

SC-400 Visual Study Guide

For the Microsoft Certified: Information Protection
and Compliance Administrator Associate Exam



Dominique Hermans

Introduction

Hi there!

Thank you for taking the time to download my Visual Study Guide for the SC-400 Microsoft Certified: Information Protection and Compliance Administrator Associate Exam.

When I started studying for the SC-400 exam last year, I noticed a gap in the available study materials. There were few clear, practical examples of how to configure the various solutions within the Purview platform and, more importantly, what the impact would be on end users. I needed to see what these configurations looked like in practice to fully understand and retain the information.

With this in mind, I began writing this study guide at the start of 2024. It's packed with screenshots (hence, 'visual study guide') and step-by-step explanations on configuring each solution in the Purview platform, with a strong focus on the end-user experience and impact. Creating this guide helped me deepen my understanding of the platform, and I found that visualizing the steps made it easier to remember. Now, I'm excited to share this knowledge with you.

I'm always open to feedback, so feel free to reach out using my contact information below.

Have fun reading!



Dominique Hermans

 <https://www.linkedin.com/in/dominiquehermans1/>

 <https://dominiquehermans.com/>

 dominiquehermans.bsky.social



Table of contents

Introduction.....	2
Disclaimer.....	6
Managing Data Security, Compliance and Governance in Microsoft 365 with Microsoft Purview	7
Overview – The big picture.....	7
Tools of the trade	8
Data Classifiers Explained.....	12
Introduction to Data Classifiers in Purview	12
Sensitive Information Types (SIT's)	12
Fingerprint based SIT's.....	14
Exact Data Match (EDM) Classification.....	17
Trainable Classifiers	25
Where can I apply all this goodness?	26
Data Lifecycle Management (DLM).....	28
Terminology.....	28
Configuring Retention Policies	28
Configuring (Retention) Labels	31
Publishing (Retention) Labels.....	33
The user experience.....	35
How to see where labels are applied?.....	36
How to see where Retention Policies are applied?	38
Records Management (RM)	40
Differences between Records Management (RM) and Data Lifecycle Management (DLM)....	40
File Plans	41
Configuring Records Management	41
The view from the user / Records Manager.....	46
Fast forward a couple of days.....	48
Differences between Retention Labels, Records and Regulatory Records	49
How to see where labels are applied?.....	50
Sensitivity Labels	51
Introduction	51
Configuring Sensitivity Labels	52
The view from the user	59
In conclusion	65
Data Loss Prevention (DLP)	66
Plan first, implement second	66

Configure Data Loss Prevention (DLP)	67
A detour into Sensitive Information Types (SIT's)	70
Back to configuration of our DLP policy	72
Reviewing simulation results.....	73
Enabling the policy	75
In conclusion	80
Adaptive Scopes	81
eDiscovery (Premium)	85
A quick note on licenses	85
eDiscovery Roles.....	86
eDiscovery Workflow	87
Analytics	95
Document Review	95
Communication Compliance	99
Communication Compliance Policy Configuration.....	99
Testing, 1, 2, 3	101
Communication Policy Match Overview	102
Notify the user.....	103
Message Removal.....	104
Resolve a message	105
Optical Character Recognition (OCR)	106
Bonus: Quick Setup: User-reported messages & Inappropriate Content	106
Insider Risk Management (IRM).....	108
Prerequisites	108
Setting up a policy	111
Good user gone bad	117
Red alert!	118
Notice templates.....	121
Case actions	122
Mastering the (Unified) Audit Log	123
Let's talk basics.....	123
What information is stored in the Unified Audit Log?	123
Using Audit Log Search	124
Utilizing the Content and Activity Explorer	127
Content Explorer	127
Activity Explorer.....	129

Configuring Alert Policies for High Risk Activities	131
A word on RBAC Permissions	131
Alert Policies Overview	131
Configuring Alert Policies	131
Triggering the alert policy	134
Examining the results	135
Summary	137
Information Barriers	138
Setting the scene.....	138
Information Barriers versioning	138
Pre-Information Barriers behavior	138
Setting up Information Barrier Basics – Segments	139
Setting up Information Barrier Basics – Policies	141
Setting Up Information Barrier Basics – Applying Policies	143
The user Experience – Teams	143
Setting up Information Barriers – Sharepoint Configuration	145
Setting up Information Barriers – OneDrive Configuration.....	149
Compliance Manager	153
Basics first, as always	153
Compliance Manager, here we come!.....	153
How to create an assessment based on a regulation	154

Disclaimer

The information contained in this ebook may be used freely, provided that proper credit is given to the author, Dominique Hermans, and his website, <https://dominiquehermans.com>.

The author and publisher make no representations or warranties with respect to the accuracy or completeness of the contents of this ebook and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate.

Neither the author nor the publisher shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

Please do note that this guide is not a replacement for other study materials but should be used as a supplement.

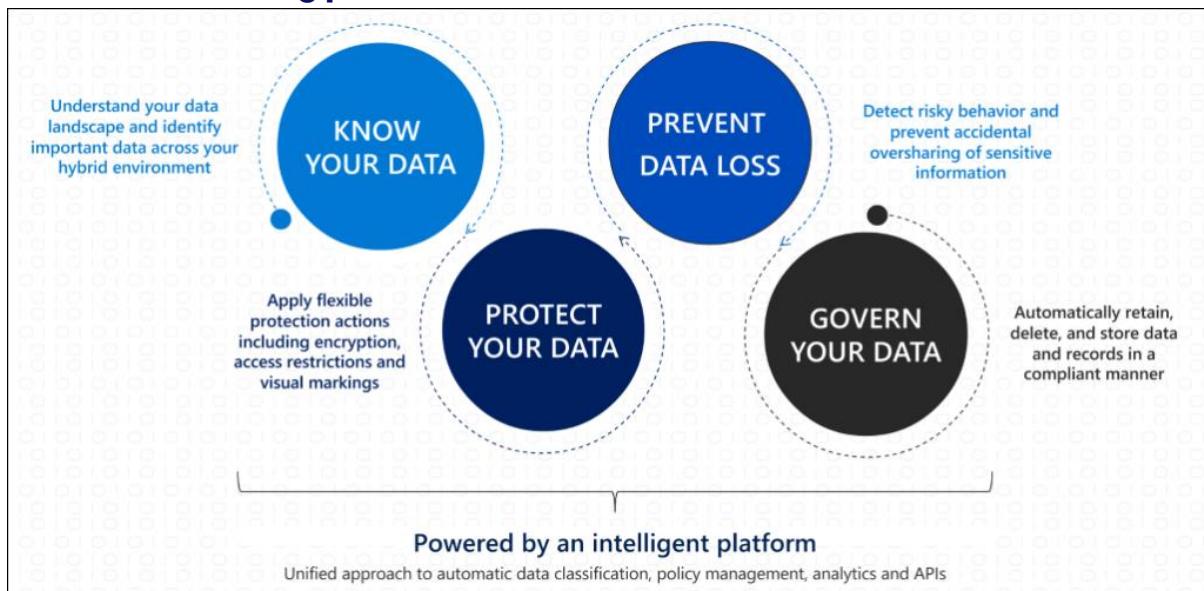
Microsoft 365, Microsoft Purview, and all affiliated products mentioned in this ebook are owned by Microsoft.

Managing Data Security, Compliance and Governance in Microsoft 365 with Microsoft Purview

Microsoft Purview. I'm certain it's features can seem a little overwhelming when first strolling through the documentation of the product. It sure was for me. In this ebook I want to give you insight in Microsoft's product for managing data security, compliance and governance for the Microsoft 365 suite, Microsoft Purview.

First things first. Microsoft Purview is a unified solution that used to consist of 2 products: Azure Purview and Microsoft 365 Compliance solutions. In this ebook the focus is on the Microsoft 365 side of the product.

Overview – The big picture



Know your data, protect your data, prevent data loss & govern your data | [Source](#)

Before I start explaining the different features that Microsoft Purview consists of, it's good to know where all the tools take their place in the "big picture" aka the Purview landscape. To implement a data security, compliance and governance strategy you'll have to walk through various steps.

Know your data

First you'll have to assess your data landscape. Where is your important data? Where is the data that poses possible privacy risks for example, and where is other sensitive data stored? These items have to be identified before you can continue to the next step. Often this step is executed by stakeholders that know their part of the data landscape well. The Purview solution can manage both cloud and hybrid environments. Want to know how your data environment matches up with current data privacy controls or standards? This is the place (or step) to be.

Protect your data

Second, data that's identified to be subject to data privacy regulation needs to be protected whether it is in the cloud or on-premise. Examples of these kinds of laws or regulatory controls in the Netherlands are NEN, ISO or BIO. This data (and of course other data you want) can be protected by applying labeling, which in turn makes it possible to apply an action based on the label your content has.

Prevent data loss

To prevent data loss, risky behavior and accidental oversharing of sensitive information needs to be prevented.

Govern your data

“Keep what you need and delete what you don’t” is a great tagline for data governing. Governing your data can be done by applying labels that make sure data isn’t being deleted if it’s supposed to be retained for a certain period or it can be deleted within a certain timeframe if it’s not meant to be retained for a certain period.

Tools of the trade

Now that we’re aware of the Purview landscape, let’s take a look at our “tools of the trade” and match them with each of the 4 categories above.

Know your data

Knowing your data starts with a non-technical process to identify where your data is stored as mentioned above. When you have gathered information on the storage locations (Which can be in the Microsoft 365 cloud, different Azure components or on-premise by using an agent), you can:

- Run a risk assessment and obtain a compliance score by using the compliance manager. The compliance manager matches your environment with various pre-defined assessments to give you insight in how compliant you are in relation to certain regulatory controls like ISO or NEN and what you have to do to improve your posture. These assessments are made available by Microsoft, but of course you can also create your own.
- You can set up sensitive information types or custom sensitive information types to classify data in your environment in various ways. An example is the use of pattern matching by using regular expressions.
- Trainable classifiers are an advanced way to identify data. It does not use pattern matching, but you can feed it samples of your data. It will then use AI to train itself and get better at matching your content.
- Content Explorer can be used to view an overview of your items that have a sensitivity label, retention label or have been classified as a sensitive information type.
- Activity Explorer takes a look at your audit logs and uses this information to show what actions are being taken on your labeled content.

Protect your data

Protecting your data can be done by:

- Applying Sensitivity Labels. These are labels that attach to your documents in clear text in the metadata of your files so they can be read by someone that might not have access to the content in the document. Actions that can be applied to labeled documents are “encrypt”, “mark content” or they can be used to protect content in containers like sites and groups. The Azure Information Protection Unified labeling client can be used to extend labeling to file explorer and powershell. To identify data on-premise, the Information Protection Scanner can be used. It will present information of your on-premise data in a unified view within the Microsoft Purview portal. Want to protect data residing in third party apps? Take a look at the Microsoft Information Protection SDK.
- Connecting on-premise Exchange, SharePoint or file-servers that run File Classification Infrastructure (FCI) for protection using Rights Management Connector.
- Implementing Office Message Encryption. This encrypts email messages (and attachments) so only authorized recipients can read the information.
- Implementing Teams and SharePoint Access Controls. Use access controls in both Teams and SharePoint to prevent oversharing of information.
- Setting up SharePoint Information Rights Management (IRM) to protect checked out documents.
- Using Microsoft Defender for Cloud Apps to label sensitive information in cloud data stores.
- Extending your labeling needs to Microsoft Purview Data Map Assets. A few examples of these assets are Azure Data Lake, Azure Files, Azure SQL DB and Azure Cosmos DB.

Prevent Data Loss

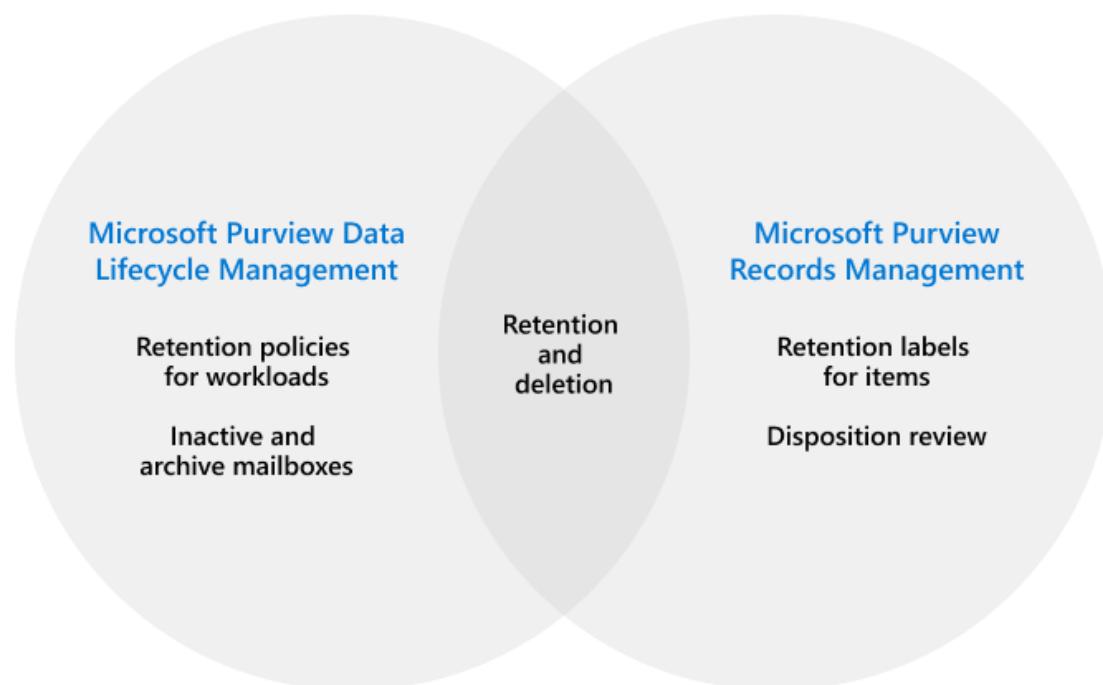
Microsoft Purview Data Loss Prevention (DLP) can be used to help unintentional sharing of sensitive items. These capabilities can be extended to:

- Browsers like Chrome and Firefox with browser extensions.
- Windows 10 endpoints by using Endpoint data loss prevention.
- Your on-premises file shares and SharePoint Document Libraries with Microsoft Purview DLP on-premises repositories.
- Microsoft Teams chat and channel messages.

Govern your Data

Governing your data within the Microsoft 365 ecosystem is done using 2 solutions. “Microsoft Purview Data Lifecycle Management” and “Microsoft Purview Records Management”. The following image shows the overlap in both solutions:

Govern your data



Data Governing Tools | [Source](#)

Remember that data governing was all about “keeping what you need and deleting what you don’t”? Here are your tools of the trade for this to happen:

Microsoft Purview Data Lifecycle Management

- Retention policies for Microsoft 365 workloads let you retain or delete content. This can be done for email, documents, Teams and Viva Engage messages.
- Inactive mailboxes lets you keep mailbox content after an employee leaves the organization.
- (Online) Archive mailboxes can be used to give users additional mailbox space.
- Import service for PST files. No explanation needed here.

Microsoft Purview Records Management

- A File plan can be used to bulk create retention labels or export information from existing labels. File plan descriptors support additional information for each label as metadata.
- Retention policies can be used to assign the same retention settings for content at a site or mailbox level, Retention labels can be assigned at item-level. See the first bullet under “Microsoft Purview Data Lifecycle Management” for more information.
- Event-driven retention gives you the capability to automatically implement labels and actions based on certain events.

- Disposition Review gives you the ability to manually review content before it's permanently deleted. Proof of disposition of records is included.

Investigate (Bonus step)

After (all of) the above is in place, there are a few features to keep in mind:

- Auditing and alert policies combined with reporting can be used to keep being informed about your environment.
- Data subject requests (by for example regulatory bodies) can be met by the content search and eDiscovery features.

Are you still with me? Good. I reckon this is a large list of features that take some time to take in, but I hope the tools make more sense now that they're categorized in the know, protect, prevent and govern categories.

Data Classifiers Explained

When talking about Microsoft Purview, it goes often hand in hand with data classification. But how can we classify our data? Microsoft Purview provides us with a few different options to do this:

1. Manually by an administrator or your users.
2. By using automated pattern-matching.
3. By using classifiers

After your data has been classified, you can take a closer look at where your sensitive data resides (for example with Data Explorer or eDiscovery) to get an overview of your information, or use the various tools in Purview to protect your sensitive data. This chapter covers classifying your data, if you want to take a closer look at the tools that protect your data, take a look at the chapters on Data Loss Prevention and Data Lifecycle Management.

Introduction to Data Classifiers in Purview

When talking about manually categorizing your content you can use pre-existing labels or sensitive information types or you can create custom ones yourself and use these to protect your data and manage it's lifecycle.

Automated pattern matching can find content automatically using the following techniques:

- Matching of keyword or metadata values using the Kusto Query Language (KQL)
- Previously identified patterns of sensitive information. Examples are credit card numbers or social security numbers which can be identified using a specific pattern. In Purview terms, this is called a Sensitive Information Type (or SIT in brief). At time of writing, 324 SIT's are preconfigured to identify sensitive information ranging from Azure Storage Account Shared Access Signatures to U.S. Physical Addresses. If these don't suit your needs, you can also create custom SIT's.
- A variation on the SIT's above are Fingerprint based SIT's, which can identify items when they are based on a template.
- Exact Data Match which searches for -you guessed it- exact data (strings) in items.

If your items aren't easily identifiable by one of the methods above, a classifier can be used as a categorization method to identify items on what the item is, instead of matching elements in the item which is done with pattern matching. These come in two flavors: pretrained classifiers that are provided by Microsoft and trainable classifiers, which can be used when pretrained classifiers don't match your needs. Trainable classifiers use AI and machine learning to recognize items.

Let's dive in and take a look at how each of these classifiers can be configured!

Sensitive Information Types (SIT's)

Let's navigate to the Purview Portal, select 'Data Classification' and 'Classifiers'. On the right hand side, let's go for 'Sensitive Info types'. Search the list for 'Credit Card Number' and open it. Now let's create a copy so we can edit it's properties to take a look at it's configuration.

Define patterns for this sensitive info type

Sensitive info types are defined by one or more patterns. Each pattern must contain a primary element and confidence level, but you can also include supporting elements and additional checks to further refine the evaluation and detection of matching items. [Learn about defining patterns](#)

Name	Confidence level
Pattern #1	High
Primary element	Function processors: Func_credit_card
Character proximity	Detect primary AND supporting elements within 300 characters
Supporting elements	<ul style="list-style-type: none"> Keyword list: Keyword_cc_verification Keyword list: Keyword_cc_name Function processors: Func_expiration_date
Pattern #2	Low
Primary element	Function processors: Func_credit_card
Character proximity	Detect primary AND supporting elements within 300 characters

As can be seen in the image above, the 'Credit Card Number' SIT checks items for 2 patterns:

- Pattern #1, which is considered a high confidence level when matched.
- Pattern #2, which is considered a low confidence level when matched.

Now let's dive in the details of Pattern #1. When editing the pattern (I will spare you all the screenshots), the following can be examined:

- Its primary element is to use a function called 'func_credit_card' which is a function written by Microsoft that looks for Credit Card Numbers. The number found must pass a test in the function so the function can be sure that it is indeed a Credit Card Number.
- In addition to the above primary element, it searches the 300 characters before and after the found Credit Card Number for an supporting element. Which in this case, is 1 of the following:
 - A keyword from the 'keyword_cc_verification' list which contains a list of variations of the word 'card identification number'.
 - A keyword from the 'keyword_cc_name' list which contains a list of Credit Card companies.
 - A function that looks for an expiration date.

When the logic from the above pattern is true and thus the primary element is found AND is accompanied by 1 of the 3 supporting elements, it's considered a match with a high confidence level.

Now, what about Pattern #2? It only uses the ‘func_credit_card’ function to look for a Credit Card Number but doesn’t use any supporting elements. When the logic from Pattern2 matches it’s also considered a match, albeit with low confidence.

When building your own SIT, all of the elements in the examples above can be manually configured and created to match your needs. Ain’t that great?

Fingerprint based SIT's

A fingerprint based SIT is somewhat different than your regular SIT. Where the regular version looks for certain elements in a document, the fingerprint based SIT has knowledge of a certain template that is being used for a certain document type.

Let’s go with an example here. Say your organization has a document template that it uses for all their invoices. You are absolutely sure that this template is used for every invoice that is being sent out the door. You can use the following procedure to create a document fingerprint of the template that you use for your invoices in this case. Let’s configure this.



Invoice from a fictional company

A demo document by DominiqueHermans.com



Invoice

Start by creating a docx file from your template if it exists as an dotx file. The one I use can be seen in the picture above. This is needed cause the dotx file type is not supported. Purview creates it's own small XML file that acts as the document fingerprint, or template if you will.

