simplifies access to Office 365 services and third-party apps to centralize collaboration and communication needs for the organization.

Every Microsoft Team is backed by an Office 365 Group. An Office 365 Group is considered the membership provider for numerous Office 365 services, including Microsoft Teams. As such, Office 365 Groups are used to securely control which users are considered members and which are owners of the group, thereby restricting the members and owners of the Team. This allows us to easily control which users have access to varying capabilities within Teams. As a result, Team members and owners may only access the capabilities that they are permitted to utilize.

A common scenario where Microsoft Teams can benefit energy organizations is collaborating with contractors or external firms as part of a field service programs such as vegetation management. Contractors are typically engaged to manage vegetation or remove trees around power system installations, and they often need to receive work instructions, communicate with dispatchers and other field service personnel, take and share pictures of external surroundings, sign off when work is complete and share data back with head office. Traditionally, these programs have been run using phone, text, paper work orders, or custom applications. This can present numerous challenges including:

- Processes are manual or analog, making metrics difficult to track
- Communications are not all captured in one place
- Data is siloed and not necessarily shared with all employees that need it
- Work may not be performed consistently or efficiently
- Custom applications are not integrated with collaboration tools, making it difficult to extract and share data or measure performance

Microsoft Teams can provide an easy-to-use collaboration space to securely share information and conduct conversations between team members and external field service contractors. Teams can be used to conduct meetings, place voice calls, centrally store and share work orders, collect field data, upload photos, integrate with business process solutions (built with Power Apps and Power Automate), and integrate line of business apps. This type of field service data may be considered low impact; however, efficiencies can be gained by centralizing communications and access data between employees and field service personnel in these scenarios.

Another example where Microsoft Teams can benefit the energy industry is when field service personnel are working to restore service during an outage. Field staff often require fast access to schematic data for substations, generating stations, or blue prints for assets in the field. This data is considered high impact and must be protected according to NERC CIP regulations. Field service work during outages requires communication between field staff and office employees, and in turn with end customers. Centralizing communications and data sharing in Microsoft Teams provides field staff with an easy method to both access critical data and communicate information or status back to head office. For example, Microsoft Teams enables field staff to join conference calls while on route to an outage. Field staff can also take photos or video of their environment and share those with head office, which is particularly important when field equipment does not match schematics. Data and status collected from the field can then be surfaced to office employees and leadership through data visualization tools such as Power BI. Ultimately, Microsoft Teams can make field staff more efficient and productive in these critical situations.

Teams: Improve collaboration and reduce compliance risk

Microsoft 365 provides common policy capabilities for Microsoft Teams through its use of Office 365 Groups as an underlying membership provider. These policies can help improve collaboration and help meet compliance needs.

Office 365 Group Naming Policies help to ensure that Office 365 Groups, and therefore Microsoft Teams, are named according to corporate policy. The name of a Team can present challenges if not named appropriately – for example, employees may not know which teams to work or share information within if they are incorrectly named. Group naming policies can enforce good hygiene and may also prevent use of specific words, such as reserved words or inappropriate terminology.

Office 365 Group Expiration Policies help to ensure that Office 365 Groups, and therefore Microsoft Teams, are not retained for longer periods of time than required by the organization. This capability helps to prevent two key information management issues:

- The proliferation of Microsoft Teams that are not necessary or used
- The over-retention of data that is no longer required by the organization

Administrators may specify an expiration period in days for Office 365 Groups, such as 90, 180 or 365 days. If a service which is backed by an Office 365 group is inactive for the expiration period, group owners are notified and if no action is taken then the Office 365 Group and all its related services including Microsoft Teams will be deleted.

The over-retention of data in a Microsoft Team can pose litigation risks to organizations, and the use of expiration policies is a recommended method for protecting the organization. Combined with built-in retention labels and policies, Microsoft 365 helps ensure that organizations are only retaining the data required to meet regulatory compliance obligations.

Teams: Integrate custom requirements with ease

Microsoft Teams enables self-service creation of Teams by default. However, many regulated organizations wish to control and understand which collaboration spaces are currently in use by employees, which spaces contain sensitive data, and who the owners are of spaces throughout their organization. To facilitate these controls, Microsoft 365 allows organizations to disable self-service Teams creation and, using built-in Microsoft 365 business process automation tools such as Power Apps and Power Automate, allows organizations to build simple processes to request a new Team. Completing an easy to use form, an approval can be automatically requested by a manager. Once approved, the Team can be automatically provisioned and the requestor is sent a link to their new Team. By building such processes, organizations may also integrate custom requirements to facilitate other business processes.

Provide Secure and Compliant Collaboration in the Energy Industry

As mentioned, Microsoft Office 365 and Office 365 U.S. Government have each achieved FedRAMP ATO at the Moderate Impact Level, and Azure and Azure Government have achieved a FedRAMP High P-ATO which represents the highest level of FedRAMP authorization. Additionally, the FedRAMP moderate control set encompasses all of the NERC CIP requirements, thereby allowing energy industry organizations ("registered entities") to leverage existing FedRAMP authorizations as a scalable and efficient approach to addressing NERC audit requirements. However, its important to note that FedRAMP is not a point-in-time certification but an assessment and authorization program that includes provisions for continuous monitoring. Although this provision applies primarily to the CSP, Microsoft customers operating Bulk Electric Systems are responsible for ensuring their own compliance with NERC CIP standards and it is generally a recommended practice to continuously monitor the organization's compliance posture to help ensure ongoing compliance with regulations.

Microsoft provides a key tool to assist with monitoring compliance with regulations over time:

- **Microsoft Compliance Manager** helps the organization understand its current compliance posture and the actions which it can take to help improve that posture. Compliance Manager calculates a risk-based score measuring progress in completing actions that help reduce risks around data protection and regulatory standards. Compliance Manager provides an initial score based on the Microsoft 365 data protection baseline. This baseline is a set of controls that includes common industry regulations and standards. While this score is a good starting point, Compliance Manager becomes more powerful once an organization adds assessments that are more relevant to their industry. Compliance Manager supports a number of regulatory standards that are relevant for NERC CIP compliance obligations, including the FedRAMP Moderate Control Set, NIST 800-53 Rev. 4, and AICPA SOC 2. Energy industry organizations may also create or import custom control sets if needed.

The workflow-based capabilities built into Compliance Manager allow energy organizations to transform and digitize their regulatory compliance processes. Traditionally, challenges faced by compliance teams in the energy industry include:

- Inconsistent reporting or tracking of progress on remediation actions
- Inefficient or ineffective processes
- Insufficient resources or lack of ownership
- Lack of real-time information and human error

By automating aspects of regulatory compliance processes through the use of Compliance Manager, organizations can reduce the administrative burden on legal and compliance functions. This tooling can help address these challenges by providing more up to date information on remediation actions, more consistent reporting, and documented ownership of actions which is linked to the implementation of actions. Organizations can automatically track remediation actions over time and see overall efficiency gains. This can in turn enable staff to focus more effort on gaining insights and developing strategies to help navigate risk more effectively.

Compliance Manager does not express an absolute measure of organizational compliance with any particular standard or regulation. It expresses the extent to which you have adopted controls which can reduce the risks to personal data and individual privacy. Recommendations from Compliance Manager should not be interpreted as a guarantee of compliance. The customer actions provided in Compliance Manager are recommendations; it is up to each organization to evaluate the effectiveness of these recommendations in their respective regulatory environment prior to implementation. Recommendations found in Compliance Manager should not be interpreted as a guarantee of compliance.

Many cyber security related controls are included in the FedRAMP Moderate Control Set and NERC CIP standards. However, key controls related to the Microsoft 365 platform include security management controls (CIP-003-6), account and access management/access revocation (CIP-004-6), electronic security perimeter (CIP-005-5), security event monitoring, and incident response (CIP-008-5). The following foundational Microsoft 365 capabilities help to address the risks and requirements included in these topics.

Secure User Identities and Control Access

Protecting access to documents and applications begins with strongly securing user identities. As a foundation, this requires providing a secure platform for the enterprise to store and manage identities and providing a trusted means of authentication. It also requires dynamically controlling access to these applications. As employees work, they may move from application to application or across multiple locations and devices. As a result, access to data must be authenticated at each step of the way. In addition, the authentication process must support a strong protocol and multiple factors of authentication (one-time SMS pass code, authenticator app, certificate, etc.) so that we can ensure that identities have not been compromised. Finally, enforcing risk-based access policies are a key recommendation to protecting data and applications from insider threats, inadvertent data leaks and data exfiltration.

Microsoft 365 provides a secure identify platform with **Azure Active Directory (Azure AD)** where identities

are centrally stored and securely managed. Azure Active Directory, along with a host of related Microsoft 365 security services, forms the basis for providing employees with the access they need to work securely while also protecting the organization from threats.

Azure AD Multi-Factor Authentication (MFA) is built into the platform and provides an additional layer of protection to help ensure users are who they say they are when accessing sensitive data and applications. Azure MFA requires at least two forms of authentication, such as a password and a known mobile device. It supports several second factor authentication options, including: the Microsoft Authenticator app, a one-time passcode delivered via SMS, receiving a phone call where a user must enter a PIN, and smart cards or certificate-based authentication. In the event a password is compromised, a potential hacker still needs the user's phone to gain access to organizational data. In addition, Microsoft 365 uses Modern Authentication as a key protocol, bringing the same strong authentication experience from web browsers to collaboration tools, including Microsoft Outlook and Microsoft Office apps.