Now let's navigate to the purview portal and navigate to 'Data Classification' and 'Classifiers'. On the right hand side, let's go for 'Sensitive Info types'. Click 'Create Fingerprint based SIT'.

Name your fingerprint based SIT

This name and description will appear in compliance policies that support sensitive info types, so be sure to enter text that helps admins easily understand what info will be detected.

Name *

Description *

Give your SIT a name and description and click next.

Upload a file to create a fingerprint for the file

Fingerprint based SITs are based on the fingerprint of a file. Please upload file for which you would like to create a fingerprint.

File name

Template.docx 1 trash

Upload file

Confidence Level 2

Low

30 & Above

Medium

50 & Above

High

80 & Above

Now, select your docx file that contains the template. Confidence levels may be adjusted if you want. The numbers here resemble the percentage of text that has to be available in the document you encounter and match against the fingerprint based SIT. So, a 30% or above text match generates a low confidence level, a 50% or above match generates a medium confidence level and so on.

Classifiers > Sensitive info types > DominiqueHermans.com Invoice FSIT

DominiqueHermans.com Invoice FSIT

[Overview](#) [Matched items](#)

Recommendation

Provide feedback to improve the SIT

Help us improve the accuracy of the sensitive information types by letting us know if detected items match this classifier or not and submitting the document to Microsoft.

[Provide feedback](#)
[Test](#) [Copy](#) [Edit](#) [Delete](#)

Feedback results

- 0 Items with feedback
- 0 Match
- 0 Not a match

Details

Description

Invoice template used by a fictional company.

Confidence level

Low: 30

Medium: 50

High: 80

Created by

[onmicrosoft.com](#)

Let's test our fingerprint based SIT. Navigate to Sensitive Info Types, scroll down and select your newly generated fingerprint based SIT. Click 'Test' and upload a document that has been created based on your template.

Match results

We have detected the following in [Invoice 0001.docx](#)

1. DominiqueHermans.com Invoice FSIT
Low - 1 unique matches
2. DominiqueHermans.com Invoice FSIT
Medium - 1 unique matches
3. DominiqueHermans.com Invoice FSIT
High - 1 unique matches

It will now show you the match results of the file you uploaded against the created fingerprint based SIT.

Exact Data Match (EDM) Classification

Let's move on to Exact Data Match (EDM) Classification. The previous classifiers we've talked about all use a mechanism to match items using a certain technique, being able to get a closest match as possible when classifying data. But, if you are sure that certain elements reside in your documents, why not let it classify data using an exact match? You guessed it, that's where Exact Data Match (EDM) comes into play.

Exact data match classification

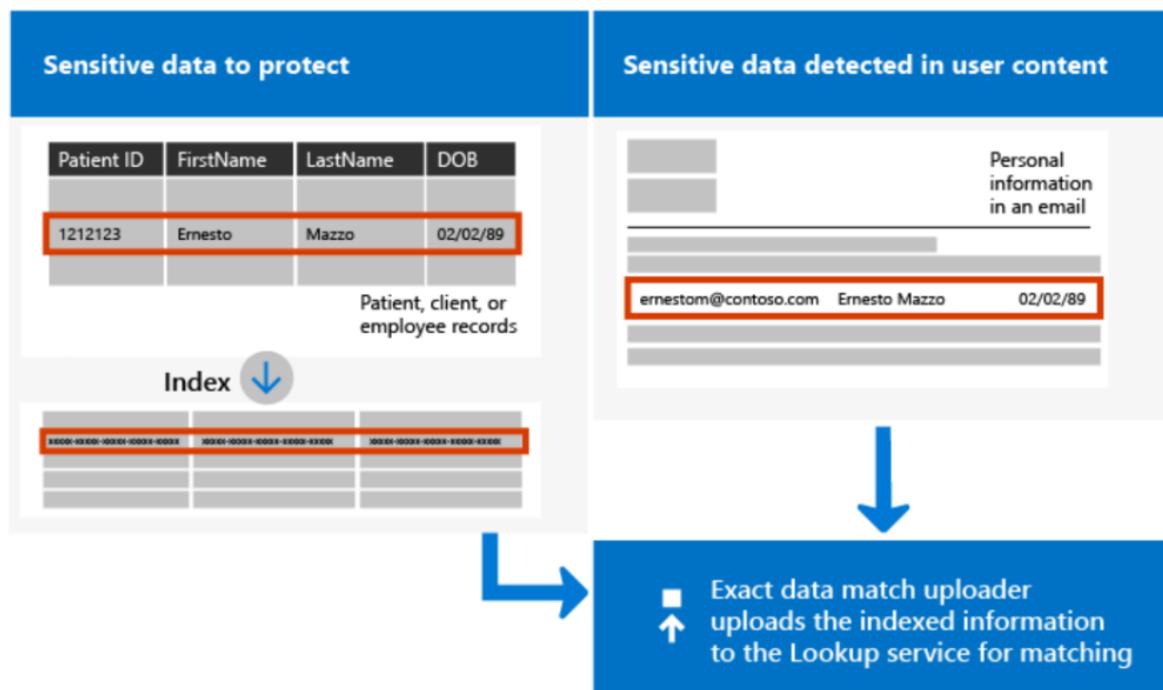


Image Source: [Microsoft](#)

With EDM, you feed Purview with a database of information that it should look for in your documents. In the example above, a list of employee records is uploaded into Purview so the service can recognize files that are present in the database. Let's take a look at the process involved.

Familiarize yourself with the steps needed to put your classifier to work

X

1. Prerequisite: Discover and prepare your sensitive data

OUTSIDE COMPLIANCE PORTAL

Before creating your EDM classifier, you'll prepare two files...one's required, the other's highly recommended.



- Required. A file containing the actual sensitive data you want your classifier to detect. For example, if you want to detect patient records, your file might contain data for "Patient ID" and "Name".
- Highly recommended. A similar file with sample data that will be used when creating the EDM classifier in the next step.

Not sure how to set these files up? [Learn how to prepare your data](#)

2. Create an EDM classifier

WITHIN COMPLIANCE PORTAL

Click "Create EDM classifier" below to open a wizard that will walk you through the steps. Process at a glance:



- Choose a method for defining the schema that's used to detect your data (we recommend uploading a file with sample data).
- Map that data to existing sensitive info types.
- Set up rules that control exactly what info will be detected in your org's content.

[Learn more about these steps](#)

3. Securely upload the file containing your org's sensitive info

OUTSIDE COMPLIANCE PORTAL

After creating the classifier, use the EDM Upload Agent tool to hash and upload the file containing your org's data. For greater security, we recommend using different computers to hash and upload separately. [Learn how to upload your data](#)



4. Test the classifier in simulation mode and publish it

WITHIN COMPLIANCE PORTAL

After the classifier is connected to your org's data file, there are a couple ways to test it out before including it in policies.



- Select the classifier from the 'Sensitive info types' page, choose 'Test', then upload a sample doc to check whether the classifier detects the elements you specified.
- Create a sensitivity auto-labeling policy that detects content matching the classifier. Run the policy in simulation mode to review matching items in your org to see if the label would be applied to the right content. As you review, you can refine the classifier and run simulation again to improve accuracy.

[Learn more about simulation mode](#)

[Create EDM classifier](#)
[Cancel](#)

Image Source: Microsoft

Creating an EDM classifier involves navigating through different steps. They are neatly outlined by Microsoft in the image above, which you are presented with when creating EDM classifiers. Let's walk through the steps one by one.

Discover and prepare your sensitive data

A	B	C	D	E	F	G	H	I
1	SSN	IndividualID	AccountID	MRN	FirstName	LastName	DOB	
2	270 34 7884	A12345678	B12345678	AB12345678	Richard	Jackson	01/01/1990	12065254152 richard@contoso.onmicrosoft.com
3	030 72 7381	A12345678	B12345678	80002910	Sarah	Jefferson	02/02/1985 1-206-234-7492	sarah.jefferson@contoso.com
4	757-85-7495	A12345678	B12345678	1482928	Bradly	McGordon	04/12/1978	2039871092 bradly.mcgordon@testcompany.com
5	781-70-8498	A12345678	B12345678	5551212	Shiva	Agarwal	11/05/2003 (425)289-5498	sargarwal@national.onmicrosoft.com
6	749 25 3016	A12345678	B12345678	AB-87654321	Flash	Gordon	06/04/1969 203-596-2109	flash.gordon@outlook.com
7	789-21-8631	12345678	12345678	12358	Andrew	O'Donnal	05/09/1994 91-421-94-9012	andrew.odonnal@contoso.onmicrosoft.com
8	182-73-1694	12345678	12345678	JZ00127498	Angelica	Jackson	02/12/1968	44823291203 angelica.jackson@msn.com
9	303-81-0470	12345678	12345678	CO-00003417	Ruchi	Gupta	12/31/2001	ruchig@office.com
10	758 13 4820	12345678	12345678	8726182745	William	Richardson	3/28/1999	2534375906 william.richardson@contoso2.onmicrosoft.com
11	111-69-8921	12345678	12345678	92298375	Joseph	Smith	10/03/1987 1(920)590-2397	joseph.smith@outlook.com

In this step you'll need to create 2 files in a supported file format (csv, tsv or pipe-formatted). The first file contains sample data, as the one that is shown in the picture above. The service uses this file to detect columns so you can specify which will be your primary fields (unique values) and secondary fields (values that are present in close proximity with your primary fields).

Next, you'll need a file that uses the same columns as the ones you used in your sample data. However this time, you need to fill the file with actual data from your environment. As this is data that can be highly sensitive, only hashes of the values in the file will be uploaded to the

Purview service. There's even a procedure to hash the values on a system that is not connected to the internet so that you can be sure the actual values are not uploaded. The file with actual values can contain up to 100 million rows of sensitive data!

Create the EDM classifier

Now it's time to dive in the Purview portal and put the data we've created in the first step to work. Navigate to 'Data Classification' and 'Classifiers'. On the right hand side, let's go for 'EDM classifiers'. Click 'Create EDM classifier'.

You'll be greeted with a familiar screen that outlines the process. Click 'Create EDM classifier' and give your EDM Classifier a name and description.

Choose a method for defining your schema

The schema defines the sensitive info you want to detect (such as patient records) and consists of one or more fields that identify the specific types of info (such as patient ID and name) related to the schema.

Upload a file containing sample data RECOMMENDED

This method is faster and helps you identify errors when mapping sensitive info types to sensitive data. We'll extract field names from your sample file to create the schema and recommend sensitive info types to map the sample field data to. [View all sample files](#) | [Learn how to format the sample file](#)

Manually define your data structure

You'll have to name the fields and choose which sensitive info types best match them. [Learn how to manually define your schema](#)

Next, we can choose 2 options. The first (recommended) one is to use our sample file to identify primary and secondary field names and create the schema for the EDM classifier. We also have to map the sample field data to a sensitive information type, which we'll get a recommendation for. The second option is to manually define your data structure. I choose for the first (recommended) approach. Click next and upload your sample file, which is then processed by the service.

Verify your sample data is correct

Review the sample data uploaded from your file to make sure it's accurate.

↑ Reupload the file 9 items

Column name	Sample data
SSN	270 34 7884, 030 72 7381, 757-85-7495, 781-70-8498, 749 25 3016 +5 more
IndividualID	A12345678, A12345678, A12345678, A12345678, A12345678, 12345678 +4 more
AccountID	B12345678, B12345678, B12345678, B12345678, B12345678, 12345678 +4 more
MRN	AB12345678, 80002910, 1482928, 5551212, AB-87654321, 12358, JZ00127498 +3 more
FirstName	Richard, Sarah, Bradly, Shiva, Flash, Andrew, Angelica, Ruchi, William, Joseph
LastName	Jackson, Jefferson, McGordon, Agarwal, Gordon, O'Donnal, Jackson, Gupta +2 more
DOB	1/1/1990, 2/2/1985, 4/12/1978, 11/5/2003, 6/4/1969, 5/9/1994, 2/12/1968 +3 more
Phone	12065254152, 1-206-234-7492, 2039871092, (425)289-5498, 203-596-2109 +5 more
Email	Richard@contoso.onmicrosoft.com, sarah.jefferson@contoso.com +8 more

Once the processing is done, you'll be presented with the columns that are identified and sample data inside those columns. I would like to stress here that you'll walk through this process with sample (not actual) data because this data is not hashed offline before uploaded to the service! Click next.

Select primary elements

Choose which fields contain the main sensitive information you're trying to detect by identifying them as "primary elements." Fields to sensitive information types, such as social security or credit card numbers, that exactly match your sample data. To further the accuracy of your matches, identify fields without a SIT as multi-token or single-token. Get tips for completing this step

Reset to original

Column name	Primary element	Match mode	Match validation
SSN	<input checked="" type="checkbox"/>	U.S. Social Security Number...	<input checked="" type="radio"/> Full match
IndividualID	<input type="checkbox"/>	Malaysia Passport Number	<input type="radio"/> Partial match
AccountID	<input type="checkbox"/>	Single-token	<input type="radio"/>
MRN	<input type="checkbox"/>	Single-token	<input type="radio"/>
FirstName	<input type="checkbox"/>	Single-token	<input type="radio"/>
LastName	<input type="checkbox"/>	Single-token	<input type="radio"/>
DOB	<input type="checkbox"/>	Single-token	<input type="radio"/>
Phone	<input type="checkbox"/>	Single-token	<input type="radio"/>
Email	<input type="checkbox"/>	Single-token	<input type="radio"/>

Summary of how "U.S. Social Security Number (SSN)" matches data in the "SSN" column

Review how well the sample data from the "SSN" column matches the sensitive info type "U.S. Social Security Number (SSN)".

Total	Match	Not a match
10	10 (100%)	0 (0%)

Sample data Matching results

270 34 7884	<input type="radio"/> Match
030 72 7381	<input type="radio"/> Match
757-85-7495	<input type="radio"/> Match
781-70-8498	<input type="radio"/> Match
749 25 3016	<input type="radio"/> Match
789-21-8631	<input type="radio"/> Match
182-73-1694	<input type="radio"/> Match
303-81-0470	<input type="radio"/> Match
758 13 4820	<input type="radio"/> Match
111-69-8921	<input type="radio"/> Match

Now for the cool part. The service takes a look at the data in the columns and recommends a SIT that would match the data in the column. In my case, the 'Social Security Number' (SSN)

column matches the 'US Social Security Number' SIT. Because I know this is always a unique value in the data I provide, I select the SSN to be the primary element and click next.

In the next screen you can select whether data in the columns is case-sensitive or whether delimiters and punctuation should be ignored. Click Next.

Configure detection rules for primary elements

Each primary element can contain up to 3 rules, each with a unique confidence level that helps determine how likely the sensitive info type detected in content exactly matches the primary element. Confidence typically increases when more supporting elements are detected within close proximity of the primary element. We added supporting elements and character proximity for high and medium confidence rules below, but you can edit them and also add a low confidence rule if needed. [Learn more about detection rules](#)

Detect supporting elements within this many characters of the primary element

Reset to original 1 item

Primary element	Confidence level
SSN	High, Medium
High confidence level	
Detect "SSN" and ANY2 supporting elements below within 300 characters	
IndividualID AccountID MRN FirstName LastName DOB Phone Email	
Medium confidence level	
Detect "SSN" and ANY supporting elements below within 300 characters	
IndividualID AccountID MRN FirstName LastName DOB Phone Email	

Remember that we talked about that secondary elements should be in close proximity to the primary element? In our case the primary element is the Social Security Number (SSN). For the service to consider the match as a 'high confidence level', the primary value should be accompanied by 2 supporting (or secondary) elements within 300 characters.

For the service to consider the match as a 'medium confidence level', the primary value should be accompanied by 1 supporting (or secondary) element within 300 characters. If you want, you can change this configuration in this screen.

Select next, review your settings and click 'Submit'.

✓ You successfully created an EDM classifier

Almost time to put your new EDM classifier to work. Last step is to upload the file containing your org's sensitive info. This is used to populate the structure you set up to create this classifier.

Next step

Use the EDM Upload Agent tool to hash and upload your data

This can be done using one computer or you can separate hashing from uploading for greater security.

[Learn more](#)

① To complete this step, you'll need to know the name of the schema that was just created. Copy the name below or find it by selecting this EDM classifier from the list and viewing the 'EDM schema name' in the details panel.

Schema name

dominiquehermanscomushealthcaredmclassifierSchema  [Copy](#)

When done, the service reminds you that you now have to upload your sensitive data and tells you the schema name it generated. The schema holds the structure for the data you upload, consisting of columns, primary and secondary elements. Copy the schema name and finish the wizard.

So, to recap what we've done so far:

Created 2 tables (CSV format); 1 with sample data to set up our database schema, for which we configured the primary and secondary columns and matching SIT for the primary column. A second table containing the actual data in our environment. In a real world scenario, this data would have been extracted from a database for example and put in the CSV.

Securely upload the values in the table with sensitive data

Next we are going to securely upload the values in the table with sensitive data to Purview. We use a tool called the 'EDM Upload Agent Tool' to hash and upload the data. Like I told earlier, you could separate this process and hash the data on a non-internet connected system but for simplicity's sake we use 1 system in this demo.

First, you make sure you meet all the [prerequisites](#). Second, create a security group in your Microsoft 365 environment and name it 'EDM_DataUploaders'. Add members of your organization that will be maintaining the database with sensitive information.

Next, download and install the EDM Upload Agent Tool from [Microsoft Learn](#) (<https://learn.microsoft.com/en-us/purview/sit-get-started-exact-data-match-hash-upload?tabs=single-computer#hash-and-upload-your-data>).

```
c:\Program Files\Microsoft\EdmUploadAgent>EdmUploadAgent.exe /Authorize
Command completed successfully.

c:\Program Files\Microsoft\EdmUploadAgent>EdmUploadAgent.exe /SaveSchema /DataStoreName dominiquehermanscomushealthcaredmclassifierSchema /OutputDir C:\Users\DominiqueHermans\EDM\
Command completed successfully.
```

Start a cmd window as administrator and logon to your tenant using the 'EdmUploadAgent.exe /Authorize' command. Next, download the XML file that contains the schema we just created using the GUI by using the following command:

```
EdmUploadAgent.exe /SaveSchema /DataStoreName
dominiquehermanscomushealthcaredmclassifierSchema /OutputDir
C:\Users\Username\EDM\
```

Next, we are running a check on our CSV file to see if it contains any characters that might be a problem. To do this, run the following command against your CSV File:

```
EdmUploadAgent.exe /ValidateData /DataFile US_Healthcare_ActualData.csv /Schema
dominiquehermanscomushealthcareedmclassifierSchema.xml
```

The command should return 'passed the schema validation.'

Now let's hash our sensitive data and upload it to the service by using the following command:

```
EdmUploadAgent.exe /UploadData /DataStoreName
dominiquehermanscomushealthcareedmclassifierSchema /DataFile
C:\Users\UserName\EDM\US_Healthcare_ActualData.csv /HashLocation
C:\Users\UserName\EDM\Hash /Schema
C:\Users\UserName\EDM\dominiquehermanscomushealthcareedmclassifierSchema.XML
/AllowedBadLinesPercentage 0
```

```
c:\Program Files\Microsoft\EdmUploadAgent>EdmUploadAgent.exe /GetDataStore
Printing list of Datastores.
Id, Name, DataLastUpdatedTime
dominiquehermanscomushealthcareedmclassifierschema, dominiquehermanscomushealthcareedmclassifierschema, 18/09/2024 09:31
:09
Command completed successfully.
```

Now when you run the 'EdmUploadAgent.exe /GetDataStore' command, you will receive a list of datastores and when the datastore was last updated.

Let's Test!

The screenshot shows the Microsoft EDM Classifier interface. On the left, there is a list of classifiers with a table header: Name, Created by, and Status. One classifier is listed: 'DominiqueHermans.com - U.S. Healthcare EHLO%20Corporation%20Inc.' with a status of 'Index complete'. On the right, a detailed view of the selected classifier is shown. The title is 'DominiqueHermans.com - U.S. Healthcare EDM Classifier'. Below the title are buttons for Edit, Delete, and a highlighted 'Test' button. The classifier details include: EDM classifier name (DominiqueHermans.com - U.S. Healthcare EDM Classifier), EDM classifier description (U.S. Healthcare EDM Classifier), EDM schema name (dominiquehermanscomushealthcareedmclassifierschema), Sensitive info types for primary elements (SSN - U.S. Social Security Number (SSN)), Schema file column settings (Data in all columns is case insensitive), and Detection rules (SSN - 2 confidence levels (High, Medium)).

Now with our EDM Classifier in place and data uploaded, we can test our EDM classifier. To do this, navigate to 'Data Classification' and 'Classifiers'. On the right hand side, let's go for 'EDM classifiers'. Select your newly created EDM classifier and select 'Test' at the right hand side.

Next, upload a document that you want to test against your EDM classifier. For this demo, I've cooked up a sample letter to one of my fictional patients.

Match results

We have detected the following in [Patient Letter 3.docx](#)

Base sensitive information type results

Name	Low	Medium	High
U.S. Social Security Number (SSN)	1	0	0

[View all results](#)

EDM sensitive information type results

Name	Low	Medium	High
DominiqueHermans.com - U.S. Healthc...	1	1	1

[View all results](#)

As you can see, the elements in the document are discovered by the SIT that's attached to the primary element, and the EDM SIT itself!

Trainable Classifiers

Now let's talk about the most extensive classification feature in Purview. Trainable classifiers. In short, they are custom built classifiers that can find (mostly unique) data that cannot be found by using other methods, like the ones we spoke about in this chapter.

They are configured in the following manner:

Navigate to Purview, data classification, classifiers, trainable classifiers and select 'start scanning process'. The process will identify the content that you have in your organization. This process can take up to 14 days to complete!

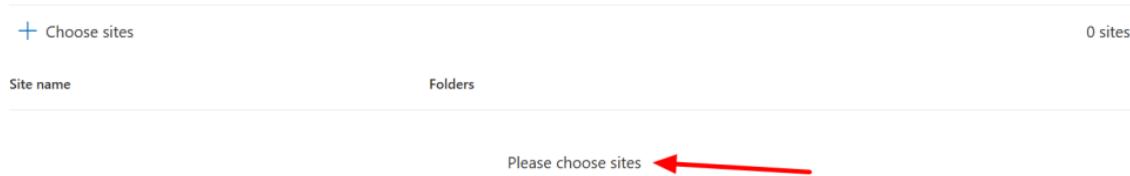
Source of the positive sample content

To learn how to classify content, the model needs samples that match the kind of document you want to classify (positive samples) that is will compare to non-matching samples (negative samples). [Learn how the model works](#) and [how to create responsible AI](#)

Each site should include between 50 and 500 files. If you include a site that contains more, we'll only process 500 of the most recently created files.

Content must be written in English. Supported file types includes: Office docs such as Word, PowerPoint, Excel, PDFs, text files, and other files types. [Learn what file types are supported](#)

Tip for choosing positive sample content: We recommend choosing specific folders for each site so we can seed the classifier with files that are most related to the type of content you want to classify.



1. Gather seed data (at least 50 samples/documents) and upload it to the trainable classifier using a SharePoint site. In this case the idiom ‘the more the merrier’ applies. The more samples you feed the trainable classifier, the more reliable it will be. The samples will then be processed by the service using Artificial Intelligence (AI) and Machine Learning (ML) techniques to identify similarities. This process can take up to 24 hours to complete. Note that you need at least 50 samples that are a strong match and at least 50 samples that are not a match.
2. Now, when the first process is done, your trainable classifier has to be trained and tested. To do this, use a separate SharePoint site in which you feed it with (at least 30, again, the more the merrier) test samples/documents. This should be positive, negative and somewhat vague items to train your classifier. After the service goes through its paces again, you will get the opportunity to review the testdata from this step in the process. Once the accuracy score for your trainable classifier stabilizes, you have the option to publish the trainable classifier.
3. You publish the Trainable classifier so it can be used in your policies.

Note that also in the case of Trainable Classifiers, Microsoft provides you with a list of pre-trained trainable classifiers that you can utilize!

Where can I apply all this goodness?

Now let's talk about where we can use all of the classifiers we talked about in this chapter.

- **Sensitive Information Types (SIT's):** Data Loss Prevention, Sensitivity Labels, Retention Labels, Insider Risk Management, Communication Compliance, Auto-labeling policies, Microsoft Priva.
- **Fingerprint Sensitive Information Types:** Data Loss Prevention
- **Exact Data Match Classifiers:** DLP Policies, Auto-labeling policies, Microsoft Defender for Cloud Apps

- **Trainable Classifiers:** Auto-labeling Office files with sensitivity labels, auto-applying a retention label policy based on a condition, communication compliance, Data Loss Prevention.

Phew, this chapter got a little longer than I initially thought. But hey, if you made it to the end, you now know all about sensitive info types and classifiers!

Data Lifecycle Management (DLM)

One of the core features of Microsoft Purview is Data Lifecycle Management (DLM), formerly known as Microsoft Information Governance. DLM is all about providing you with the tools you need to keep information that you need, and delete the information that you don't. This process is of great importance for compliance with regulations, risk management and liability management.

Data Lifecycle Management in Microsoft Purview has a great overlap with Records Management, although they also have their differences. This chapter will explain how to configure the basics of Data Lifecycle Management, and will show you the end-user experience.

Terminology

Let's start with some terminology:

- **Retention Policies** are applied to containers by DLM administrators. Examples of a container can be a SharePoint site, Microsoft 365 group, a team in Microsoft Teams or a mailbox in Exchange Online. All information within these containers will inherit the retention policy that is set at the container level. Note that retention policies are not visible to your end-users. The recommended approach here is that Retention Policies should be your foundational solution, where labels should be used to configure documents that are an exception to this rule.
- **Labels** are applied at the item level, such as Word, Excel or PowerPoint files within a container. They can be auto-applied (for example by using pattern matching) or manually applied by your end-users.
- **Label Policies** determine where labels are published. For example a SharePoint Site, Exchange Mailbox or OneDrive accounts.

Configuring Retention Policies

With the terminology out of the way, let's dive in and configure a retention policy.

Data lifecycle management

Navigate to the Microsoft Purview portal at <https://compliance.microsoft.com> and navigate to “Data Lifecycle Management” and “Microsoft 365”. Click “New Retention Policy”.

Name your retention policy

Name *

RP - Retain 7 days then Delete

Description

This policy will retain data to which it is applied for 7 days. After this period, the data will be deleted.

Give your retention policy a name that makes sense and is an explanation to what the retention policy will achieve. Also provide a description with the same thing in mind. In the next screen, admin units can be provided if you configured your environment to be divided in those units. Each unit can be managed by a specific set of administrators. However, this is beyond the scope of this chapter so we'll go with the standard "Full Directory" admin unit.

Choose the type of retention policy to create

Locations can be specified dynamically with an adaptive scope using attributes or properties, or if you know the specific target locations, you can select them individually from a list. An advantage of using an adaptive scope to determine target locations is that it will automatically update where it's applied based on the attributes or properties you define.

Adaptive

After selecting adaptive policy scopes, which consist of attributes or properties (e.g. 'Department' or 'Site URL') that define the users, groups, or sites in your org, you'll choose supported locations containing the content you want to retain. The policy will automatically update to match the criteria defined in the scopes.

Static

You'll choose locations containing the content you want to retain. If locations change after this policy is created (for example if a SharePoint site is added or removed), you'll need to manually update the policy.

Next, we can choose the type of retention policy that should be created. An adaptive policy scope can be used to select locations containing the content you want to retain or delete in a dynamic way. When new locations are added that match the adaptive policy scope, they are automatically added. This is, in fact, the recommended way that Microsoft wants you to assign your retention policy. However, I would like to make clear how retention policies work before throwing in adaptive scopes in the equation. So for now, I'll select "Static" which lets you manually select the location to where the retention policy should be applied. As the descriptive text explains, when new locations are added or existing ones change, you'll have to manually update the policy to include these locations.

Status	Location	Applicable Content
Off	Exchange mailboxes	Items in user, shared, and resource mailboxes, calendar items with an end date, no end date. Doesn't apply to items in SharePoint communication sites. More details
Off	SharePoint classic and communication sites	Files in classic sites or communication sites that aren't connected to a Microsoft 365 Group. Doesn't apply to all document libraries (including default Assets). More details
On	Microsoft 365 Group mailboxes & sites	Items in the Microsoft 365 Group mailbox and corresponding group-connected SharePoint communication sites or SharePoint connected to Microsoft 365 Groups

In the next screen we can select the locations to where this policy will apply. Here I've chosen "Microsoft 365 Group Mailboxes & Sites" and further narrowed it down to just the "Mark 8 Project Team" Microsoft 365 group mailbox & site.

Decide if you want to retain content, delete it, or both

Retain items for a specific period
Items will be retained for the period you choose.

Retain items for a specific period
of years months days
Custom

Start the retention period based on

At the end of the retention period
 Delete items automatically
 Do nothing

Retain items forever
Items will be retained forever, even if users delete them.

Only delete items when they reach a certain age
Items won't be retained, but when they reach the age you choose, we'll delete them from where they're stored.

The next screen is where the action is at. Here we can choose what needs to happen with the content we specified in the previous screen. Items can be retained for a specific period after which they can be automatically deleted, they can be retained for ever or can be deleted when they reach a certain age. In this example, I choose to retain them for 7 days after which they will be deleted automatically. Note that this is just for demonstration purposes, in a live environment a period of a couple of years would be more realistic.

Review and finish

It will take up to a week to apply this policy to the locations you selected.

Policy name

RP - Retain 7 days then Delete

[Edit](#)

Description

RP - Retain 7 days then Delete

[Edit](#)

Locations to apply the policy

Microsoft 365 Group mailboxes & sites (1 Group)

[Edit](#)

Retention settings

Retain items for 7 days based on when they were last modified

Delete items at end of retention period

[Edit](#)

⚠ Items that are currently older than 7 days will be permanently deleted after you turn on this policy.

Take a look at the review screen and note that “items that are currently older than 7 days will be permanently deleted after you turn on this policy”. Also take a look at the line on the top of the summary mentioning “it will take up to a week to apply this policy to the locations you selected”. You will be notified of this fact once more when you end the wizard to create the policy.

Name	Modified	Modified By
MARK8-ElevatorPitch.pptx	12 hours ago	Debra Berger
XT1050 Usability test 2.3.docx	12 hours ago	Debra Berger
Usability Testing Priorities.docx	12 hours ago	Debra Berger
ELEVATOR-PITCH.1.jpg	February 12	MOD Administrator
ELEVATOR-PITCH.pdf	February 12	MOD Administrator
ELEVATOR-PITCH.jpg	February 12	MOD Administrator
XT1050 Marketing Collateral Timelines_V2.docx	February 12	MOD Administrator
marketing-initiatives-FY17.xlsx	February 12	MOD Administrator
bearing-44DDF-stress-test.xlsx	February 12	MOD Administrator

To close out this chapter, take a look at the “design” channel of the “Mark 8 Project Plan” team to which the retention policy was applied. Files in the Teams channel / SharePoint teamsite are all older then 7 days, with the exception of the files that are highlighted in the screenshot above.

Remember, retention policies are the foundation of your DLM configuration!

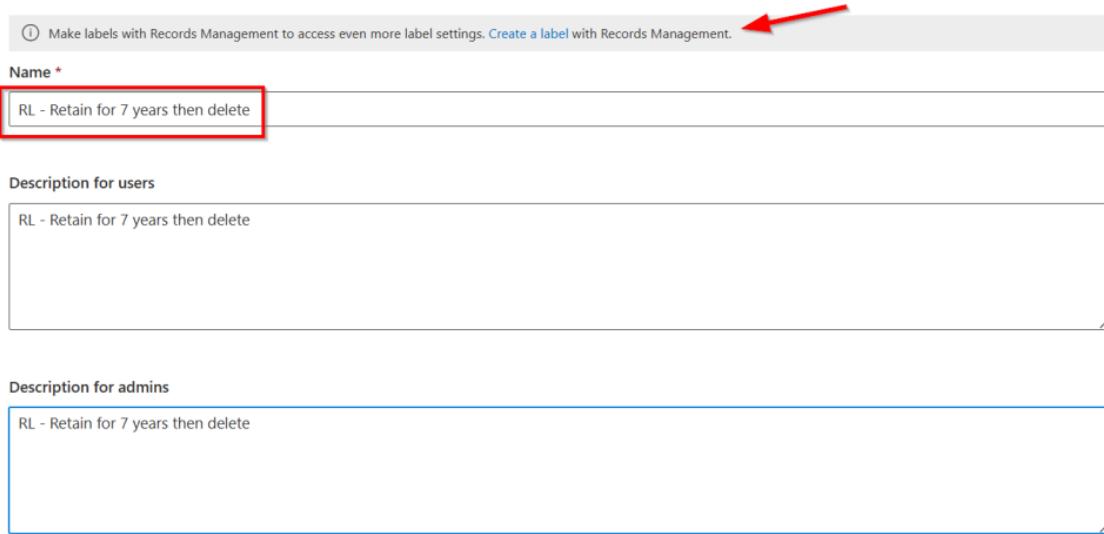
Configuring (Retention) Labels

Now that the foundation is in place with retention policies, let’s focus on the progress of configuring (retention) labels so that they can be used as an exception to the configured retention policy.

Navigate to “Data Lifecycle Management” and “Microsoft 365”. Click the “Labels” tab and click “Create a label”.

Name your retention label

This is the name of the label your users will see in the apps where it's published (like Outlook, SharePoint, and OneDrive). So be sure to come up with a name that helps them understand what it's used for.



① Make labels with Records Management to access even more label settings. [Create a label with Records Management](#).

Name *

RL - Retain for 7 years then delete

Description for users

RL - Retain for 7 years then delete

Description for admins

RL - Retain for 7 years then delete

Again, create a name for your label that makes sense and can be used later to distinguish the label from other labels. Also, create a description for users and admins that make clear where the label should be used for.

Define label settings

We'll apply the settings you choose to labeled items

- Retain items forever or for a specific period**
Items won't be retained but when they reach the age you specify, they'll be deleted from where they are stored.
- Enforce actions after a specific period**
Labeled items won't be retained. You can decide whether they should be deleted, or relabeled when the period you specify in the next step ends.
- Just label items**
Choose this setting if you only want to classify labeled items. The items won't be retained and your users won't be restricted from editing, moving, or deleting them.

In the next screen, the label settings can be configured. You can configure to retain items for a certain period, enforce actions after a specific period or just label items with no action applied. In this demonstration I'll retain items for a specific period.

Define the retention period

Specify how long the retention period should be.

Retain items for

7 years

Start the retention period based on

When items were last modified

[+ Create new event type](#)

The retention period has to be configured next. Here I choose to retain items for 7 years. The retention period is based on the last modification date of the item.

Choose what happens after the retention period

These settings determine what happens to items when the retention period ends.

- Delete items automatically**
We'll permanently remove labeled items from wherever they're stored.
- Start a disposition review**
Let the disposition reviewers you assign in the next step decide if items can be safely deleted or whether other actions (such as changing the retention period) should be taken. [Learn more](#)
- Change the label**
You can extend the period by choosing an existing label to replace this one with. [Learn more about relabeling items](#)
- Run a Power Automate flow**
Customize what happens to labeled items with a Power Automate flow. You can run a flow to meet a specific business need, such as moving labeled items to a certain location or sending email notifications. [Learn more about running a Power Automate flow](#)
- Deactivate retention settings**
Labeled items won't be retained or deleted when their retention settings are deactivated. You'll have to manually remove any items that you want deleted.

What has to happen after the retention period should be configured next. To keep this progress simple I've chosen to "delete items automatically". Review the summary next and click "create label". A new wizard appears telling you that your label has been created but that one of 2 actions have to happen next:

1. Publish the label to Microsoft 365 locations. This way, users can use the label to manually apply it to their content.
2. Auto-apply this label to a specific type of content. Using this path, you can create conditions that are used to match your content to which the label is then automatically applied.

For now, select "do nothing" and finish the wizard. This way, we can see each step of the process.

Publishing (Retention) Labels

As I would like to keep auto-applying of labels for another time, let's publish the label to a Microsoft 365 location to see the (end-user) impact there. Let's go!

Back in the Purview portal, make sure the “Label policies” tab is selected and choose “Publish Labels”. In the first pane of the wizard click “Choose labels to publish” and select the label that was just created. In my case that’s “RL – Retain for 7 years then delete”. In the admin units screen, click “next”.

Choose the type of retention policy to create

Locations can be specified dynamically with an adaptive scope using attributes or properties, or if you know the specific target locations, you can select them individually from a list. An advantage of using an adaptive scope to determine target locations is that it will automatically update where it's applied based on the attributes or properties you define.

Adaptive

After selecting adaptive policy scopes, which consist of attributes or properties (e.g. 'Department' or 'Site URL') that define the users, groups, or sites in your org, you'll choose supported locations containing the content you want to retain. The policy will automatically update to match the criteria defined in the scopes.

Static

You'll choose locations containing the content you want to retain. If locations change after this policy is created (for example if a SharePoint site is added or removed), you'll need to manually update the policy.

The next screen asks us again whether it should use a adaptive or static scope, just as with the creation of our retention policy (even the title is the same 😊). Choose static for the sake of simplicity and in the next screen -which is also the same as with the creation of our retention policy- choose “Microsoft 365 group mailboxes & sites” and select “Mark 8 project Team”. In the last screen, give your label policy a name and description that makes sense. In my case it's “LP – 7 Years then delete – Mark 8 Project Team SharePoint site”.

Finish

⚠ Most labels will become available to your users within a week. Labels will appear in Outlook and Outlook on the web only for mailboxes that have at least 10 MB of data.

Choose labels to publish

1 label(s) will be published (made available) so your users can classify their content
RL - Retain for 7 years then delete 7 years keep + delete

[Edit](#)

Applies to content in these locations

Microsoft 365 Group mailboxes & sites (1 Group)

[Edit](#)

Name

LP - 7 Years then Delete - Mark 8 Project Team Sharepoint Site

[Edit](#)

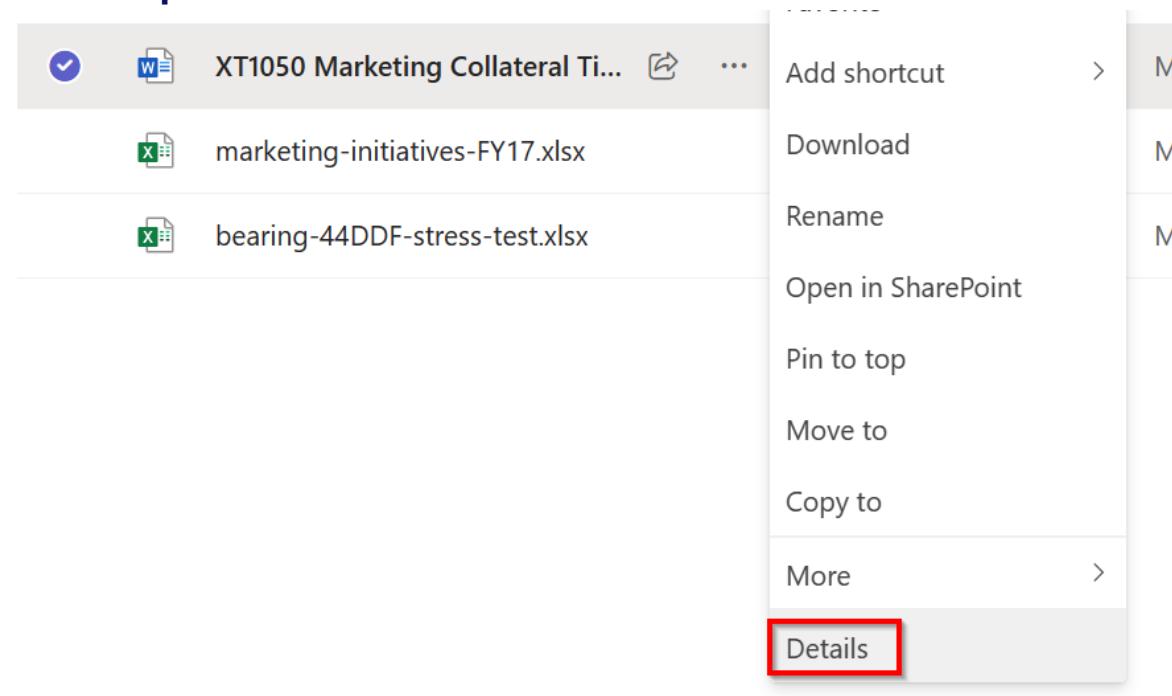
Description

LP - 7 Years then Delete - Mark 8 Project Team Sharepoint Site

[Edit](#)

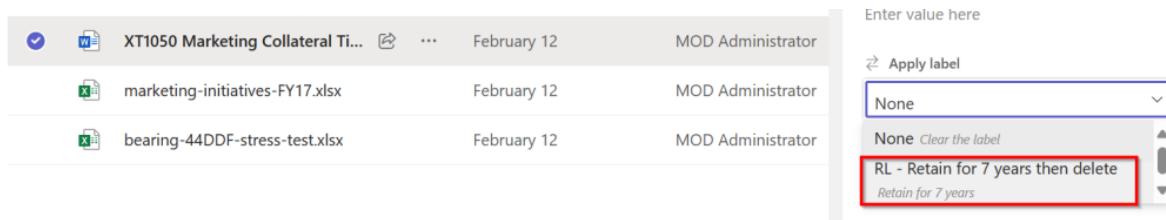
In the summary, be sure to notice that it can take up to a week for your labels to become available in the chosen location. If this label was published to a mailbox, the mailbox needs to have 10 MB of data before the label will appear. Now your retention label is published to the Mark 8 project team SharePoint site / Teams channel.