Azure AD Conditional Access provides a robust solution for automating access control decisions and enforcing policies to protect company assets. A common example is when an employee wishes to access an application containing sensitive customer data, and they are automatically required to perform a multi-factor authentication to specifically access that application. Azure Conditional Access brings together signals from a user's access request, such as properties about the user, their device, location, network, and the app or repository they are trying to access. It can dynamically evaluate every attempt to access the application against configured policies. If user or device's risk is elevated, or if other conditions are not met, Azure AD can automatically enforce policy such as dynamically requiring MFA, restricting or even blocking access. This helps ensure that sensitive assets are protected in dynamically changing environments.

Microsoft Defender for Office 365 provides an integrated service to protect organizations from malicious links and malware delivered through email. One of the most common attack vectors impacting users today is email phishing attacks. These attacks can be carefully targeted at specific high-profile employees and can be crafted to be very convincing. They typically contain some call to action requiring a user to click a malicious link or open an attachment with malware. Once infected, an attacker can steal a user's credentials and move laterally across the organization. They can also exfiltrate emails and data looking for sensitive information. Microsoft Defender for Office 365 evaluates links at click-time for potentially malicious sites and blocks them. Email attachments are opened in a protected sandbox prior to delivering them to a user's mailbox.

Microsoft Cloud App Security (MCAS) provides organizations with the ability further enforce policies at a granular level and detect behavioral anomalies based on individual user profiles that are automatically defined using Machine Learning. MCAS can build on Azure Conditional Access policies, to further protect sensitive assets by evaluating additional signals related to user behavior and properties of the documents being accessed. Over time, MCAS will learn what is considered typical behavior for each employee, with regard to the data they access and the applications they use. Based on learned behavioral patterns, policies can automatically enforce security controls if an employee goes outside of that behavioral profile. For example, if an employee typically accesses an accounting app from 9am to 5pm, Monday to Friday, but that same user begins to access that application heavily on a Sunday evening, MCAS can dynamically enforce policies to require the user to re-authenticate. This helps ensure that credentials have not been compromised. In addition, MCAS can help discover and identify Shadow IT in the organization, helping InfoSec teams ensure that employees are using sanctioned tools when working with sensitive data. Finally, MCAS can protect sensitive data anywhere in the Cloud, even outside of the Microsoft 365 platform. It allows organizations to sanction (or un-sanction) specific external Cloud apps, controlling access and monitoring when users work in those applications.

Azure Active Directory, and the related Microsoft 365 security services, provide the foundation upon which a modern cloud collaboration platform can be rolled out to energy industry organizations so that access to data and applications can be strongly secured and regulatory compliance obligations can be met. To summarize, these tools provide the following key capabilities:

- Centrally store and securely manage user identities
- Use a strong authentication protocol including multi-factor authentication to authenticate users on access

requests and provide a consistent and robust authentication experience across any application
- Dynamically validate policies on all access requests, incorporating multiple signals into the policy decision making process, including identity, user/group membership, application, device, network, location, and real-time risk score
- Validate granular policies based on user behavior and file properties and dynamically enforce additional security measures when required
- Identify shadow IT in the organization and allow InfoSec teams to sanction or block cloud applications
- Monitor and control access to Microsoft and non-Microsoft cloud applications
- Proactively protect against email phishing and ransomware attacks

Identify Sensitive Data and Prevent Data Loss

The FedRAMP Moderate Control Set and NERC CIP standards also include information protection as a key control requirement (CIP-011-2). These requirements specifically address the need to identify information related to BES (Bulk Electric System) Cyber System Information, the protection and secure handling of that information, including storage, transit, and use. Specific examples of BES Cyber System Information may include security procedures or security information about systems that are fundamental to operating the bulk electric system (BES Cyber Systems, Physical Access Control Systems, and Electronic Access Control or monitoring systems) that is not publicly available and could be used to allow unauthorized access or unauthorized distribution. However, the same need exists to identify and protect customer information that is critical to the day to day operations of energy organizations.

Microsoft 365 allows sensitive data to be identified and protected within the organization through a combination of powerful capabilities, including:

- **Microsoft Information Protection (MIP)** for both user-based classification and automated classification of sensitive data
- **Office 365 Data Loss Prevention (DLP)** for automated identification of sensitive data using sensitive data types (i.e. regular expressions) and keywords, and policy enforcement

Microsoft Information Protection (MIP) allows employees to classify documents and emails with sensitivity labels. Sensitivity labels can be applied manually by users to documents within the Microsoft Office apps and to emails within Microsoft Outlook. Sensitivity labels can in turn automatically apply document markings, protection through encryption, and enforce rights management. They may also be applied automatically, by configuring policies which use keywords and sensitive data types (credit card numbers, social security numbers, identity numbers, etc.) to automatically find and classify sensitive data.

In addition, Microsoft provides trainable classifiers which use machine learning models to identify sensitive data based on what the content is, as opposed to simply through pattern matching or by the elements within the content. A classifier learns how to identify a type of content by looking at numerous examples of the content to be classified. Training a classifier begins by providing it with examples of content within a particular category. Once it processes those examples, the model is tested by providing it with a mix of both matching and non-matching examples. The classifier then predicts whether a given example falls into the category or not. A person then confirms the results, sorting the positives, negatives, false positives, and false negatives to help increase the accuracy of the classifier's predictions. When the trained classifier is published, it processes content in SharePoint Online, Exchange Online, and OneDrive for Business, and automatically classifies the content.

Applying sensitivity labels to documents and emails will embed metadata within the object which identifies the chosen sensitivity, thereby allowing the sensitivity to travel with the data. As a result, even if a labeled document is stored on a user's desktop or within an on-premise system, it is still protected. This in turn enables other Microsoft 365 solutions such as Microsoft Cloud App Security or network edge devices to identify sensitive data and automatically enforce security controls. Sensitivity labels have the added benefit of educating employees as to which data within an organization is considered sensitive, and how to handle that data should they receive it.

Office 365 Data Loss Prevention (DLP) will automatically identify documents, emails and conversations which contain sensitive data by scanning them for sensitive data types and then enforcing policy on those objects. Policies are enforced on documents within SharePoint and OneDrive for Business. They are also enforced when users send email, and in Microsoft Teams within chat and channel conversations. Policies may be configured to look for keywords, sensitive data types, retention labels and whether data is shared within the organization or externally. Controls are provided to help organizations fine-tune DLP policies to better avoid false positives. When sensitive data is found, customizable policy tips may be displayed to users within Microsoft 365 applications, informing them that their content contains sensitive data and/or proposing corrective actions. Policies can also prevent users from accessing documents, sharing documents, or sending emails which contain certain types of sensitive data. Microsoft 365 supports over 100 built-in sensitive data types, and organizations can configure custom sensitive data types to meet their policies.

Rolling out MIP and DLP policies to organizations requires careful planning and typically a user education program so that employees understand the organization's data classification schema and which types of data are considered sensitive. Providing employees with tools and education programs that help them identify sensitive data and help them understand how to handle it makes them part of the solution for mitigating information security risks.

Govern Data by Effectively Managing Records

Regulations require many organizations to manage the retention of key organizational documents according to a managed corporate retention schedule. Organizations face regulatory compliance risks if data is under-retained (deleted too early), or legal risks if data is over-retained (kept too long). Effective records management strategies help to ensure that organizational documents are retained according to predetermined retention periods which are designed to minimize risk to the organization. Retention periods are prescribed in a centrally managed organizational record retention schedule, and they are based on the nature of each type of document, on regulatory compliance requirements for retaining specific types of data, and on the defined policies of the organization.

Assigning record retention periods accurately across organizational documents may require a granular process which assigns retention periods uniquely to individual documents. The vast number of documents within energy industry organizations, coupled with the fact that in many cases retention periods can be triggered by organizational events (such as contracts expiring or an employee leaving the organization), make applying record retention policies at scale challenging for many organizations.

Microsoft 365 provides capabilities for defining retention labels and policies to easily implement records management requirements. A record manager defines a retention label, which represents a "record type" in a traditional retention schedule. The retention label contains settings which define:

- How long a record is retained for
- The concurrency requirements or what occurs when retention period expires (delete the document, start a disposition review, or take no action)
- What triggers the retention period to start (created date, last modified date, labeled date, or an event), and
- If the document or email is a record (meaning it cannot be edited or deleted)

Retention labels are then published to SharePoint or OneDrive sites, Exchange mailboxes, and Office 365 Groups. Users may then apply retention labels manually to documents and emails, or record managers can use rules to automatically apply retention labels. Auto-apply rules can be based on keywords or sensitive data found within documents or emails, such as credit card numbers, social security numbers or other personally identifiable information (PII) or they can be based on SharePoint metadata.

The FedRAMP Moderate Control Set and NERC CIP standards also include Asset Reuse and Disposal as a key control requirement (CIP-011-2). These requirements once again specifically address BES (Bulk Electric System) Cyber System Information. However, other jurisdictional regulations will require energy industry organizations

to manage and dispose of records effectively for numerous types of information. This will include financial statements, capital project information, budgets, customer data, etc. In all cases, energy organizations will be required to maintain robust records management programs and evidence related to the defensible disposition of corporate records.