The user experience



So, what about the user experience? As stated in the beginning of the chapter, labels that are published to a location using a label policy can be manually applied by users or automatically applied by using pattern matching for example. In this example, we've published the labels so that a user can apply them manually to items. These items would be an exception to the retention policy we applied at the container level.

In the above screenshot, I've logged in to Microsoft Teams with user "Debra Berger". Debra has permissions to apply retention labels. In the "Mark 8 Project Team" team, she selects a file and clicks "details".



Here we can apply our published retention label "RL – Retain for 7 years then delete".

Name	Modified	Modified By	Retention label
MARK8-ElevatorPitch.pptx	Monday at 11:57 AM	Debra Berger	
Usability Testing Priorities.docx	Monday at 11:56 AM	Debra Berger	
XT1050 Marketing Collateral Timelines_V2.docx	February 12	MOD Administrator	RL - Retain for 7 years then delete
XT1050 Usability test 2.3.docx	Monday at 11:57 AM	Debra Berger	

Fast forward a few days for all policies to take effect and notice how our files view in the “Mark 8 Project Team” – “Design” channel has changed. I’ve added a column to easily see which retention label is applied to which file. Remember the files that were shown at the start of this chapter? A lot of them are gone because:

1. We applied a retention policy at the container level (Team / SharePoint teamsite) as an admin that deleted items that were older than 7 days.
2. We manually applied a retention label that we published to the Team / SharePoint teamsite and added to the “XT101 Marketing Collateral Timelines_V2.docx” file.
3. So only the files that are not older than 7 days and the file that was manually tagged remain in place.

How to see where labels are applied?

If you want to get a complete view of where your retention labels are applied, you can do so as follows:

Navigate to the Purview Portal and select “Content Search”, “New Search”. Give the search an appropriate name and description so you can easily find and re-run the search later.

New search

Locations

Specific locations

Status	Location	Included	Excluded
Off	Exchange mailboxes Microsoft 365 Groups Teams Yammer user messages		
On	SharePoint sites OneDrive sites Microsoft 365 Group Sites Team Sites Yammer Networks	All Choose sites	None
Off	Exchange public folders		

Add App Content for On-Premises Users. [Learn more](#)

Specify locations. Here, I’ll use “SharePoint Sites”.

New search

- Name and description
- Locations
- Conditions**
- Review your search

Define your search conditions

Query language-country/region: None [? F](#)

Query builder KQL editor

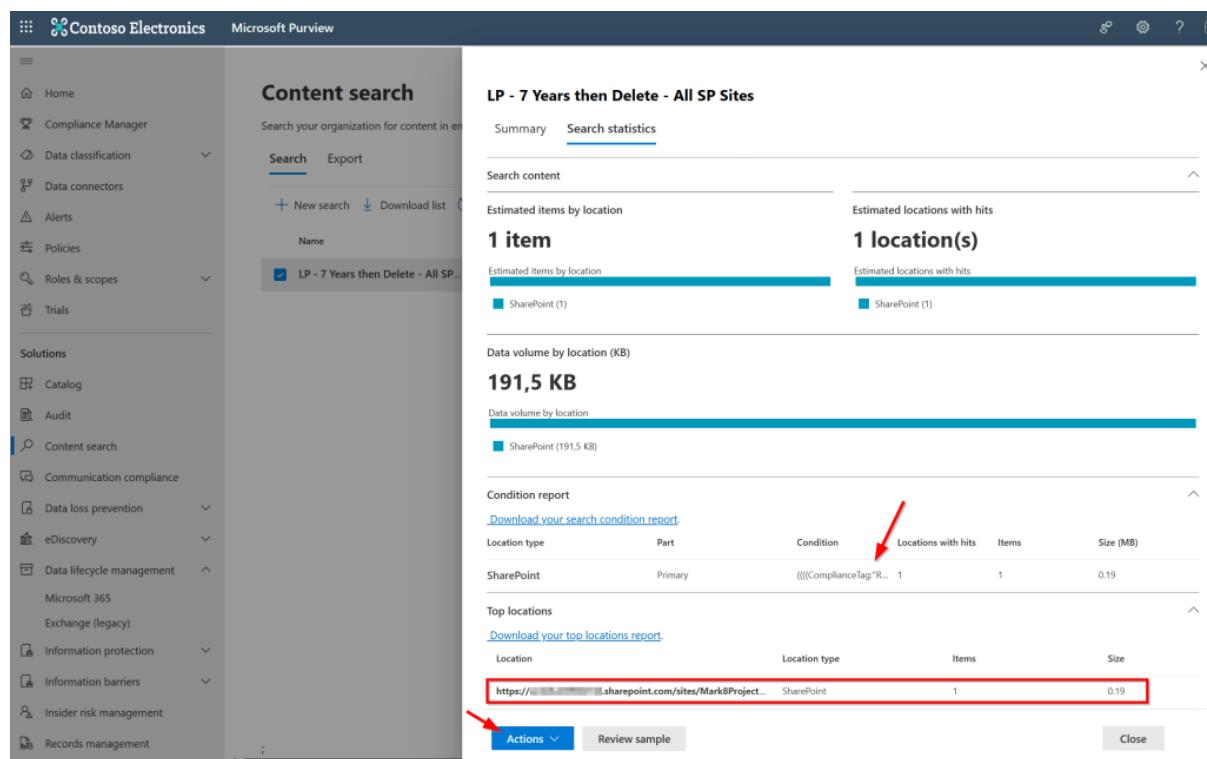
Retention label

Equals any of

RL - Retain for 7 years then delete

[+ Add condition](#)

Define your search condition, in this example I'll use “Retention Label”, “Equals any of”, “RL – Retain for 7 years then delete”. Review the summary screen and finish the wizard.



Content search

Search your organization for content in enterprise search

Search Export

+ New search [Download list](#)

Name

LP - 7 Years then Delete - All SP Sites

LP - 7 Years then Delete - All SP Sites

Summary [Search statistics](#)

Search content

Estimated items by location

1 item

Estimated items by location

SharePoint (1)

Estimated locations with hits

1 location(s)

Estimated locations with hits

SharePoint (1)

Data volume by location (KB)

191,5 KB

Data volume by location

SharePoint (191,5 KB)

Condition report

[Download your search condition report](#)

Location type	Part	Condition	Locations with hits	Items	Size (MB)
SharePoint	Primary	(((ComplianceTag:"R... 1		1	0.19

Top locations

[Download your top locations report](#)

Location	Location type	Items	Size
https://.../sharepoint.com/sites/MarkBProject... SharePoint	SharePoint	1	0.19

[Actions](#) [Review sample](#) [Close](#)

When looking in the details of the search result, you can see exactly where your labels are applied!

When only looking for documents that are labeled from within SharePoint Online, you can use the following trick.

Navigate to Sharepoint Online and in the search bar type:

ComplianceTag:"Your TagName here"

Documents labeled with your retention label will automatically show up! Note that this only shows documents that are within SharePoint online and are accessible by the user that performs the search!

How to see where Retention Policies are applied?

So as we saw above, retention labels can be easily seen by the user -whether it's in Teams or in SharePoint- by looking at the details of a file or adding the "retention label" column to your view. Administrators can also utilize Content Search or users may use SharePoint Search.

So what about retention policies? As stated earlier, they are invisible to the end user as they are applied by your admin or lifecycle management department at the container level. So we can use "policy lookup" as admin or someone with the appropriate role group to check where retention policies are applied.

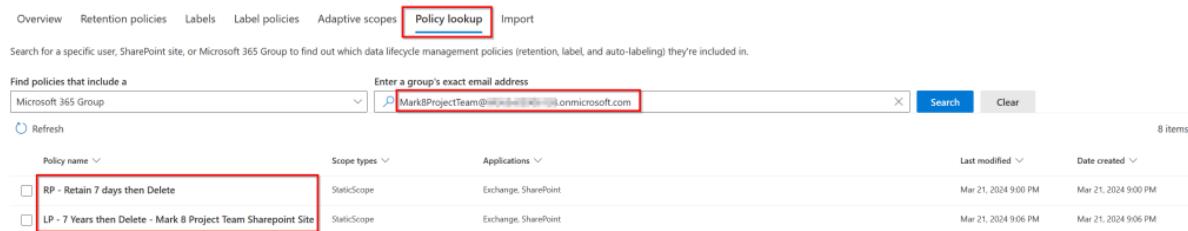
Policy Lookup – Search within a Team / Microsoft 365 group

In the DLM toolbox, we can use "Policy Lookup" to see where certain retention policies are applied. First, navigate to your Microsoft 365 admin center and navigate to "Teams & Groups", "Active Teams and Groups".

Name	Email	Sync status	Teams status
Mark 8 Project Team	Mark8ProjectTeam@contoso.onmicrosoft.com		

Look for the Team that you want to look up retention policies for and copy its email address.

Data lifecycle management



Overview Retention policies Labels Label policies Adaptive scopes **Policy lookup** Import

Search for a specific user, SharePoint site, or Microsoft 365 Group to find out which data lifecycle management policies (retention, label, and auto-labeling) they're included in.

Find policies that include a

Enter a group's exact email address

Microsoft 365 Group

Policy name ▾ Scope types ▾ Applications ▾ Last modified ▾ Date created ▾

Policy name	Scope types	Applications	Last modified	Date created
<input type="checkbox"/> RP - Retain 7 days then Delete	StaticScope	Exchange, SharePoint	Mar 21, 2024 9:00 PM	Mar 21, 2024 9:00 PM
<input type="checkbox"/> LP - 7 Years then Delete - Mark B Project Team Sharepoint Site	StaticScope	Exchange, SharePoint	Mar 21, 2024 9:06 PM	Mar 21, 2024 9:06 PM

8 items

Now navigate to the Purview portal, Data Lifecycle Management, Microsoft 365, Policy lookup and select “Microsoft 365 Group”. Next, paste your copied group email address and click “search”. There you have it, it shows all your retention policies AND label policies that are applied!

Records Management (RM)

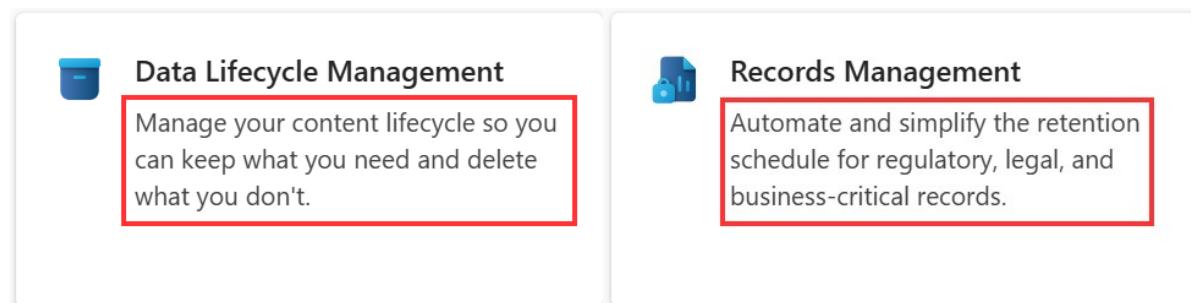
Records Management in Microsoft Purview can be used to:

- Setup a retention schedule for your files or folders (Just as with [Purview Data Lifecycle Management](#))
- Mark items such as Word, Excel or Powerpoint files as records.

When an item such as a file or folder becomes a record, the item or its contents cannot be changed any more. This is often done to comply with legal requirements, such as those that require a certain company to retain their documents for a certain period of time and during that time, the files (that have become records) cannot be altered by anyone. This chapter will explain how to configure the basics of Records Management, and will show you the end-user experience.

Differences between Records Management (RM) and Data Lifecycle Management (DLM)

Let's start by looking at the differences between Data Lifecycle Management and Records Management.



The new Microsoft Purview portal explains the differences pretty nice actually. Both features have a certain amount of overlap as they both use retention labeling to manage retention of your items. However, it's the purpose of configuring retention labeling where both features differ. Data Lifecycle Management is used to manage your content so you keep what you need, and delete what you don't. Items that are not present in your environment can not be misused, right?

Records Management on the other hand is used to automate and simplify the retention schedule for regulatory, legal, and business-critical records. And as said earlier, "records" are your items (or files) now turned into a static object called a record. They still look like a Word, Excel or PowerPoint file for example, but they can't be changed in any way. In other words, marking items as a records makes the item immutable. The following actions are prevented when an item is marked as a record:

- Changing the retention of an item.
- Change the contents or metadata of the item.
- Deletion of an item or removal of the retention label that is applied to the item.
- Moving a file between SharePoint libraries.

File Plans

The Microsoft Purview Records Management solution uses a file plan to manage retention labels. Using the File Plan, you can (bulk) create labels and publish or auto-apply labels to a certain location, such as Exchange mailboxes, SharePoint sites and OneDrive Accounts.

Publishing to a certain SharePoint site for example can be done manually by selecting the site, or by using [Adaptive Scopes](#). When auto-applying labels, you can target your label to be automatically applied to content that matches a certain filter, for example content that contains sensitive info, content that contains specific words or phrases, or has certain properties.

Other properties that your file plan includes are:

- Status of your labels (Active, or Inactive)
- Where your label is based on (for example “when created” or “last modified”)
- Whether your label marks item as a record.
- Whether your records should be unlocked by default.
- Whether content should be relabeled.
- The duration of the retention.
- The disposition type that applies when items are to be deleted (for example “No Action”, “Auto-delete” or “Review required”).
- Metadata of your labels (not required).
- Last modification date of the label and who was responsible for the modification.

Configuring Records Management

Let's configure some labels from within the Records Management Console. Let's say for this demo we want to create a retention label that is published to a certain SharePoint site so users can label folders (and items within that folder) to become a record. After a couple of days (because I don't want to wait a year before showing you the result) the items marked as records are placed on a disposition list for review after which they are deleted. Ready? Let's go!

The screenshot shows the Microsoft Purview Records Management console. At the top, there is a navigation bar with the following items: Overview, File plan (which is highlighted with a red box), Label policies, Adaptive scopes, Policy lookup, Events, and Disposition. Below the navigation bar, there is a descriptive text block: "Take advantage of a more flexible and comprehensive way to manage your business-critical data. Our file plan lets you apply item-level retention settings, import or export templates of your content management plan, define whatever label settings meet your needs, review detailed label analytics, and more. [Learn about using a file plan](#)". At the bottom of the screenshot, there is a toolbar with the following buttons: Create a label (highlighted with a red box), Publish labels, Import, Export, Refresh, a search bar with the text "9 items" and a magnifying glass icon, and a filter icon.

Navigate to the Microsoft Purview portal and click “Records Management” on the left hand side. Within Records Management, navigate to file plan (remember the name?) and click “Create a label”.

Create retention label

- Name
- File plan descriptors
- Label Settings
- Period
- Finish

Name your retention label

This is the name of the label your users will see in the apps where it's published (like Outlook, SharePoint, and OneDrive). So be sure to come up with a name that helps them understand what it's used for.

Name *

Mark Item as Record - Retain for 2 Days

Description for users

Marks items as a record so they become immutable. These records are retained for 2 days after which they are placed on a disposition list for review before they are deleted.

Description for admins

Marks items as a record so they become immutable. These records are retained for 2 days after which they are placed on a disposition list for review before they are deleted.

Give your label a logical name, as it will show up in your users Teams or SharePoint environment alongside the description for users. So this will be equally important.

Create retention label

- Name
- File plan descriptors
- Label Settings
- Period
- Finish

Define file plan descriptors for this label

By default, this label will be included in your file plan. To help organize this label, choose any values related to the default descriptor columns included in your file plan.

Reference ID

Business function/department No data available

Choose

Category No data available

Choose

Sub category No data available

Choose

Authority type No data available

Choose

Provision/citation No data available

Choose

In the next screen, you can configure file plan descriptors (metadata of your label), which I've discussed earlier in this chapter. I'll ignore them for now.

Create retention label

- Name
- File plan descriptors
- Label Settings
- Period
- Finish

Define label settings

We'll apply the settings you choose to labeled items

Retain items forever or for a specific period

Items won't be retained but when they reach the age you specify, they'll be deleted from where they are stored.

Enforce actions after a specific period

Labeled items won't be retained. You can decide whether they should be deleted, or relabeled when the period you specify in the next step ends.

Just label items

Choose this setting if you only want to classify labeled items. The items won't be retained and your users won't be restricted from editing, moving, or deleting them.

In the next screen we can configure what will be done to files or items that are labeled with this label. For this demo, I choose to go with “Retain items forever or for a specific period” because I want to show you the disposition list feature in Records Management.

Create retention label

Define the retention period

Specify how long the retention period should be.

Retain items for
of years months days

Custom

Start the retention period based on
When items were labeled

Create new event type

So as I told earlier, I'll go with a retention period of 2 days and that retention period will start when the label is applied.

Create retention label

Choose what happens during the retention period

These settings control what users can do to retained items.

During the retention period

Retain items even if users delete
Users will be able to edit items and change or remove the label. If they delete items, we'll keep copies in a secure location. [Learn more](#)

Mark items as a record
Users won't be able to edit or delete items, and only admins will be able to change or remove the label. For SharePoint or OneDrive files, actions are blocked or allowed based on whether the item's record status is locked or unlocked. [Learn more about records](#)

Unlock this record by default
 Choose this option if you want to allow users to edit items before locking the record. Users won't be able to move items across OneDrive accounts or SharePoint document libraries while the record is unlocked.

First, we'll have to choose what happens with your items (files) during the retention period. Since we want to check out record management, we'll mark the items as a record.

Create retention label

Choose what happens after the retention period

These settings determine what happens to items when the retention period ends.

Delete items automatically
We'll permanently remove labeled items from wherever they're stored.

Start a disposition review
Let the disposition reviewers you assign in the next step decide if items can be safely deleted or whether other actions (such as changing the retention period) should be taken. [Learn more](#)

Stage 1

Name	Alias
MOD Administrator	admin@M365x05906138.onmicrosoft....

[Edit stages, reviewers, and settings](#)

Change the label
You can extend the period by choosing an existing label to replace this one with. [Learn more about relabeling items](#)

Run a Power Automate flow
Customize what happens to labeled items with a Power Automate flow. You can run a flow to meet a specific business need, such as moving labeled items to a certain location or sending email notifications. [Learn more about running a Power Automate flow](#)

Deactivate retention settings
Labeled items won't be retained or deleted when their retention settings are deactivated. You'll have to manually remove any items that you want deleted.

Second, you get to choose what happens with your items after the retention period. I want to show you what a disposition review does since it's a feature specific to records management. With disposition reviews, certain users or departments in your organization are tasked with reviewing items after the retention period is over, but before the items are deleted. The items will show up in a disposition list that needs to be reviewed by one or multiple users. This is what is set up in this screen. Here you can configure 1 or multiple stages and add certain people or groups to each stage. When you configure multiple stages, a user in each stage should review an item before it is actually deleted or for example is configured with another label that can extend the lifetime of an item.

After the disposition review has been configured, take a look at the summary screen and finish the wizard.

Create retention label

Your retention label is created

Creating the label is just the first step in classifying and governing content. To make this label available to users in your organization, publish it in select locations or auto-apply it to specific content.

Next steps

Publish this label to Microsoft 365 locations
You'll create a label policy to make this label available in locations like Exchange and OneDrive. When published, users can manually apply it to their content or set it as the default label for content containers (such as SharePoint document libraries or email folders).

Auto-apply this label to a specific type of content
You'll create an auto-labeling policy to apply the label to content matching certain conditions, such as content containing specific sensitive info.

Do Nothing
You can publish or auto-apply it to content later.

As with Data Lifecycle Management, labels have to be published to a location before it can be used by your users. It can also be auto-applied as stated earlier, but this is out of scope for this chapter.

I'll select "publish this label to Microsoft 365 locations" so the next wizard to publish the new label is automatically started.

Publish labels so users can apply them to their content.

- Choose labels to publish
- Administrative Units
- Scope
- Name your policy
- Finish

Choose labels to publish

Choose the labels you want to publish to your organization's apps so users can apply them to their content. If you don't see the labels you want, you'll be able to create one from scratch.

Publish these labels (1 label(s))

Name	Retention
Mark Item as Record - Retain for 2 Days	2 days keep + review + delete

The label that's just created is already selected. In the policy scope screen, we are going to use the default of "Full Directory".

Publish labels so users can apply them to their content.