With each retention label, Office 365 allows record managers to determine if a disposition review is required. Then when those record types come up for disposition, after their retention period has expired, a review must be conducted by the designated disposition reviewers before content is deleted. Once the disposition review is approved, content deletion will proceed, however evidence of the deletion, the user that performed the deletion and date/time in which it occurred is still retained for multiple years (as a certificate of destruction). If organizations require longer or permanent retention of certificates of destruction, Azure Sentinel may be used for long term cloud-based storage of log and audit data. Azure Sentinel gives organizations full control over the long term storage and retention of activity data, log data and retention/disposition data.

Comply with FERC and FTC Regulations for Energy Markets

The U.S. Federal Energy Regulatory Commission (FERC) oversees regulations related to energy markets and trading for the electric energy and natural gas markets. The U.S. Federal Trade Commission (FTC) oversees similar regulations in the petroleum market. In both cases these regulatory bodies set out rules and guidance to prohibit energy market manipulation. FERC, for example, recommends that energy organizations invest in technology resources to monitor trading, trader communications, and compliance with internal controls. They further recommend that energy organizations evaluate, on a regular basis, the ongoing effectiveness of the organization's compliance program.

Traditionally, communication monitoring solutions are costly, and they can be complex to configure and manage. Also, organizations can experience challenges with monitoring the numerous, varying communication channels available to employees. Microsoft 365 provides several built-in robust capabilities for monitoring employee communications, supervising employee activities, and helping to comply with FERC regulations for energy markets.

Implement Supervisory Control

Microsoft 365 enables organizations to configure supervision policies which capture employee communications (based on configured conditions) and allow these to be reviewed by designated supervisors. Supervision policies can capture internal/external email and attachments, Microsoft Teams chat and channel communications, Skype for Business Online chat communications and attachments, along with communications through third-party services (such as Facebook or Dropbox).

The comprehensive nature of communications that may be captured and reviewed within an organization, and the extensive conditions with which policies may be configured, allow Microsoft 365 Supervision Policies to help organizations comply with FERC energy market regulations. Supervision policies may be configured to review communications for individuals or groups. In addition, supervisors may be configured to be individuals or groups. Comprehensive conditions may be configured to capture communications based on inbound or outbound messages, domains, retention labels, keywords or phrases, keyword dictionaries, sensitive data types, attachments, message size, or attachment size. Reviewers are provided with a dashboard where they can review flagged communications, act on communications that potentially violate policies, or mark flagged items as resolved. They may also review the results of previous reviews and items that were have been resolved.

Microsoft 365 provides reports which allow supervision policy review activities to be audited based on the policy and the reviewer. The available reports may be used to validate that supervision policies are working as defined by the organizations written supervision policies. They may also be used to identify communications requiring review and which communications are not compliant with corporate policy. Finally, all activities related to configuring supervision policies and reviewing communications are audited in the Office 365 unified audit log.

Microsoft 365 Supervision Policies allow organizations to monitor communications for compliance with

corporate policies, such as human resources harassment violations and offensive language in company communications. It also allows organizations to reduce risk, by monitoring communications when organizations are undergoing sensitive organizational changes, such as mergers and acquisitions, or leadership changes.

Communication Compliance

With many communication channels available to employees, organizations increasingly require effective solutions for monitoring or supervising communications in regulated industries such as energy trading markets. The recently launched Communication Compliance solution built into Microsoft 365 helps organizations overcome common challenges such as increasing numbers of communication channels and message volume, as well as the risk of potential fines for policy violations.

Communication Compliance can monitor multiple communication channels and use machine learning models to identify potential policy violations, including Office 365 email, Microsoft Teams, Skype for Business Online, Facebook, Twitter and Bloomberg instant messages. Communication Compliance helps compliance teams effectively and efficiently review messages for potential violations of:

- Corporate Policies, such as acceptable use, ethical standards, and corporate specific policies
- sensitivity or sensitive business disclosures, such as unauthorized communications about sensitive projects like upcoming acquisitions, mergers, earnings disclosures, reorganizations, or leadership team changes
- Regulatory compliance requirements, such as employee communications regarding the types of businesses or transactions in which an organization engages in compliance with FERC regulations for energy markets

Communication Compliance provides built-in threat, harassment, and profanity classifiers to help reduce false positives when reviewing communications. This saves reviewers time during the investigation and remediation process. It helps reviewers focus on specific messages within long threads that have been highlighted by policy alerts. This helps compliance teams more quickly identify and remediate risks. It provides compliance teams with the ability to easily configure and fine-tune policies, adjusting the solution to the organization's specific needs and reducing false positives. Communication Compliance can also track user behavior over time, highlighting potential patterns in risky behavior or policy violations. Finally, it provides flexible built-in remediation workflows so that reviewers can quickly take action and escalate to legal or human resources teams according to defined corporate processes.

Protect Against Data Exfiltration and Insider Risk

A common threat to enterprises is data exfiltration, or the act of extracting data from an organization. This can be a significant concern for energy organizations due to the sensitive nature of the information that may be accessed by employees or field service staff day to day. This includes both BES (Bulk Electric System) Cyber System information as well as business related information and customer data. With the increasing methods of communications available and numerous tools for moving data, advanced tools are typically required to mitigate risks of data leaks, policy violations and insider risk.

Insider Risk Management

Enabling employees with online collaboration tools that may be accessed anywhere inherently brings risk to an organization. Employees may inadvertently or maliciously leak data to attackers or to competitors. Alternatively, they may exfiltrate data for personal use or take data with them to a future employer. These scenarios' present serious risks to organizations from a security and a compliance standpoint. Identifying these risks when they occur and quickly mitigating them requires both intelligent tools for data collection and collaboration across departments such as legal, human resources, and information security.

Microsoft 365 has recently launched the Insider Risk Management console which uses signals across Microsoft 365 services and machine learning models to monitor user behavior for signs of insider risk. This tool presents data to investigators so that they can easily identify risky behavioral patterns and escalate cases based on pre-determined workflows.

For example, Insider Risk Management can correlate signals from a user's Windows 10 desktop, such as copying

files to a USB drive or emailing a personal email account, with activities from online services such as Office 365 email, SharePoint Online, Microsoft Teams, OneDrive for Business, etc. to identify data exfiltration patterns. It can also correlate these activities with employees leaving an organization which is a common behavioral pattern associated with data exfiltration. It can monitor multiple activities and behavior over time, and when common patterns emerge, it can raise alerts and help investigators focus on key activities to verify a policy violation with a high degree of confidence. Insider Risk Management can also obfuscate data from investigators to help meet data privacy regulations, while still surfacing key activities that help them efficiently perform investigations. When ready, it allows investigators to package and securely send key activity data to human resources and legal departments, following common escalation workflows for raising cases for remediation action.

Insider Risk Management is a significant increase in capabilities in Microsoft 365 for monitoring and investigating insider risks, while allowing organizations to still meet data privacy regulations and follow established escalations paths when cases require higher-level action.

Conclusion

Microsoft 365 provides an integrated and comprehensive solution which enables easy-to-use cloud-based collaboration across the enterprise with Microsoft Teams. Microsoft Teams also enables better communication and collaboration with field service staff, helping energy organizations to be more efficient and effective. Better collaboration across the enterprise and with field staff can ultimately help energy organizations to better serve customers.

Energy industry organizations must comply with strict regulations related to how they store, secure, manage, and retain information related to their operations and customers. They must also comply with regulations related to how they monitor and prevent the manipulation of energy markets. Microsoft 365 provides robust security controls for protecting data, identities, devices, and applications from risks and complying with strict energy industry regulations. Built-in tools are provided to help energy organizations assess their compliance, as well as take action and track remediation activities over time. These tools also provide easy to use methods for monitoring and supervising communications. The Microsoft 365 platform is built on foundational components like Microsoft Azure and Azure Active Directory, helping to secure the overall platform and helping the organization meet compliance requirements for FedRAMP Moderate and High control sets. This in turn contributes to an energy organization's ability to meet NERC CIP standards.

Overall, Microsoft 365 helps energy organizations to better protect the organization, to have more robust compliance programs, and to enable staff to focus on gaining better insights and implementing strategies to better reduce risk.

Exchange Online mail encryption with AD RMS

11/2/2020 • 3 minutes to read • Edit Online

To help prevent information leakage, Exchange Online includes Information Rights Management (IRM) functionality that provides online and offline protection of email messages and attachments. You can configure Exchange Online IRM to use on-premises Active Directory Rights Management Service (AD RMS), if needed, to satisfy your organization requirements. This is not common. If you do not have a requirement to use AD RMS, use Office 365 Message Encryption instead.

IRM protection can be applied by users in Microsoft Outlook or Outlook on the web, and it can be applied by administrators using transport protection rules or Outlook protection rules. IRM helps you and your users control who can access, forward, print, or copy sensitive data within an email.

Changes to how IRM works with Office 365 Message Encryption (OME) and Azure Active Directory

As of September 2017, when you set up the new Office 365 Message Encryption capabilities for your organization, you also set up IRM for use with Azure Rights Management (Azure RMS). You no longer set up IRM with Azure RMS separately. Instead, OME and rights management work seamlessly together. For more details about the new capabilities, see Office 365 Message Encryption FAQ. If you're ready to get started using the new OME capabilities within your organization, see Set up new Office 365 Message Encryption capabilities built on top of Azure Information Protection.