- Choose labels to publish
- Administrative Units
- Scope
- Name your policy
- Finish

Choose the type of retention policy to create

Locations can be specified dynamically with an adaptive scope using attributes or properties, or if you know the specific target locations, you can select them individually from a list. An advantage of using an adaptive scope to determine target locations is that it will automatically update where it's applied based on the attributes or properties you define.

Adaptive

After selecting adaptive policy scopes, which consist of attributes or properties (e.g. 'Department' or 'Site URL') that define the users, groups, or sites in your org, you'll choose supported locations containing the content you want to retain. The policy will automatically update to match the criteria defined in the scopes.

Static

You'll choose locations containing the content you want to retain. If locations change after this policy is created (for example if a SharePoint site is added or removed), you'll need to manually update the policy.

The publication of labels isn't different than if you would use Data Lifecycle Management, so to keep this chapter as simple as possible our new label will be targeted to 1 site. If you would like to know more about adaptive scopes, I would urge you to take a look at the chapter 'Adaptive Scopes'.

Publish labels so users can apply them to their content.

- Choose labels to publish
- Administrative Units
- Scope
- Publish to users and groups
- Name your policy
- Finish

Choose where to publish labels

When published, users in your organization will be able to apply this label to items in the locations you choose.

You can set up data connectors to import content from non-Microsoft apps like Slack, WhatsApp and many more, for use with this solution. [Set up now](#)

All locations. Includes content in Exchange email, Office 365 groups, OneDrive and SharePoint documents.

Let me choose specific locations.

Status	Location	Included	Excluded
<input type="radio"/> Off	Exchange mailboxes		
<input type="radio"/> Off	SharePoint classic and communication sites		
<input type="radio"/> Off	OneDrive accounts		
<input checked="" type="radio"/> On	Microsoft 365 Group mailboxes & sites	1 microsoft 365 group	None
		Edit	

Select “let me choose specific locations” and in the “included” column, select your desired SharePoint site or the Microsoft 365 group where it’s based upon. I’ve included my “High Five Sky High” Microsoft 365 Group / SharePoint site.

Next up, name your policy and take a look at the review screen. Finish the wizard.

The view from the user / Records Manager

Let’s take a look on the user side of things. Ideally, the upcoming tasks will be performed by someone of your records management department. For this demo, I’ve created a folder called “Project Codename Kitesurfing” within a standard channel with an equal name in the SharePoint Site / Team “High Five Sky High”.

The screenshot shows a SharePoint list view for a 'Standard Channel'. A folder named 'Project Codename Kitesurfing' is selected. A red arrow points to the three-dot menu icon next to the folder. The 'Properties' pane on the right displays the folder's details, including a dropdown for 'Label' with the option 'Mark Item as Record - Retain for 2 Day' highlighted with a red box. A tooltip at the bottom of the pane states: 'Marks items as a record so they become immutable. These records are retained for 2 days after which they are placed on a disposition list for review before they are deleted.'

Since the project has come to an end I want to mark the folder and its contents as a record. To do this, I click the 3 dots, choose “details” so the flyout menu on the right side opens. Under “Label”, our label “Mark Item as Record – Retain for 2 Days” (note that it also shows the configured description) can now be selected.

The top screenshot shows a 'Standard Channel' list view with a folder named 'Project Codename Kitesurfing'. A red arrow points to the padlock icon on the right of the folder's row, indicating it is now a record. The bottom screenshot shows the 'Project Codename Kitesurfing' folder expanded, displaying three documents: 'Kitesurfing Project Explanation.docx', 'Project Budget.xlsx', and 'Secret Presentation.pptx'. Each document has a red box around its padlock icon, and a red arrow points to the padlock icon of the first document, further confirming they are all marked as records.

Immediately, the folder gains a padlock icon on the right, telling us that the folder and its contents are now marked as record. So let's bring some of the theory above into practice. Let's try to edit the contents of the document.

A screenshot of a Microsoft Word document titled 'Kitesurfing Project Explanation'. A red arrow points to a yellow status bar at the bottom left that says 'Read-only This document is read-only.'

As expected, this fails. The yellow ribbon tells us the document is read-only.

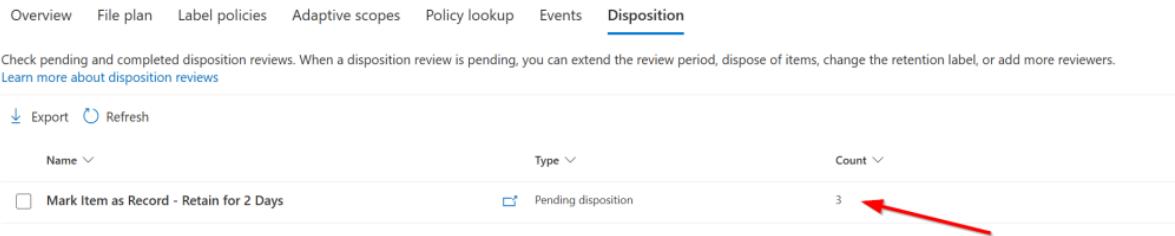
A screenshot of a 'Standard Channel' list view showing the 'Project Codename Kitesurfing' folder. A red box highlights the 'Kitesurfing Project Explanation.docx' file. A red box also highlights a tooltip that appears when the file is deleted, stating '1 item wasn't deleted from Project Codename Kitesurfing' and 'Kitesurfing Project Explanation.docx This operation has been disabled by policy'.

Trying to delete the document also fails with error "This operation has been disabled by policy". Ain't that great?

Fast forward a couple of days...

...And the files will show up on your disposition list. The disposition list can be seen by navigating to the Microsoft Purview portal and clicking the “Records Management” button on the left hand side. Within Records Management, navigate to “Disposition”.

Records management

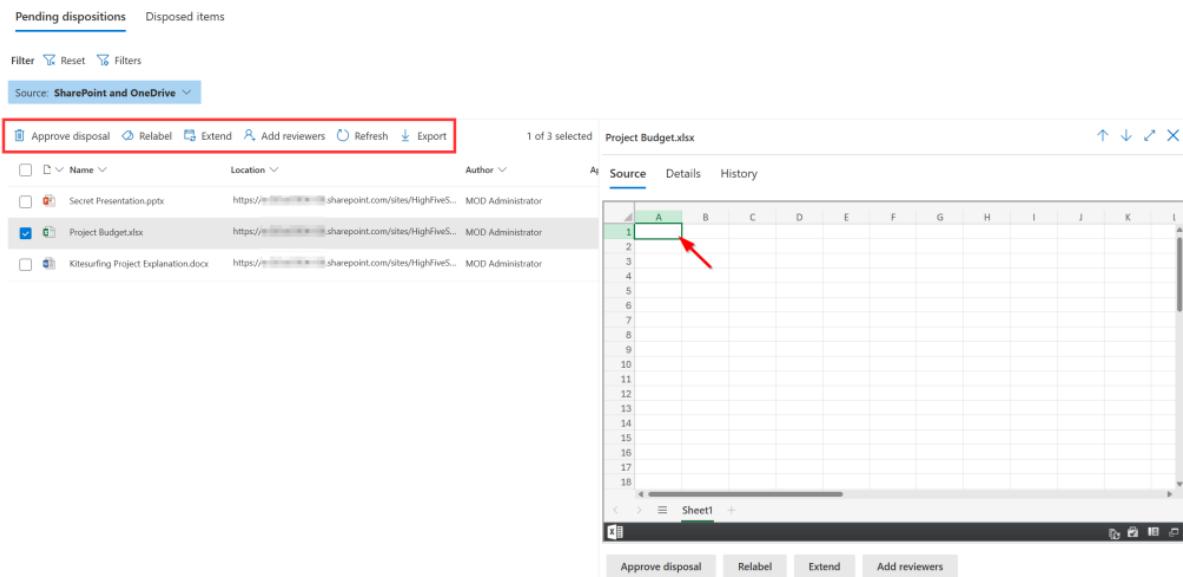


Check pending and completed disposition reviews. When a disposition review is pending, you can extend the review period, dispose of items, change the retention label, or add more reviewers. [Learn more about disposition reviews](#)

Export Refresh

Name	Type	Count
Mark Item as Record - Retain for 2 Days	Pending disposition	3

As can be seen, the name of your file plan and the current number of documents included on it's disposition list is present.



Pending dispositions Disposed items

Filter Reset Filters

Source: SharePoint and OneDrive

Approve disposal Relabel Extend Add reviewers Refresh Export

Name	Location	Author
Secret Presentation.pptx	https://.../sites/HighFiveS...	MOD Administrator
<input checked="" type="checkbox"/> Project Budget.xlsx	https://.../sites/HighFiveS...	MOD Administrator
<input type="checkbox"/> Kitesurfing Project Explanation.docx	https://.../sites/HighFiveS...	MOD Administrator

1 of 3 selected Project Budget.xlsx

Source Details History

1 A B C D E F G H I J K L

Approve disposal Relabel Extend Add reviewers

Now when you click your file plan followed by the “open in a new window” button you’ll see an overview of the items currently pending disposition. You can filter the list by using the filter button on the top of the screen.

Note that if you want to see previews of your items you’ll have to be a member of the “Content Explorer Content Viewer” role group, which can be found in the Purview portal, Roles & Scopes, Permissions. If you only want to see a file’s label but not its contents, you can add a user to the “Content Explorer List Viewer” role group.

In the portal, you are able to execute the following actions on the items in the disposition list:

- Approve disposal, after which the files will be removed. This process can take up to 15 days before the files are permanently deleted from the location where they’re stored.
- Relabel the files, after which a new label will be applied (with different settings).
- Extend the disposition date, after which the files will show up on the disposition list again.

- Add reviewers, to select someone from your organization to also review the pending disposition.
- Export a comma separated file that contains the files name, location, title or subject, tagname (label name), comment, name of the account that deleted the content and the delete date.

Differences between Retention Labels, Records and Regulatory Records

Action	Retention label	Record - locked	Record - unlocked	Regulatory record
Edit contents	Allowed	Blocked	Allowed	Blocked
Edit properties, including rename	Allowed	Allowed ¹	Allowed	Blocked
Delete	Allowed ²	Blocked	Blocked	Blocked
Copy	Allowed	Allowed	Allowed	Allowed
Move within container ³	Allowed	Allowed	Allowed	Allowed
Move across containers ³	Allowed	Allowed if never unlocked	Blocked	Blocked
Open/Read	Allowed	Allowed	Allowed	Allowed
Change label	Allowed	Allowed - container admin only	Blocked	Blocked
Remove label	Allowed	Allowed - container admin only	Blocked	Blocked

Image source: Microsoft

As can be seen in the table above, records have certain allowed actions but also restrictions. Where all actions are allowed with retention labels (that are configured using Data Lifecycle Management), items that have the “record” status cannot -for example- be edited or deleted. Of course this make sense, because the exact function of declaring an item as a record is making sure its contents cannot be altered anymore.

But even when an item is declared a record, properties can still be edited by default. If you don’t want this, there is a tenant option to block this behavior. This option can be found at Records management – Records management settings – Retention labels – Allow editing of record properties.

Another -more restrictive option- is to declare items as a regulatory record. The table above shows exactly how restrictive this option is. When comparing regulatory records to normal records, items declared as regulatory records properties cannot be changed, the items declared as regulatory record cannot be moved across containers (SharePoint sites, OneDrive accounts, and Exchange mailboxes), but the most restrictive option is that labels cannot be changed or removed anymore (not even by global administrators!).

Also good to note about regulatory records is:

- Retention period of an item declared as regulatory record cannot be made shorter, but only be extended.
- Retention labels that configure items as regulatory records do not support auto-labeling policies.
- Retention labels that configure items as regulatory records cannot be applied to items that are checked-out.

When you've gone through all of the above and you still want to declare items as regulatory records, you will have to enable the option first by connecting to the Office 365 Security & Compliance PowerShell and running the following cmdlet after which you can select marking content as regulatory record in the retention label wizard:

```
Set-RegulatoryComplianceUI -Enabled $true
```

Be sure to take a look at the following Microsoft Learn articles before you decide you need to have regulatory records:

- <https://learn.microsoft.com/en-us/purview/declare-records>
- <https://learn.microsoft.com/en-us/purview/records-management#compare-restrictions-for-what-actions-are-allowed-or-blocked>

How to see where labels are applied?

If you want to get an overview of where your retention labels are applied, take a look at the Data Lifecycle Management (DLM) chapter as the steps to get an overview of locations where your record labels are applied is the same as with DLM!

Sensitivity Labels

One of the most distinct features of Purview is of course sensitivity labeling, which is part of the information protection section in the Purview portal. Before we head off to configuring sensitivity labeling and dive into what it looks like from a users perspective, let's first talk about what sensitivity labels are.

Introduction

You can think of a sensitivity label like a stamp, which you can apply to content like documents, email and meetings. The cool thing is that the sensitivity label is added in clear text to the metadata of the files, so it travels together with the content (hence the reference to the stamp 😊). Because it's stored in clear text, applications and services can use the sensitivity label to apply logic to it. Examples of this logic is adding a watermark to a document, protecting content from being openend by unauthorized people or content being protected from being sent outside your organization. This protection part can be done by Microsoft 365 or a third-party application. But a sensitivity label by itself can inform users of the sensitivity level of a certain item.

There are various automatic methods of applying labels to your content, but for this chapter we'll focus on manually adding labels to content so we understand how the basic process works before we move on to some form of automatic labeling.

Labels should go with the flow of your document in a certain process. An example process could be like this.

1. You create a document.
2. You apply the label "internal" which shows users that the document is not ready to be sent outside your organization.
3. You work with internal colleagues on the document.
4. The document is ready to be distributed outside your organization to be edited by external coworkers.
5. You apply a new label "external".
6. You finish working on the document and publish it to it's final destination. Depending on the process, you apply another label or remove the sensitivity label.

When looking at the implementation of sensitivity labels, there's a certain learning curve that you must guide your users through. Because of this, I recommend to first introduce sensitivity labels to your organization without applying restrictions based on it to give your users a visual introduction that a sensitivity label is applied and learn how to work with them. After this, it's possible to extend the labels functionality with content protection or some other restriction.

In short, this would be a great way to introduce sensitivity labels to your organization:

1. Create sensitivity labels (less is more).
2. Implement sensitivity labels and guide your users on how to apply and use them.
3. Monitor your environment and keep educating your users

4. Collect feedback from your users and use it to optimize labels.
5. When labeling is well embedded in your organization, start restricting content based on your sensitivity labels.

Lastly, you should remember the following about sensitivity labels:

- Content can hold 1 sensitivity label, however it can be combined with a retention label.
- Sensitivity labels can be seen by members of your organization but are not visible to guests or users from other organizations.

Configuring Sensitivity Labels

Configuration of sensitivity labels that can be manually applied consists of steps:

1. Create and configure your sensitivity label.
2. Publish the label to end-users.

Create and configure your sensitivity label

To start, navigate to the Microsoft Purview portal and navigate to Information Protection, Labels.

① Your organization has not turned on the ability to process content in Office online files that have encrypted sensitivity labels applied and are stored in OneDrive and SharePoint. You can turn on here, but note that additional configuration is required for Multi-Geo environments. [Learn more](#)

② You can now create sensitivity labels with privacy and access control settings for Teams, SharePoint sites, and Microsoft 365 Groups. To do this, you must first [complete these steps](#) to enable the feature.

Right away, we are greeted with 2 messages where number 1 is telling us to turn on the ability to process content in Office online files that have encrypted sensitivity labels applied and are stored in OneDrive and SharePoint. When you don't use a multi-geo Microsoft 365 environment, you can safely enable the feature. Feature 2 however makes it possible to create sensitivity labels and access control settings for Teams, SharePoint sites and Microsoft 365 Groups. However in this chapter, we stick to the basics and go with manual labeling.

	Name	Priority	Scope	Created by	Last modified
<input type="checkbox"/>	Personal	0	File, Email		Jun 23, 2024 4:04:56 AM
<input type="checkbox"/>	Public	1	File, Email		Jun 23, 2024 4:04:57 AM
<input type="checkbox"/>	General	2	File, Email		Jun 23, 2024 4:05:00 AM
<input type="checkbox"/>	Confidential	3	File, Email		Jun 23, 2024 4:05:04 AM
<input type="checkbox"/>	Highly Confidential	4	File, Email		Jun 23, 2024 4:05:13 AM
<input type="checkbox"/>	Confidential - Finance	5	File, Email, Meetings	Megan Bowen	Jun 24, 2024 11:50:14 AM

Click "Create a label" in the "Labels" screen.

Provide basic details for this label

The protection settings you choose for this label will be immediately enforced on the items or content containers to which it's applied. Labeled files will be protected wherever they go, whether they're saved in the cloud or downloaded to a computer.

Name * ⓘ

Display name * ⓘ

Label priority ⓘ

ⓘ By default, this label will be assigned the highest priority, but you can change this after it's created. X

Description for users * ⓘ

Description for admins ⓘ

Label color ⓘ

Dark
Light
Grey
Pink
Blue
Green
Yellow
Orange
Red

Here we can configure the basics of our label: it's name, display name, priority, a description for users and administrators and a label color. They're all quite basic, however we need to discuss the label priority. If you take a look back to the second screenshot of this chapter you'll see that there are various sensitivity labels already present. They all have a priority set ranging from 0 to 5, where 5 is the highest priority (most restrictive) and 0 being the lowest (least restrictive). In an upcoming step in the label creation wizard, we're able to set an option to make it mandatory for users to provide a justification when they lower the sensitivity label on an item. This is why you should think the prioritization of your labels through before going all-in on configuration mode 😊. You can change the priority of a label using the 3-dotted-menu next to a label to move it up or down.

This behavior however, doesn't apply to sublabels. Sublabels are taken into account with automatic labeling, which is for another time.

Define the scope for this label

Labels can be applied directly to items (such as files, emails, meetings), containers like SharePoint sites and Teams, Fabric and Power BI items, schematized data assets, and more. Let us know where you want this label to be used so you can configure the applicable protection settings. [Learn more about label scopes](#)

Items
Be aware that restricting the scope to only files or emails might impact access control settings and where the label can be applied. [Learn more](#)

Files
Protect files created in Word, Excel, PowerPoint, and more.

Emails
Protect messages sent from all versions of Outlook.

Meetings
Protect calendar events and meetings scheduled in Outlook and Teams.

ⓘ Parent label will automatically inherit meeting scope from sub labels

Groups & sites
Configure privacy, access control, and other settings to protect labeled Teams, Microsoft 365 Groups, and SharePoint sites.
 ⓘ To apply sensitivity labels to Teams, SharePoint sites, and Microsoft 365 Groups, you must first [complete these steps](#) to enable the feature.

Schematized data assets (preview)
Apply labels to files and schematized data assets in Microsoft Purview Data Map. Schematized data assets include SQL, Azure SQL, Azure Synapse, Azure Cosmos, AWS RDS, and more.
 ⓘ To apply this label to schematized data assets, you must first turn on labeling for Microsoft Purview Data Map. You can do this from the Labels page. [Learn more about labeling for Microsoft Purview Data Map](#)

On with configuration of our label. Here we set where we want our labels to be available. I've chosen for items, which are files, emails and meetings. Groups and sites is not yet possible because -as discussed at the beginning of this chapter- we have some configuration to do for this to be possible. Schematized data assets can be used to apply labels to files and schematized data assets in SQL, Azure SQL, Azure Synapse, Azure Cosmos, AWS RDS and more and hence, are out of scope for this ebook.

Choose protection settings for the types of items you selected

The protection settings you configure will be enforced when the label is applied to items in Microsoft 365.

Control access
Control who can access and view labeled items.

Apply content marking
Add custom headers, footers, and watermarks to labeled items.

Protect Teams meetings and chats
Configure protection settings for Teams meetings and chats.
 ⓘ To protect Teams meetings and chats, your org must have a Teams Premium license. [Learn more about Teams Premium](#)

In the next screen, we can choose the protection settings we want, which can be “control access”, “apply content marking” or “protect teams meeting and chats”. Note that it isn’t mandatory to pick an option here. Sensitivity labels can easily be applied without having protection settings added to them, which will only show you the document has a sensitivity label applied showing its name and description. For this demo however I’d like to add a little bit of content marking so we can see what that looks like.

Content marking

Add custom headers, footers, and watermarks to content that has this label applied. [Learn more about content marking](#)

ⓘ All content marking will be applied to documents but only the header and footer will be applied to email messages. If you chose to configure meeting settings for this label, the header and footer will also be applied to meeting invites.

Content marking



Add a watermark

Customize text

Add a header

Customize text

Add a footer

Customize text

Highly Confidential - For you eyes only!

Content marking

Add custom headers, footers, and watermarks to content that has this label applied.

ⓘ All content marking will be applied to documents but only the header and footer will be applied to email messages. If you chose to configure meeting settings for this label, the header and footer will also be applied to meeting invites.

Content marking

Add a watermark

Customize text

Add a header

Customize text

Add a footer

Customize text

Customize footer text

This text will appear as a footer on labeled email messages and documents.