How IRM works with Exchange Online and Active Directory Rights Management Services

Exchange Online IRM uses on-premises Active Directory Rights Management Services (AD RMS), an information protection technology in Windows Server 2008 and later. IRM protection is applied to email by applying an AD RMS rights policy template to an email message. Rights are attached to the message itself so that protection occurs online and offline and inside and outside of your organization's firewall.

Users can apply a template to an email message to control the permissions that recipients have on a message. Actions, such as forwarding, extracting information from a message, saving a message or printing a message can be controlled by applying an AD RMS rights policy to the message.

You can configure IRM to use an AD RMS server running Windows Server 2008 or later. You can use this AD RMS server to manage the AD RMS rights policy templates for your cloud-based organization. Outlook also relies on the AD RMS server to enable users to apply IRM protection to messages they send. For details, see Configure IRM to use an on-premises AD RMS server.

After it's enabled, IRM protection can be applied to messages as follows:

- **Users can manually apply a template using Outlook and Outlook on the web.** Users can apply an AD RMS rights policy template to an email message by selecting the template from the **Set permissions** list. When users send an IRM-protected message, any attached files that use a supported format also receive the same IRM protection as the message. IRM protection is applied to files associated with Word, Excel, and PowerPoint, as well as .xps files and attached email messages.
- **Administrators can use transport protection rules to apply IRM protection automatically to both Outlook and Outlook on the web.** You can create transport protection rules to IRM-protect messages. Configure the transport protection rule action to apply an AD RMS rights policy template to

messages that meet the rule condition. After you enable IRM, your organization's AD RMS rights policy templates are available to use with the transport protection rule action called **Apply rights protection to the message with**.

- **Administrators can create Outlook protection rules.** Outlook protection rules automatically apply IRM-protection to messages in Outlook 2010 (not Outlook on the web) based on message conditions that include the sender's department, who the message is sent to, and whether recipients are inside or outside your organization. For details, see Create an Outlook Protection Rule.

Configure IRM to use an on-premises AD RMS server

5/8/2020 • 6 minutes to read • Edit Online

For use with on-premises deployments, Information Rights Management (IRM) in Exchange Online uses Active Directory Rights Management Services (AD RMS), an information protection technology in Windows Server 2008 and later. IRM protection is applied to email by applying an AD RMS rights policy template to an email message. Rights are attached to the message itself so that protection occurs online and offline and inside and outside of your organization's firewall.

This topic shows you how to configure IRM to use an AD RMS server. For information about using the new capabilities for Office 365 Message Encryption with Azure Active Directory and Azure Rights Management, see the Office 365 Message Encryption FAQ.

To learn more about IRM in Exchange Online, see Information Rights Management in Exchange Online.

What do you need to know before you begin?

- Estimated time to complete this task: 30 minutes
- You need to be assigned permissions before you can perform this procedure or procedures. To see what permissions you need, see the "Information Rights Management" entry in the Messaging policy and compliance permissions topic.
- The AD RMS server must be running Windows Server 2008 or later. For details about how to deploy AD RMS, see Installing an AD RMS Cluster.
- For details about how to install and configure Windows PowerShell and connect to the service, see Connect to Exchange Online Using Remote PowerShell.
- For information about keyboard shortcuts that may apply to the procedures in this topic, see Keyboard shortcuts for the Exchange admin center in Exchange Online.

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

How do you do this?

Step 1: Use the AD RMS console to export a trusted publishing domain (TPD) from an AD RMS server

The first step is to export a trusted publishing domain (TPD) from the on-premises AD RMS server to an XML file. The TPD contains the following settings needed to use RMS features:

- The server licensor certificate (SLC) used for signing and encrypting certificates and licenses
- The URLs used for licensing and publishing
- The AD RMS rights policy templates that were created with the specific SLC for that TPD

When you import the TPD, it's stored and protected in Exchange Online.

1. Open the Active Directory Rights Management Services console, and then expand the AD RMS cluster.
2. In the console tree, expand **Trust Policies**, and then click **Trusted Publishing Domains**.
3. In the results pane, select the certificate for the domain you want to export.
4. In the **Actions** pane, click **Export Trusted Publishing Domain**.
5. In the **Publishing domain file** box, click **Save As** to save the file to a specific location on the local computer. Type a file name, making sure to specify the `.xml` file name extension, and then click **Save**.
6. In the **Password** and **Confirm Password** boxes, type a strong password that will be used to encrypt the trusted publishing domain file. You will have to specify this password when you import the TPD to your cloud-based email organization.