Footer text * ⓘ

Highly Confidential - For you eyes only!

Font size

10

Font color

Red

Align text

Center

To be precise, I've chosen to add a footer, which can be displayed on documents, or email. Please note that a watermark will not be present on email messages, should you configure this.

Auto-labeling for files and emails

When users edit Office files or compose, reply to, or forward emails from Outlook that contain content matching the conditions you choose here, we'll automatically apply this label or recommend that they apply it themselves. [Learn more about auto-labeling for Microsoft Purview](#)

ⓘ To automatically apply this label to files that are already saved (in SharePoint and OneDrive) or emails that are already processed by Exchange, you must create an auto-labeling policy. [Learn more about auto-labeling policies](#)

Auto-labeling for files and emails



In the next screen auto-labeling (or semi-auto labeling that can give users some recommendation to apply a label) can be configured for files and emails. We'll keep this as is for now since we want to take a look at what manual labeling looks like.

Define protection settings for groups and sites

These settings apply to teams, groups, and sites that have this label applied. They don't apply directly to the files stored in those containers. [Learn more about these settings](#)

- Privacy and external user access**
Control the level of access that internal and external users will have to labeled teams and Microsoft 365 Groups.
- External sharing and Conditional Access**
Control external sharing and configure Conditional Access settings to protect labeled SharePoint sites.
- Private teams discoverability and shared channel settings**
Decide whether private teams will be discoverable in searches and control the types of teams that can be invited to shared channels.

This next screen is of use when labeling for groups and sites is configured (remember action number 2 that popped up when opening the labeling console?). So, skip for now.

Auto-labeling for schematized data assets (preview)

Automatically apply this label to schematized data assets in Microsoft Purview Data Map that contain the sensitive info types you choose here. You can automatically label database columns in SQL, Azure SQL, Azure Synapse, Azure Cosmos, AWS RDS, and various other data sources governed by Microsoft Purview Data Map. [Learn more about auto-labeling for schematized data assets](#)

Auto-labeling for schematized data assets (preview)



Last but not least, auto-labeling for schematized data assets can be enabled. As discussed, we'll leave this out of scope for now. After this, check the settings displayed in the summary and click "create label" when you are satisfied with its contents.

✓ Your sensitivity label was created

Creating the label is just the first step in labeling and protecting content. To make this label available to users in your org, you can auto-apply it to specific content and publish it to users' apps.

Next steps

- Publish label to users' apps
Create a publishing policy to show the label in Office apps, SharePoint, Teams, and Microsoft 365 Groups so users can apply it to content themselves. [Learn more about publishing labels](#)
- Don't create a policy yet
You can publish or auto-apply this label later.

Recommended resources based on your settings

- [Review prerequisites](#) to get the most out of your access control settings
- [Review prerequisites](#) for applying sensitivity labels to Fabric and Power BI content.
- [Review a Microsoft Purview Data Map tutorial](#) on how to start scanning assets and automatically apply this label

When your label is created, the wizard will ask you whether to publish the label to users' apps or to do this later. To see all of the wizard pages, I've opted to not create a label policy yet.

Label policies

Create sensitivity label policies to publish one or more labels to your users' Office apps (like Outlook and Word), SharePoint sites, and Office 365 groups. Once published, users can apply the labels to protect their content. [Learn more about sensitivity label policies](#)



Name	Priority	Created by	Last modified	3 items
<input type="checkbox"/> Global sensitivity label policy	0 - lowest		Jun 23, 2024 4:05 AM	
<input type="checkbox"/> Confidential-Finance Policy	1	Megan Bowen	Jun 24, 2024 11:50 AM	
<input type="checkbox"/> Highly Confidential Policy	2 - highest	Megan Bowen	Jun 24, 2024 11:50 AM	

Instead, point your browser to Information Protection, Label Policies and click "publish label".

Choose sensitivity labels to publish

When published, the labels you choose here will be available in specified users' Office apps (Word, Excel, PowerPoint, and Outlook), SharePoint and Teams sites, and Microsoft 365 Groups.

Sensitivity labels to publish

Classified Information

[Edit](#)

In the first step of the wizard, select the label that we've just created. In my case, it's the "Classified Information" label. We'll skip the admin units screen for now and continue to the next screen.

Publish to users and groups

The labels you selected will be available for the users, distribution groups, mail-enabled security groups, and Microsoft 365 Groups you choose here.

Location	Scope	
<input checked="" type="checkbox"/> Users and groups	All users & groups	Edit Edit Scope for users and groups

Here, we can select the users and groups we want the label to be available for. In my case I select "all users and groups". You can however, narrow this down if you'll want.

Policy settings

Configure settings for the labels included in this policy.

<input checked="" type="checkbox"/> Users must provide a justification to remove a label or lower its classification
Users will need to provide a justification before removing a label or replacing it with one that has a lower-order number. You can use activity explorer to review label changes and justification text.
<input type="checkbox"/> Require users to apply a label to their emails and documents
Users will be required to apply labels before they can save documents or send emails (only if these items don't already have a label applied). <small> ⓘ Support and behavior for this setting varies across apps and platforms. Learn more about managing sensitivity labels</small>
<input type="checkbox"/> Require users to apply a label to their Fabric and Power BI content
Users will be required to apply labels to unlabeled content they create or edit in Fabric and Power BI. Learn more about mandatory labeling in Fabric and Power BI
<input checked="" type="checkbox"/> Provide users with a link to a custom help page
If you created a website dedicated to helping users understand how to use labels in your org, enter the URL here. Learn more about this help page <input type="text" value="https://www.dominiquehermans.com"/>

Now an interesting part. In this next screen we can make sure users have to provide a justification when they remove a label or lower its classification (which we talked about earlier in this chapter). This I'll select. It's also possible to require users to apply a label to their emails, documents or Fabric and Power BI content. This I will not select for this demo because not each item will have to have a sensitivity label in my case. Lastly, we can provide users with a link to a custom help page, which I configured with the URL of my blog so you'll get the idea. Ideally you'll provide users to a place on your intranet for example where you guide them through the process of labeling items with a set of best practices that are present for your company.

Default settings for documents

Apply a default label to documents

The label you choose will automatically be applied to Word, Excel, and PowerPoint documents when they're created or modified. Users can always select a different label to better match the sensitivity of their document. [Learn which Office app versions support this setting](#)

Default label

▼

In this next screen another important choice should be made. Will you automatically apply a default sensitivity label to (Word, Excel and PowerPoint) **documents** when they're created or modified or do you want your items to be “label less” by default? I opt for the latter in this demo.

Default settings for emails

Apply a default label to emails

The label you choose will automatically be applied to new and existing, unlabeled emails. Users can always change the default label before they send the message. [Learn which Outlook versions support these settings](#)

Default label

▼

Require users to apply a label to their emails (i)

Inherit label from attachments

If a label is applied to an email, then an attachment with a higher priority label is added to an email, this setting replaces the existing label with the label from the attachment. If multiple labeled attachments are added, the highest priority label will be applied. If the email isn't already labeled, it will inherit the highest priority label from any attachments. [Learn more about label inheritance](#)

- Email inherits highest priority label from attachments
- Recommend users apply the attachment's label instead of automatically applying it

Next, a default label can be selected for **emails** if you want. Also, you can configure whether an email message inherits the highest priority label from the attachments it contains (handy stuff!) and whether you want to give users a recommendation to apply the labels to the email message or to automatically apply it. I choose to give the user a recommendation to label the email message with the highest priority label of its attachments.

Default settings for meetings and calendar events

Apply a default label to meetings and calendar events

The label you choose will automatically be applied to new and existing, unlabeled meetings and calendar events. Users can always change the default label before they create the event or meeting. [Learn which Teams and Outlook clients support these settings](#)

Default label

▼

Require users to apply a label to their meetings and calendar events (i)

Another default action to be considered is up next. Do you want to apply a default label to **meetings and calendar events**? I don't in my case but think this through for your environment. Another screen appears to select a default label for **Fabric and Power BI content** which we'll won't use in this demo.

Name your policy

Name *

Policy - Classified Information

Enter a description for your sensitivity label policy

Policy - Classified Information

Next, name your policy, review the summary and "submit" when you're happy with the results it shows!

✓ New policy created

It can take up to 24 hours to publish the labels to the selected users' apps.

Next steps

[Review data classification reports](#) to see how labels are being used

[Read guidance](#) on how to educate users about sensitivity labels

Do take note of the message stating that it can take up to 24 hours to publish the labels to the selected users' apps.

The view from the user

Files

Now that configuration of our sensitivity label is done and 24 hours have passed, let's take a look at the user side of things.

The screenshot shows the Microsoft Teams interface on the left and a SharePoint document library on the right. The SharePoint library is titled 'General' and contains several documents. The document 'Ad Slogans.docx' is selected, and its details are shown on the right. A red box highlights the document's row in the list, and a red arrow points to the three-dot menu icon next to it. Another red box highlights the 'Sensitivity' field, which is currently set to 'None'.

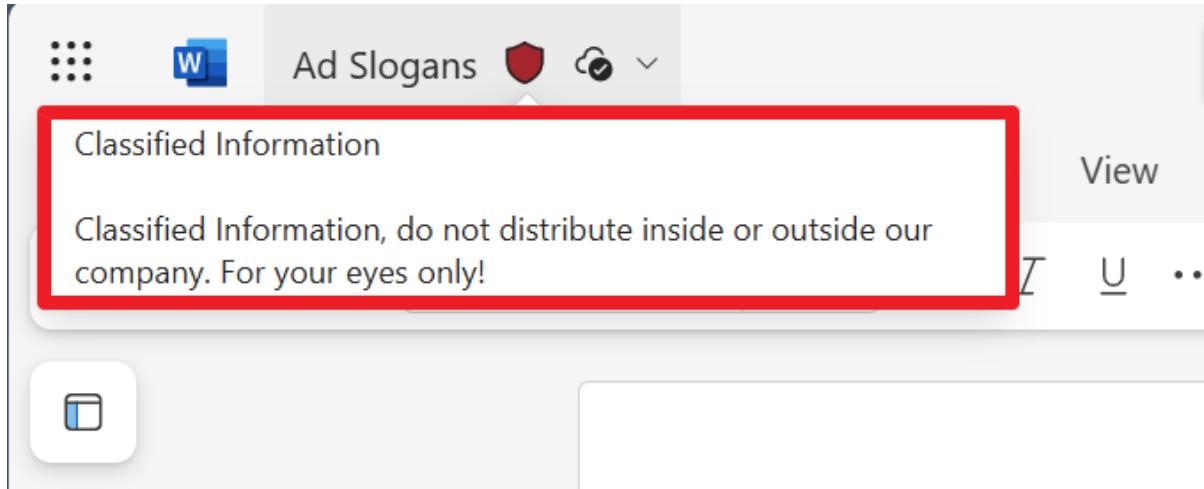
In this screen I've opened Teams as a user in my tenant and navigated to the team that holds a document called "Ad Slogans.docx". When you press the 3-dotted menu to the right and select "details" you immediately see that at this point, there is no sensitivity label applied.

The screenshot shows a SharePoint document list with the 'Ad Slogans.docx' document selected. A red arrow points to the three-dot menu icon next to the document. A dropdown menu for sensitivity labels is open on the right, showing various options like 'None', 'Personal', 'Public', and 'Classified Information'. The 'None' option is currently selected. A red box highlights the 'Classified Information' option in the list.

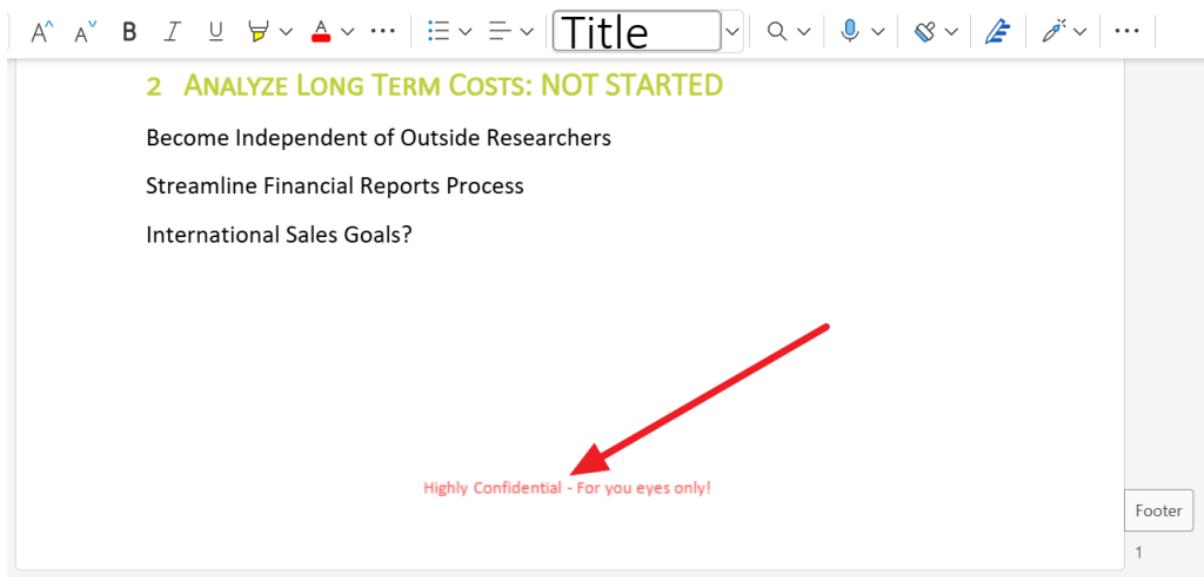
It's possible to add a sensitivity label to the document from here, but let's open the document and check out the view from Word Online.

The screenshot shows the Microsoft Word Online editor with the document 'Ad Slogans' open. The ribbon at the top shows various formatting options. A red box highlights a note at the bottom of the page that reads 'Classified Information, do not distribute inside or outside our company. For your eyes only!'. A red arrow points to this note. A sensitivity label dropdown menu is open on the right, showing options like 'Personal', 'Public', 'General', 'Confidential', 'Highly Confidential', and 'Confidential - Finance'. The 'Classified Information' option is highlighted with a red box. A red box also highlights the 'Classified Information' label in the note.

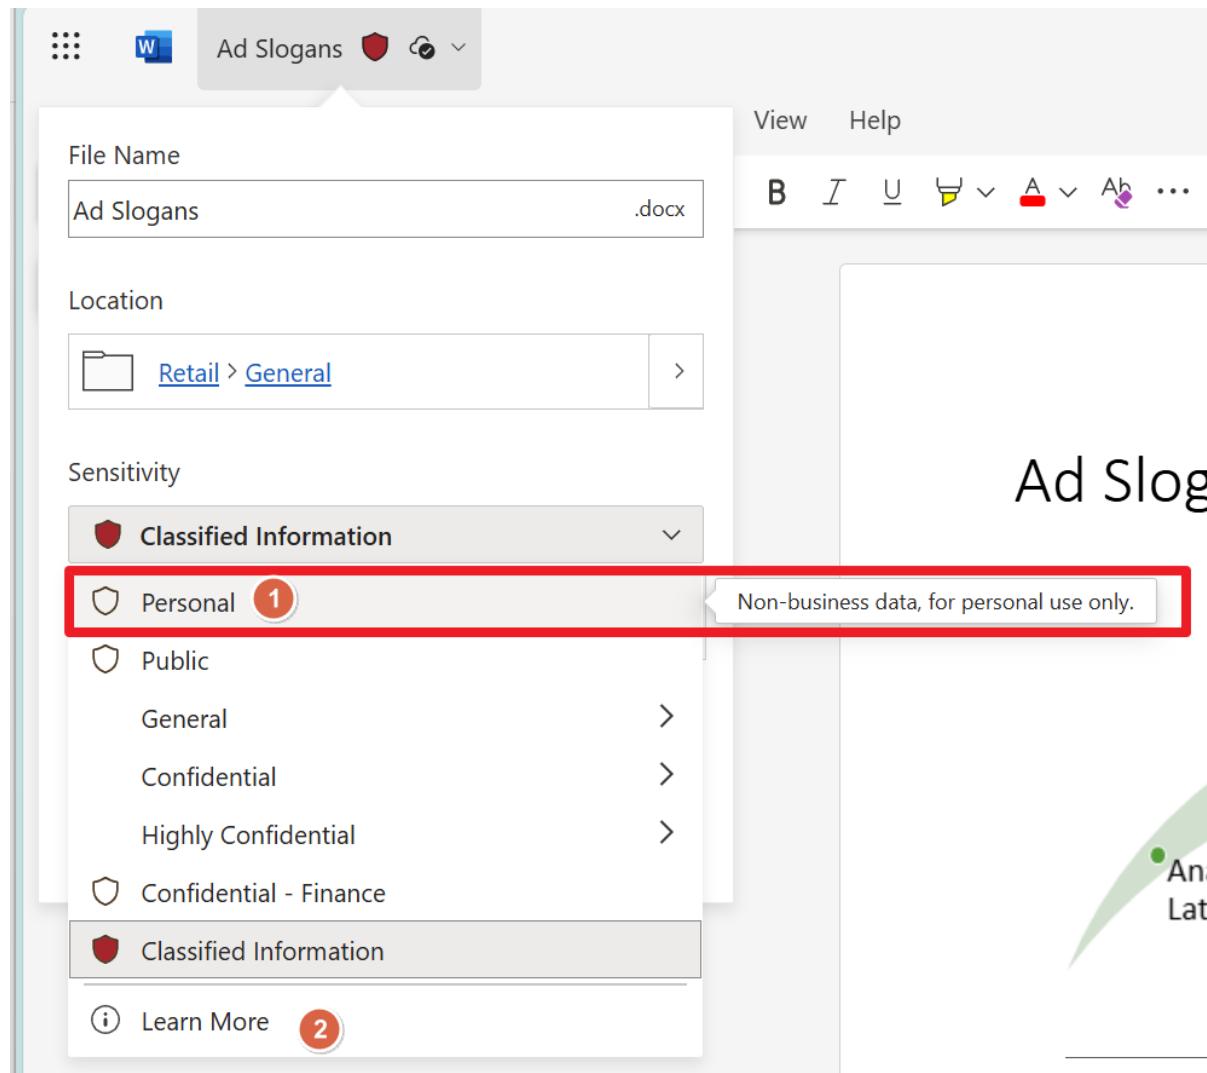
When we open the document in Word (Online in my case, however the desktop app responds just the same), you will see a “stamp” icon that can be used to stamp a sensitivity label to the document. In my case there are several, and amongst them is the “classified information” label we created earlier. Notice the red color we selected and when we hover over it, the description for the user is shown. Let’s click it to apply the label to the file.



When the label is applied (or stamped if you will), the top bar immediately shows the updated sensitivity information (in red, like we selected) and here also the description for the user is shown. Great!