Step 2: Use the Exchange Management Shell to import the TPD to Exchange Online

After the TPD is exported to an XML file, you have to import it to Exchange Online. When a TPD is imported, your organization's AD RMS templates are also imported. When the first TPD is imported, it becomes the default TPD for your cloud-based organization. If you import another TPD, you can use the **Default** switch to make it the default TPD that is available to users.

To import the TPD, run the following command in Windows PowerShell:

```
Import-RMSTrustedPublishingDomain -FileData $([byte[]](Get-Content -Encoding byte -Path <path to exported
TPD file> -ReadCount 0)) -Name "<name of TPD>" -ExtranetLicensingUrl <URL> -IntranetLicensingUrl <URL>
```

You can obtain the values for the *ExtranetLicensingUrl* and *IntranetLicensingUrl* parameters in the Active Directory Rights Management Services console. Select the AD RMS cluster in the console tree. The licensing URLs are displayed in the results pane. These URLs are used by email clients when content has to be decrypted and when Exchange Online needs to determine which TPD to use.

When you run this command, you'll be prompted for a password. Enter the password that you specified when you exported the TPD from your AD RMS server.

For example, the following command imports the TPD named Exported TPD using the XML file that you exported from your AD RMS server and saved to the desktop of the Administrator account. The Name parameter is used to specify a name to the TPD.

```
Import-RMSTrustedPublishingDomain -FileData $([byte[]](Get-Content -Encoding byte -Path
C:\Users\Administrator\Desktop\ExportTPD.xml -ReadCount 0)) -Name "Exported TPD" -ExtranetLicensingUrl
https://corp.contoso.com/_wmcs/licensing -IntranetLicensingUrl https://rmsserver/_wmcs/licensing
```

For detailed syntax and parameter information, see Import-RMSTrustedPublishingDomain.

How do you know this step worked?

To verify that you have successfully imported the TPD, run the **Get-RMSTrustedPublishingDomain** cmdlet to retrieve TPDs in your Exchange Online organization. For details, see the examples in Get-RMSTrustedPublishingDomain.

Step 3: Use the Exchange Management Shell to distribute an AD RMS rights policy template

After you import the TPD, you must make sure an AD RMS rights policy template is distributed. A distributed template is visible to Outlook on the web (formerly known as Outlook Web App) users, who can then apply the templates to an email message.

To return a list of all templates contained in the default TPD, run the following command:

```
Get-RMSTemplate -Type All | fl
```

If the value of the *Type* parameter is `Archived`, the template isn't visible to users. Only distributed templates in the default TPD are available in Outlook on the web.

To distribute a template, run the following command:

```
Set-RMSTemplate -Identity "<name of the template>" -Type Distributed
```

For example, the following command imports the Company Confidential template.

```
Set-RMSTemplate -Identity "Company Confidential" -Type Distributed
```

For detailed syntax and parameter information, see Get-RMSTemplate and Set-RMSTemplate.

The Do Not Forward template

When you import the default TPD from your on-premises organization into Exchange Online, one AD RMS rights policy template named **Do Not Forward** is imported. By default, this template is distributed when you import the default TPD. You can't use the **Set-RMSTemplate** cmdlet to modify the **Do Not Forward** template.

When the **Do Not Forward** template is applied to a message, only the recipients addressed in the message can read the message. Additionally, recipients can't do the following:

- Forward the message to another person.
- Copy content from the message.
- Print the message.

> **IMPORTANT**
> The **Do Not Forward** template can't prevent information in a message from being copied with third-party screen capture programs, cameras, or users manually transcribing the information

You can create additional AD RMS rights policy templates on the AD RMS server in your on-premises organization to meet your IRM protection requirements. If you create additional AD RMS rights policy templates, you have to export the TPD from the on-premises AD RMS server again and refresh the TPD in the cloud-based email organization.

How do you know this step worked?

To verify that you have successfully distributed and AD RMS rights policy template, run the **Get-RMSTemplate** cmdlet to check the template's properties. For details, see the examples in Get-RMSTemplate.

Step 4: Use the Exchange Management Shell to enable IRM

After you import the TPD and distribute an AD RMS rights policy template, run the following command to enable IRM for your cloud-based email organization.

```
Set-IRMConfiguration -InternalLicensingEnabled $true
```

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

How do you know this step worked?

To verify that you have successfully enabled IRM, run the Get-IRMConfiguration cmdlet to check IRM

configuration in the Exchange Online organization.

How do you know this task worked?

To verify that you have successfully imported the TPD and enabled IRM, do the following:

- Use the **Test-IRMConfiguration** cmdlet to test IRM functionality. For details, see "Example 1" in Test-IRMConfiguration.
- Compose a new message in Outlook on the web and IRM-protect it by selecting **Set permissions** option from the extended menu (•••).