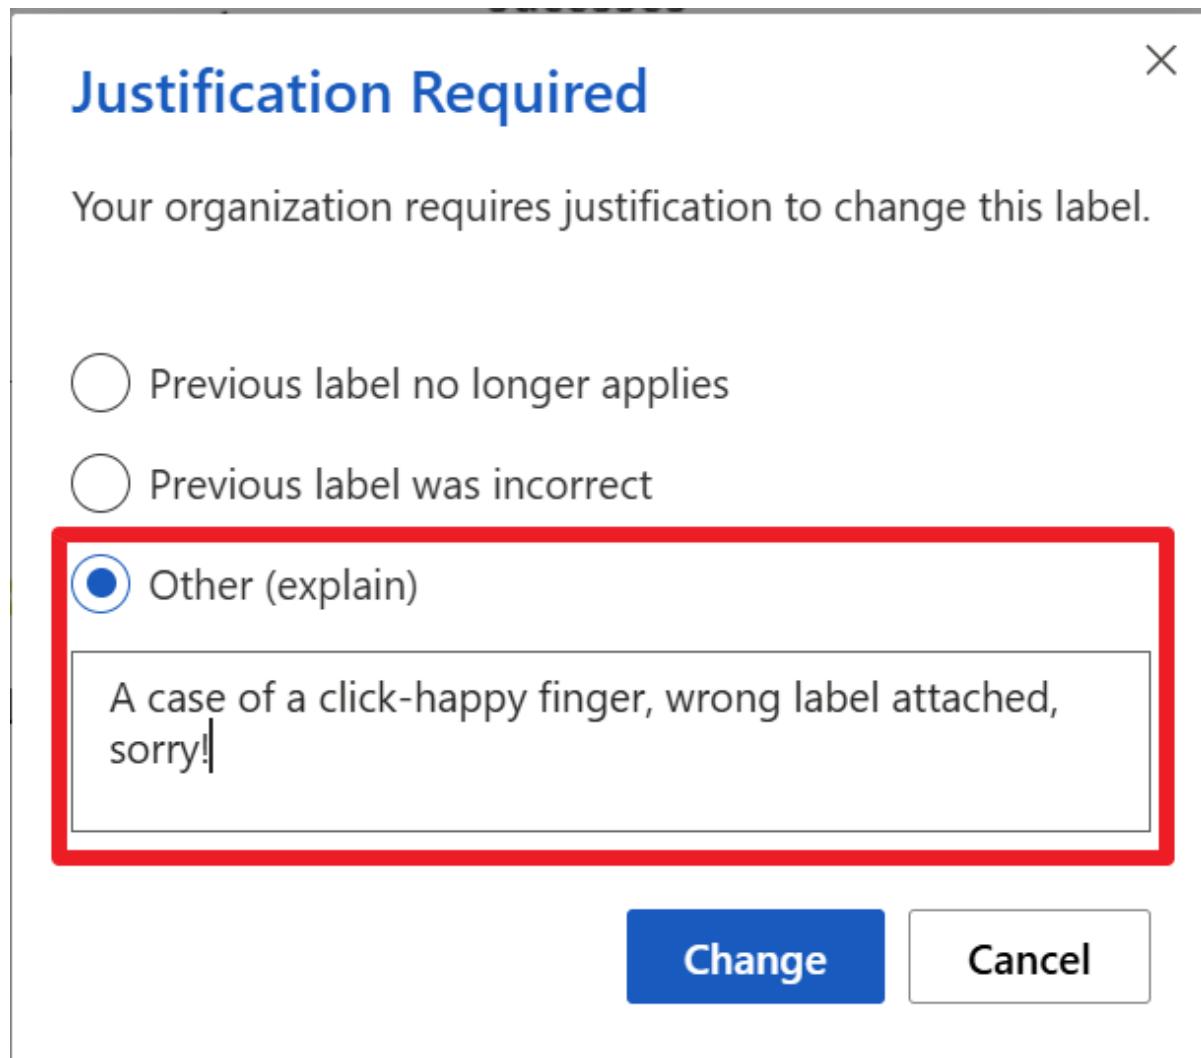


Remember we configured the label to show content marking in the form of a footer? Here you go, that’s also present!



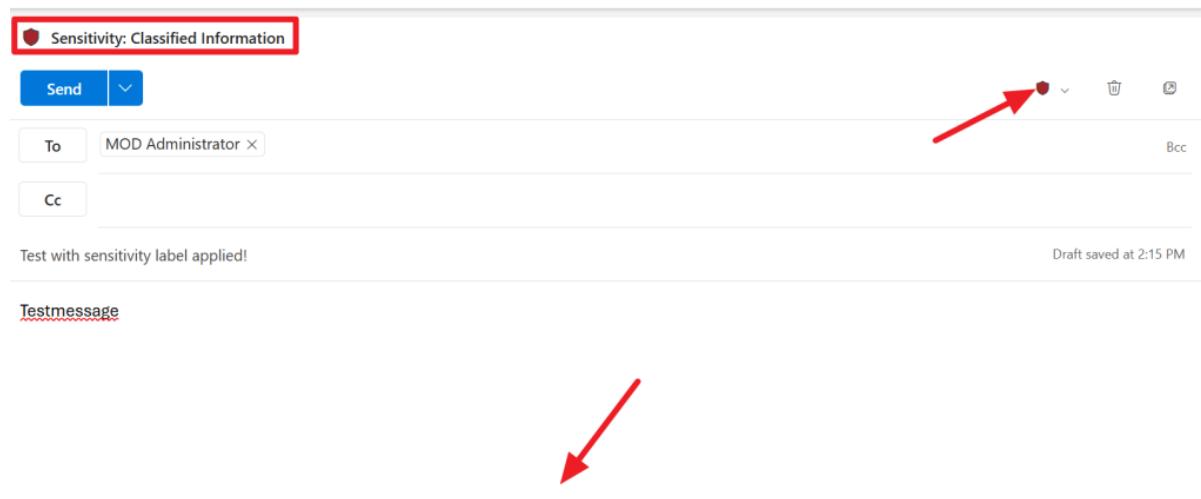
When you click the sensitivity label on the top left, you can click the “Learn More” (2) link. This will guide you to the page that you set up to give your users more information on your labeling process (in this demo, I configured it to navigate to the URL of my blog).

When we want to lower the sensitivity to -for example- the “personal” sensitivity label, the following message appears:



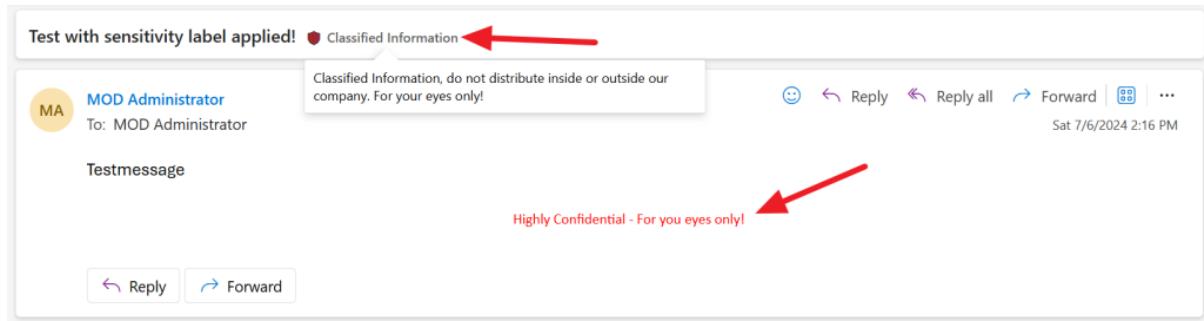
Remember we configured the option to require users to justify their action if they lower the document's sensitivity? That's the setting that shows you this dialog box.

Emails



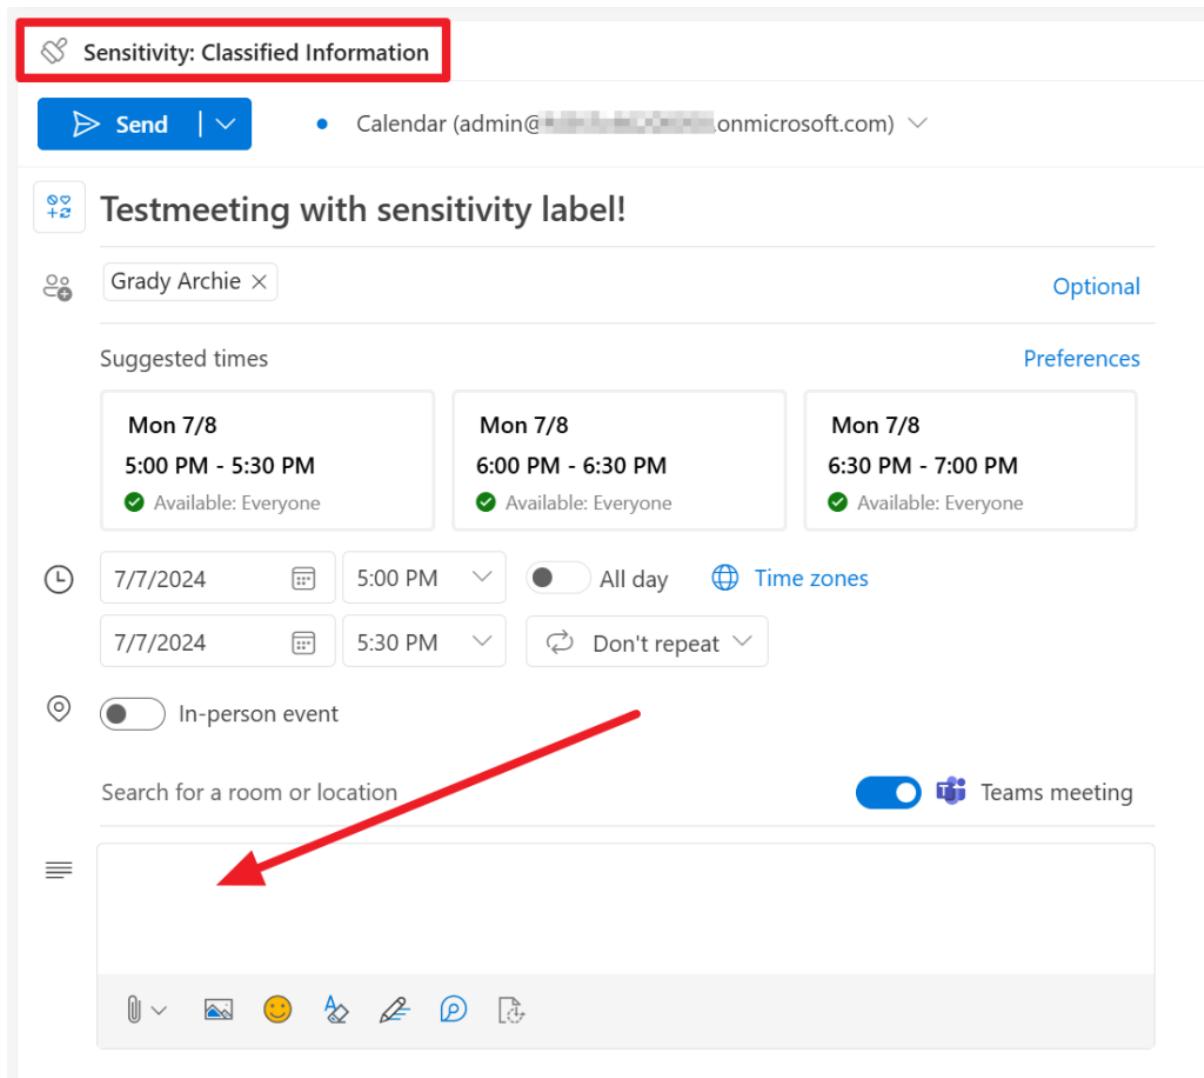
When drafting an email and configuring it with our label by clicking on the button on the top right, we can see immediately that our label is stamped on the email, but the configured markup

(footer) is not shown. This is because it's applied after the message is being sent, and can be seen when the message has been delivered to the recipient:



When the message is delivered, the footer is attached and the sensitivity label of course remains on the email message.

Meetings



The fact that markup is not applied on email also goes for meetings. After stamping the meeting with a sensitivity label and sending it out, the recipient does have the markup attached to the message and meeting details:

The screenshot shows a Microsoft Teams meeting details page and an adjacent email organizer. The meeting details include the title 'Testmeeting with sensitivity label!', a date and time of 'Sun 7/7/2024 8:00 AM - 8:30 AM', and a note 'MOD Administrator invited you. Didn't respond 1'. The 'Meeting Details' tab is selected. The email organizer shows an email from 'MOD Administrator' with the subject 'Testmeeting with sensitivity label!' and a note 'Sent on Saturday, 7/6/2024 at 5:23 AM'. The note in the organizer is highlighted with a red arrow and the text 'Highly Confidential - For your eyes only!'. The 'Email organizer' tab is selected.

There you have it! As a closing note, always take a look at the “Minimum versions for sensitivity labels in Office apps” document on [Microsoft Learn](https://learn.microsoft.com/en-us/purview/sensitivity-labels-versions) (<https://learn.microsoft.com/en-us/purview/sensitivity-labels-versions>) to check whether your required behavior is supported by the app you are using. As an example, Outlook for the Web does not support the option to inherit an attachments sensitivity label and apply it to your email message.

In conclusion

- You can think of a sensitivity label as a **stamp**, that you “stamp” onto documents to give it a certain level of “sensitivity”. The higher the number, the higher the sensitivity.
- It teaches your users to think about with who to share documents that have a certain sensitivity.
- The label travels with your documents.
- You can attach **logic** to sensitivity labels. Examples of this logic is adding watermarks or protecting content from being openend by unauthorized people.
- Create a **sensitivity label** to provide a name, priority, description for users and admins, color, scope and protection settings for your label.
- Create a **label policy** to publish your label, provide additional settings like the fact that users should provide a justification to remove or lower a labels classification, require users to apply labels in certain locations or to provide users with a link to a custom help page. You can also use the policy to apply a default label to documents, emails, meeting and calendar events and whether attachments in an email should dictate the label of the email.

Did you make it till the end? Then now you can call yourself an expert on configuring sensitivity labels!

Data Loss Prevention (DLP)

Data Loss Prevention (DLP) in Microsoft Purview can be used to prevent your users from oversharing information. Oversharing information is the process of accidentally or purposely sharing information with recipients that are not allowed to have or view this information.

While there are various ways to implement DLP with Microsoft Purview, one of the main ones is by leveraging DLP Policies. When taking a look at the DLP Policies pages in Purview, Microsoft gives us the following introductory text:

Use data loss prevention (DLP) policies to help identify and protect your organization's sensitive info. For example you can set up policies to help make sure information in email and docs isn't shared with the wrong people.

DLP makes use of so-called Sensitive info types, often referred to as SIT's. Microsoft includes an enormous list of SIT's you can use out-of-the-box. You can look at SIT's as the engine in DLP, as each SIT holds a pattern and/or logic for recognizing content. A few examples of these classifiers are:

- Credit Card Number
- U.K. Physical Addresses
- User Login Credentials

If you want to take a look at the entire list I would recommend to navigate to the Purview portal, Data classification, Classifiers, Sensitive info types. At the time of writing this ebook the list consists of 324 items. If the pattern/logic for classifying a piece of information in your environment isn't present, you also have the option to create a SIT yourself.

Plan first, implement second

A few questions that you should ask yourself before heading out and start configuring DLP enthusiastically:

- Which stakeholders do I have to interview or include in my team to select the right types of sensitive information for my company?
- How do I validate my setup before enforcing policies on users?
- What is my scope? What is included in my scope and what is not?
- What is my business planning and what is my planning on technology?
- How do I introduce DLP to my end-users. Should I include training or adoption?

Generally, the following step-by-step action plan would give you the opportunity to get some insights and let your users get acquainted with the introduction of DLP in their day to day jobs.

1. Design the DLP policies you would like to configure.
2. Make a deployment plan.
3. Configure your policies to first run in simulation mode.

4. Use statistics and information from simulation mode to finetune your policies.
5. Place policies in production according to your deployment plan.

Configure Data Loss Prevention (DLP)

Policies

 Data loss prevention settings

Use data loss prevention (DLP) policies to help identify and protect your organization's sensitive info. For example you can set up policies to help make sure information in email and docs isn't shared with the wrong people. [Learn more about DLP](#)

 If your role group permissions are restricted to a specific set of users or groups, you'll only be able to manage policies for those users or groups. [Learn more about role group permissions](#) X

View role groups

 Create policy	 Export	 Refresh		5 items	 Search	 Customize columns
<input type="checkbox"/> Name			Priority	Last modified	Status	
<input type="checkbox"/> Default Office 365 DLP policy			: 0	Jun 23, 2024 4:04 AM	On	
<input type="checkbox"/> Default policy for Teams			: 1	Jun 23, 2024 4:04 AM	On	
<input type="checkbox"/> Default policy for devices			: 2	Jun 23, 2024 4:04 AM	On	
<input type="checkbox"/> U.S. Financial Data			: 3	Jun 23, 2024 9:06 AM	On	
<input type="checkbox"/> General Data Protection Regulation (GDPR)			: 4	Jun 23, 2024 9:07 AM	On	

After you've done your homework we can start configuring the DLP policies. To start, we move to the Purview portal and select 'Data Loss Prevention', 'Policies'. Here, let's click 'Create policy'.

Start with a template or create a custom policy

Choose an industry regulation to see the DLP policy templates you can use to protect that info or create a custom policy start from scratch. If you need to protect labeled content, you'll be able to choose labels later. [Learn more about DLP policy templates](#)

 **Check out our new enhanced policy templates.** These enhanced templates extend several of the original templates by also detecting named entities (such as full names and physical addresses). Just look for the templates labeled 'Enhanced' to start protecting even more personal data.

 Search for specific templates	All countries or regions ▼	
<p>Categories</p> <ul style="list-style-type: none">  Enhanced  Financial  Medical and health  Privacy  Custom 	<p>Regulations</p> <ul style="list-style-type: none"> Australia Financial Data Canada Financial Data France Financial Data Germany Financial Data Israel Financial Data Japan Financial Data PCI Data Security Standard (PCI DSS) Saudi Arabia - Anti-Cyber Crime Law Saudi Arabia Financial Data U.K. Financial Data 	<p>U.S. Financial Data</p> <p>Helps detect the presence of information commonly considered to be financial information in United States, including information like credit card, account information, and debit card numbers.</p> <p>Protect this information:</p> <ul style="list-style-type: none"> • Credit Card Number • U.S. Bank Account Number • ABA Routing Number

In the first screen of the wizard, we are able to select one of the regulations that are provided by Microsoft and are divided into various categories. In this example, we want to protect U.S.

Financial info that is found in data. As said, if there's no category or regulation that fits your needs, you can also create your own custom rules to match data in your environment.

Name your DLP policy

Create a DLP policy to detect sensitive data across locations and apply protection actions when the conditions match.

Name *

U.S. Financial Data - DominiqueHermans.com

Description

Helps detect the presence of information commonly considered to be financial information in United States, including information like credit card, account information, and debit card numbers.

In the next step, name your policy and create a matching description and select Admin Units if you use them.

Choose where to apply the policy

We'll apply the policy to data that's stored in the locations you choose.

① Protecting sensitive info in on-premises repositories (SharePoint sites and file shares) is now in preview. Note that there are prerequisite steps needed to support this new capability. [Learn more about the prerequisites](#)

Location	Scope	
<input type="checkbox"/> Exchange email	Turn on location to scope	
<input checked="" type="checkbox"/> SharePoint sites	All sites	Edit
<input type="checkbox"/> OneDrive accounts	Turn on location to scope	
<input checked="" type="checkbox"/> Teams chat and channel messages	All users & groups	Edit
<input type="checkbox"/> Devices	Turn on location to scope	
<input type="checkbox"/> Instances	Turn on location to scope	
<input type="checkbox"/> On-premises repositories	Turn on location to scope	
<input type="checkbox"/> Power BI workspaces	Turn on location to scope	

Select where you want your policy to apply. In this demo, I choose to go with all SharePoint sites and Teams chat and channel messages for all users and groups. If you need, you can filter down this list.

Define policy settings

Decide if you want to use the default settings from the template you selected to quickly set up a policy or configure custom rules to refine your policy further.

Review and customize default settings from the template. [\(1\)](#)

Credit Card Number

U.S. Bank Account Number

ABA Routing Number

Create or customize advanced DLP rules [\(2\)](#)

Now for the fun part. The template you choose in the first step of the wizard came with a predefined set of rules and conditions to match data in your environment. You can go with the defaults here, but it's also possible to edit them. Let's go with 'create or customize advanced DLP rules' to see what's in store for us in these policies.

Customize advanced DLP rules

The rules here are made up of conditions and actions that define the protection requirements for this policy. You can edit existing rules or create new ones.

[+ Create rule](#)

2 items

Name	Status
1 Low volume of content detected U.S. Financial Data - DominiqueHe	On
2 High volume of content detected U.S. Financial Data - DominiqueH	On

1 [Low volume of content detected U.S. Financial Data - DominiqueHe](#)

Conditions
Content contains any of these sensitive info types:
Credit Card Number
U.S. Bank Account Number
ABA Routing Number

Content is shared from Microsoft 365 with people outside my organization

Actions
Notify users with email and policy tips
Send alerts to Administrator

2 [High volume of content detected U.S. Financial Data - DominiqueH](#)

Conditions
Content contains any of these sensitive info types:
Credit Card Number
U.S. Bank Account Number
ABA Routing Number

Content is shared from Microsoft 365 with people outside my organization

Actions
Notify users with email and policy tips
Restrict access to the content for external users
Send incident reports to Administrator
Send alerts to Administrator

The rules that are included in the 'U.S. Financial data' policy match data on the following sensitive info types, or SIT's, that can be seen in the screenshot above tagged with label '1' and '3':

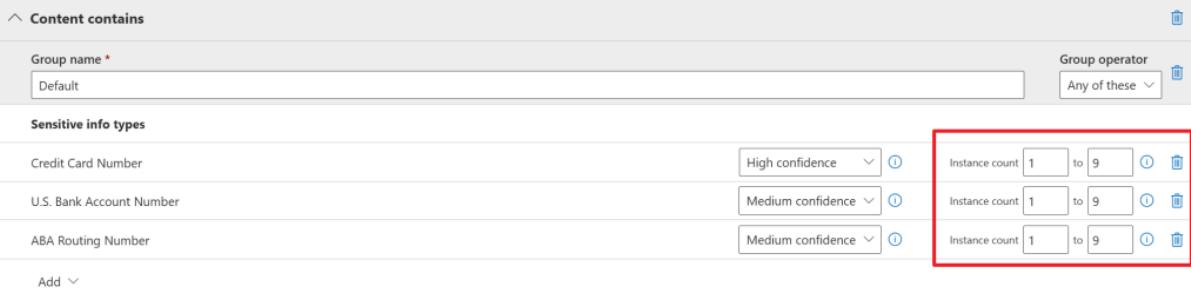
- Credit Card Number OR
- U.S. Bank Account Number OR
- ABA Routing Number

AND

- Content is **shared** from Microsoft 365 with people **outside my organization**.

As you can see, the policy is configured twice, with the difference being that the first one is for ‘low volume of content detected U.S. Financial data’ and the second being ‘High volume of content detected U.S. Financial Data’.

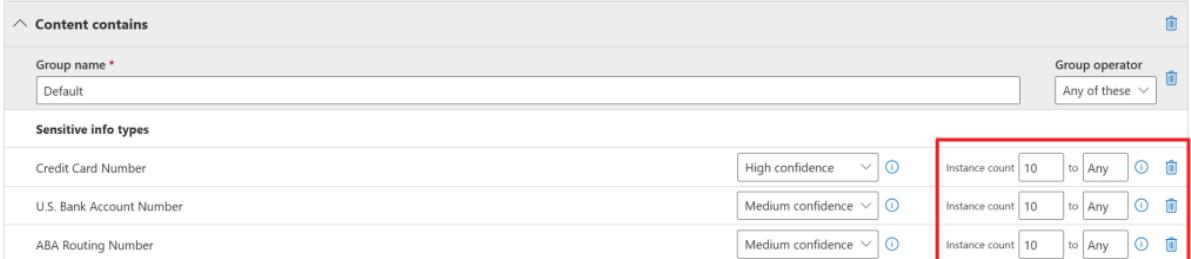
When we press the edit button on the policies, we can see the exact difference between ‘high volume’ and ‘low volume’.



The screenshot shows the 'Content contains' policy configuration. The 'Group name' is 'Default' and the 'Group operator' is 'Any of these'. The 'Sensitive info types' section lists three items: 'Credit Card Number' (High confidence), 'U.S. Bank Account Number' (Medium confidence), and 'ABA Routing Number' (Medium confidence). Each item has an 'Instance count' field with the value '1 to 9' highlighted with a red box. There is also an 'Add' button.

When looking at the ‘low volume’ policy properties first, you’ll see that a sensitive info type has to be found in a certain document between 1 and 9 times for it to fall into this category and for the policy to be applied. When looking back at the screenshot before this one, actions applied with the ‘low volume’ policy are:

- Notify users with email and policy tips
- Send alerts to administrator



The screenshot shows the 'Content contains' policy configuration. The 'Group name' is 'Default' and the 'Group operator' is 'Any of these'. The 'Sensitive info types' section lists three items: 'Credit Card Number' (High confidence), 'U.S. Bank Account Number' (Medium confidence), and 'ABA Routing Number' (Medium confidence). Each item has an 'Instance count' field with the value '10 to Any' highlighted with a red box. There is also an 'Add' button.

Now let’s take a look at the properties for the ‘high volume’ policy. We can see that a SIT has to be found between 10 and ‘any’ times for this policy to apply. Actions applied for this policy are of course much stricter:

- Notify users with email and policy tips
- Restrict access to the content for external users
- Send incident reports to administrator
- Send alerts to administrator

A detour into Sensitive Information Types (SIT’s)

For additional knowledge, let’s take a look at one of the SIT’s used in the policies, ‘Credit Card Number’.

Credit Card Number

Description

Detects credit card numbers for American Express, Diner's Club, Discover Card, JCB, BrandSmart, Mastercard, and Visa.

High confidence, Medium confidence (recommended)

Primary element

Function processor: Func_credit_card

Supporting element

Minimum 1 matches should be found from following elements:

Keyword from keyword list: Keyword_cc_verification

Keyword from keyword list: Keyword_cc_name

Function processor: Func_expiration_date

Low confidence

Primary element

Function processor: Func_credit_card

The 'Credit Card Number' SIT consists of the following elements:

- Name
- Description
- Primary and supporting elements that define the logic to find that a number found is actually a Credit Card Number. The more elements that are matched, the more certain Purview is that a match is actually a Credit Card Number in this case. If all of the above elements match, it's considered to be a High confidence or Medium confidence match. If only the 'Function processor: Func_credit_card' matches, it's considered to be a Low confidence match. You can see that this confidence level is also used in the policy properties where we looked at 'high volume' and 'low volume' properties.

Back to configuration of our DLP policy

Policy mode

You can test this policy before turning it on to verify if it needs improvement or if it meets all your objectives. If you turn the policy on right away, you can edit it later and safely test those changes in simulation mode.

- Run the policy in simulation mode**
We'll show you items that match the policy's conditions to help you evaluate its impact. Your data won't be affected; the policy stays off while in simulation mode. [Learn more about simulation mode](#)
 - Show policy tips while in simulation mode.
 - Turn the policy on if it's not edited within fifteen days of simulation
- Turn the policy on immediately**
After the policy is created, it'll take up to an hour before any changes are enforced.
- Leave the policy turned off**
Decide to test or activate the policy later.

In the 'policy mode' screen we can configure the policy to:

- Run in simulation mode. This doesn't enable the policy but will run a scan that you can evaluate before enabling the policy. This gives you the opportunity to fine-tune your policy before placing it into production. We'll take a look at this later. While in simulation, you can show policy tips to users, which I'll enable. It's also possible to turn the policy on automatically if it's not edited within 15 days. I'll skip this for now.
- Turn the policy on right away (not recommended).
- Leave the policy turned off.

Let's go with the first option, and do show policy tips in simulation mode.

All set! Review the summary and finish the wizard.

Policies Data loss prevention settings

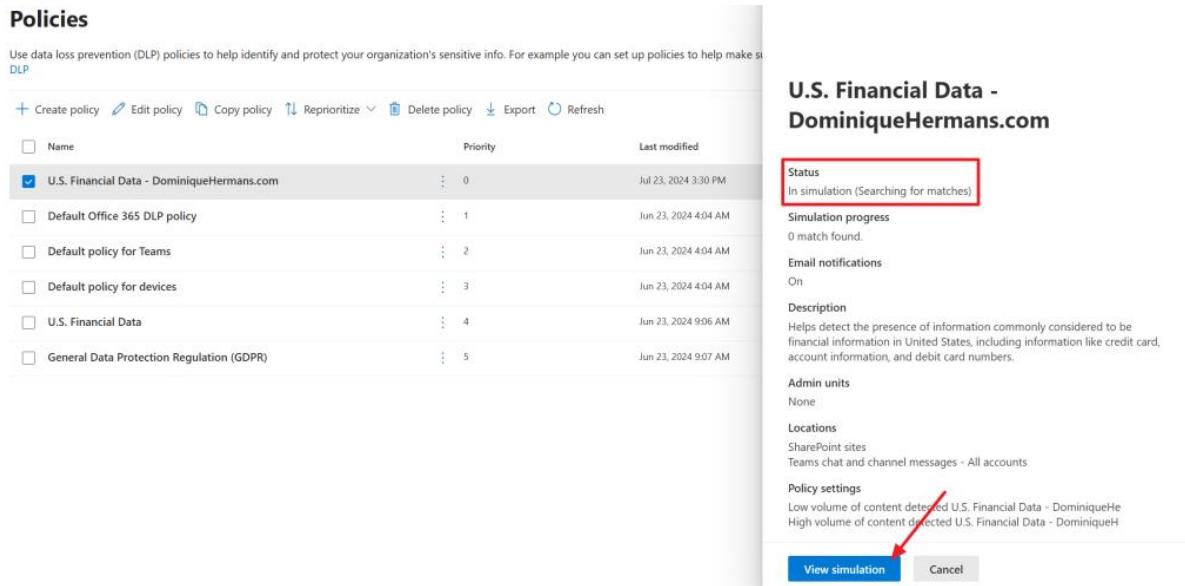
Use data loss prevention (DLP) policies to help identify and protect your organization's sensitive info. For example you can set up policies to help make sure information in email and docs isn't shared with the wrong people. [Learn more about DLP](#)

[+ Create policy](#) [Export](#) [Refresh](#) 6 items Search [Customize columns](#)

<input type="checkbox"/> Name	Priority	Last modified	Status
<input type="checkbox"/> U.S. Financial Data - DominiqueHermans.com	0	Jul 23, 2024 3:30 PM	In simulation with notifications
<input type="checkbox"/> Default Office 365 DLP policy	1	Jun 23, 2024 4:04 AM	On
<input type="checkbox"/> Default policy for Teams	2	Jun 23, 2024 4:04 AM	On
<input type="checkbox"/> Default policy for devices	3	Jun 23, 2024 4:04 AM	On
<input type="checkbox"/> U.S. Financial Data	4	Jun 23, 2024 9:06 AM	On
<input type="checkbox"/> General Data Protection Regulation (GDPR)	5	Jun 23, 2024 9:07 AM	On

Back in the policies screen, I configure the new policy to have top priority for it to be applied before any other matching policy is applied.

Reviewing simulation results



The screenshot shows the Microsoft 365 DLP Policies page. On the left, a list of policies is displayed with columns for Name, Priority, and Last modified. The 'U.S. Financial Data - DominiqueHermans.com' policy is selected and highlighted. On the right, a detailed view of this policy is shown, including its status as 'In simulation (Searching for matches)', simulation progress (0 match found), and policy settings (Low volume of content detected). A red box highlights the 'Status' section, and a red arrow points to the 'View simulation' button.

Name	Priority	Last modified
U.S. Financial Data - DominiqueHermans.com	0	Jul 23, 2024 3:30 PM
Default Office 365 DLP policy	1	Jun 23, 2024 4:04 AM
Default policy for Teams	2	Jun 23, 2024 4:04 AM
Default policy for devices	3	Jun 23, 2024 4:04 AM
U.S. Financial Data	4	Jun 23, 2024 9:06 AM
General Data Protection Regulation (GDPR)	5	Jun 23, 2024 9:07 AM

U.S. Financial Data - DominiqueHermans.com

Status
In simulation (Searching for matches)

Simulation progress
0 match found.

Email notifications
On

Description
Helps detect the presence of information commonly considered to be financial information in United States, including information like credit card, account information, and debit card numbers.

Admin units
None

Locations
SharePoint sites
Teams chat and channel messages - All accounts

Policy settings
Low volume of content detected U.S. Financial Data - DominiqueHermans.com
High volume of content detected U.S. Financial Data - DominiqueHermans.com

View simulation Cancel

When clicking the newly created policy, we can see the status is 'In simulation (searching for matches)'. Now it's time to take a well-deserved break and enjoy your weekend and come back a few days later. The exact time it takes to run the simulation depends on the size of your environment and data it contains.

Now before continuing, make sure to grant yourself the 'Data Classification Content Viewer' role by adding yourself to the 'Content Explorer Content Viewer' role group. If you don't, you won't be able to look into sensitive info details.

When returning and matches are actually found, click the 'view simulation' button.

U.S. Financial Data - DominiqueHermans.com

In progress

Simulation overview Items for review Alerts

Simulation progress

We're scanning specific locations, such as SharePoint, Teams in real-time for items that match the policy's conditions. Predicted matches will appear on the [Items for review](#) page when ready.

[Learn more about simulation mode](#)

Total matches

2 matches found

Scanning per location

Location	Status
Share Point	Real-time
Teams	Real-time

[Review matching items](#)

View insights from

Share Point Teams

Share Point

Scan status
2
Matching files found

Matching items

Please review the file selected for you from the recent matches below.

Files

Contoso_Online_Class_Registration.xlsx

Contoso_Online_Presales_Report_M400.xlsx

Top sites with most matches

<https://m365...sharepoint.com/sites/newempl...> 2

Here you'll find an overview of the matches found in your organization. In my case, 2 matching files are found which both are found in 1 SharePoint site.

U.S. Financial Data - DominiqueHermans.com

In progress

[Turn the policy on](#) [Download report](#) [Edit the policy](#) [Delete the policy](#) [Restart the simulation](#)

Simulation overview Items for review Alerts

Review items that match your policy to decide whether it will be applied to the right content. Files listed are a sample of the total matching files from each site included in the policy (approximately 100 files per site for each policy rule). Matching emails will continue to appear here as they're sent to recipients.

Filter [Reset](#) [Filters](#)

Date match was detected: 7/23/2024-8/3/2024

Location: SharePoint sites, +1

Rules: Any

[Export](#) [Refresh](#) 1 of 2 selected [Customize columns](#)

File name Rule Location Sensitive info t

Contoso_Online_Class_Registration.xlsx

[↑](#) [↓](#) [↗](#) [↘](#) [X](#)

Source Match summary Metadata

The content in this file matches 1 of 2 policy rules.

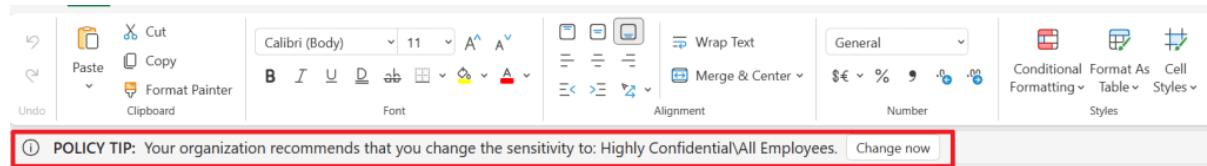
High volume of content detected U.S. Financial Data - DominiqueH

Sensitive info type High confidence Medium confidence Low confidence

Credit Card Number 100 100 100

File name	Rule	Location	Sensitive info t	Source	Match summary	Metadata
<input checked="" type="checkbox"/> Contoso_Online_Class_Regist...	High volume of content detected	SharePoint	Credit Card Num			
<input type="checkbox"/> Contoso_Online_Presales_Re...	High volume of content detected	SharePoint	Credit Card Num			

When you click on the 'items for review' tab followed by the 'match summary' tab, you can select an item. At the right hand side, the portal shows which sensitive information type was found in your document by confidence level. So it could be true that a SIT was identified using the 'low confidence' logic, but not by the medium or high confidence logic that was configured in the SIT. In this document however, the configured SIT (Credit Card Number) was found 100 times by low, medium and high logic defined in the SIT.



The screenshot shows the Microsoft Word ribbon with a policy tip message in the status bar: 'POLICY TIP: Your organization recommends that you change the sensitivity to: Highly Confidential\All Employees. Change now.' The message is highlighted with a red box.

M83	X	✓	fx								
A	B	C	D	E	F	G	H	I	J	K	
1	Order Number	Order Total	Order Date	Order Time	Class Registration	Customer	Zip	Credit Card	Security Type	CVV2/CID	Expiration
2	976890907	\$ 99.00	6/1/2015	5:46:31 PM	Yes	Berg, Karen	11716	3734-6623-7621-570	CID	3287	10/2020
3	976890831	\$ 198.00	6/1/2015	7:15:40 PM	Yes	Tollevsen, Bjorn	32102	4539-6333-7119-3889	CVV2	263	01/2020
4	976890396	\$ 149.00	6/2/2015	4:17:43 PM	Yes	Ceckova, Lucia	52643	6441-2087-8163-6199	CID	736	09/2017
5	976890436	\$ 99.00	6/4/2015	1:32:48 PM	Yes	Saraiva, Jose	72441	4539-5907-4789-7710	CVV2	361	01/2018
6	976890764	\$ 49.00	6/5/2015	1:11:51 PM	Yes	Balasubramanian, Dev	30447	4556-9525-3115-1167	CVV2	525	01/2018
7	976890381	\$ 99.00	6/5/2015	4:18:17 PM	Yes	Andrews, Ben	34178	3447-9311-2642-660	CID	9558	03/2019

Let's take a look at the 'Contoso_Online_Class_Registration.xlsx' sheet. Here, credit card info is found. Actually a list of 100 entries (which I shortened for readability) can be found in the sheet which matches the score above. As configured, the policy tip is also shown to the user. It's also possible to customize the policy tip if necessary.

At this point you would review some items and if all seems ok, you can enable the policy so it becomes active in your environment.

Enabling the policy

U.S. Financial Data - DominiqueHermans.com In progress

→ Turn the policy on Download report Edit the policy Delete the policy Restart the simulation

We can enable the policy by going back to the 'view simulations' screen and selecting 'Turn the policy on'. Confirm the dialog box and refresh the page. According to the documentation, it can take about 1 hour for changes to apply across your environment.

- **Verifying results from the users perspective**

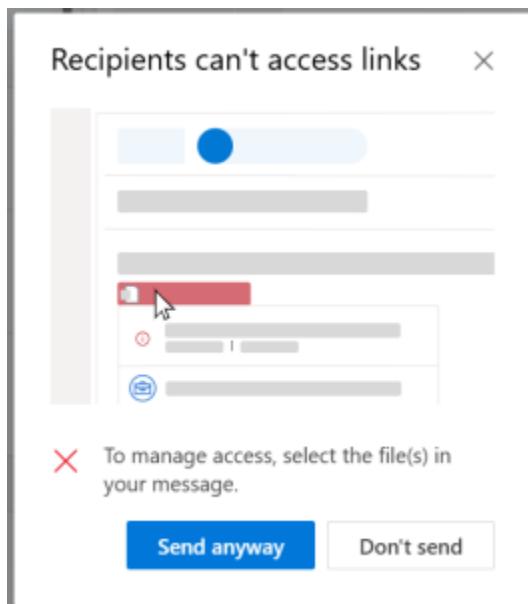
Now, let's see how the protective actions of our policy look like when viewed from a users point of view. Let's do this by trying to share a document that's considered to be a sensitive information type (Credit Card Number) via SharePoint.

The screenshot shows a SharePoint list titled 'Training audit'. A file sharing dialog is open, showing a recipient 'christiec@...' and a warning message: 'This item contains sensitive information. It can't be shared with people outside your organization. View policy tip'. The list table shows various items with status columns like 'Review process', 'Review date', and 'Region'.

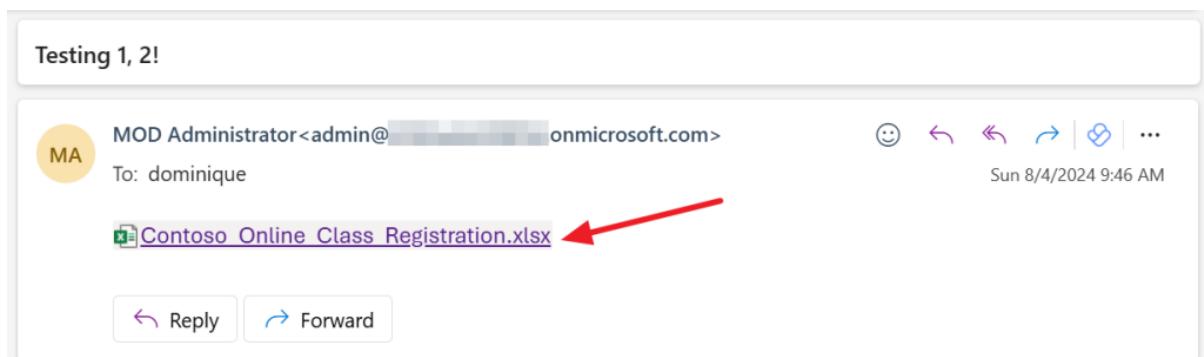
As you can see, the file can't be shared because it contains sensitive information, which was found using the SIT we configured in our DLP policy!

The screenshot shows an Outlook for Web email draft. The 'To' field is populated with 'dominique@...'. A red box highlights a warning message: 'The following recipient is outside your organization: dominique@... onmicrosoft.com.'. Below the message is a 'Send' button and a red box highlights a warning message for a file attachment: 'People outside your organization won't be able to open Contoso_Online_Class_Registration.xlsx. Manage access | Ignore'. The attachment menu is open, showing options like 'Open in Web app', 'Attach as a copy', 'Change to full URL', 'Edit hyperlink', and 'Copy link'.

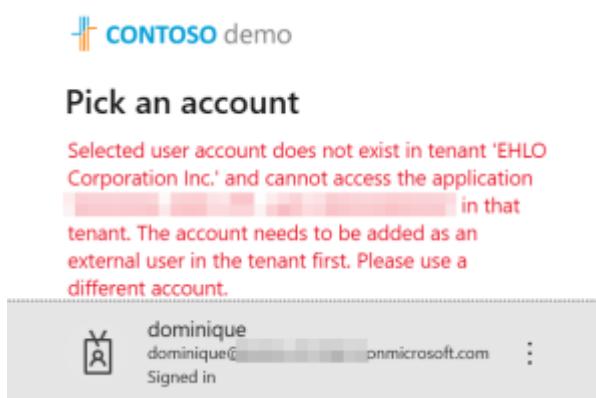
Now let's try to work around this by using Outlook for Web. When adding a link to the document, here also a message appears that people outside your organization won't be able to open the file.



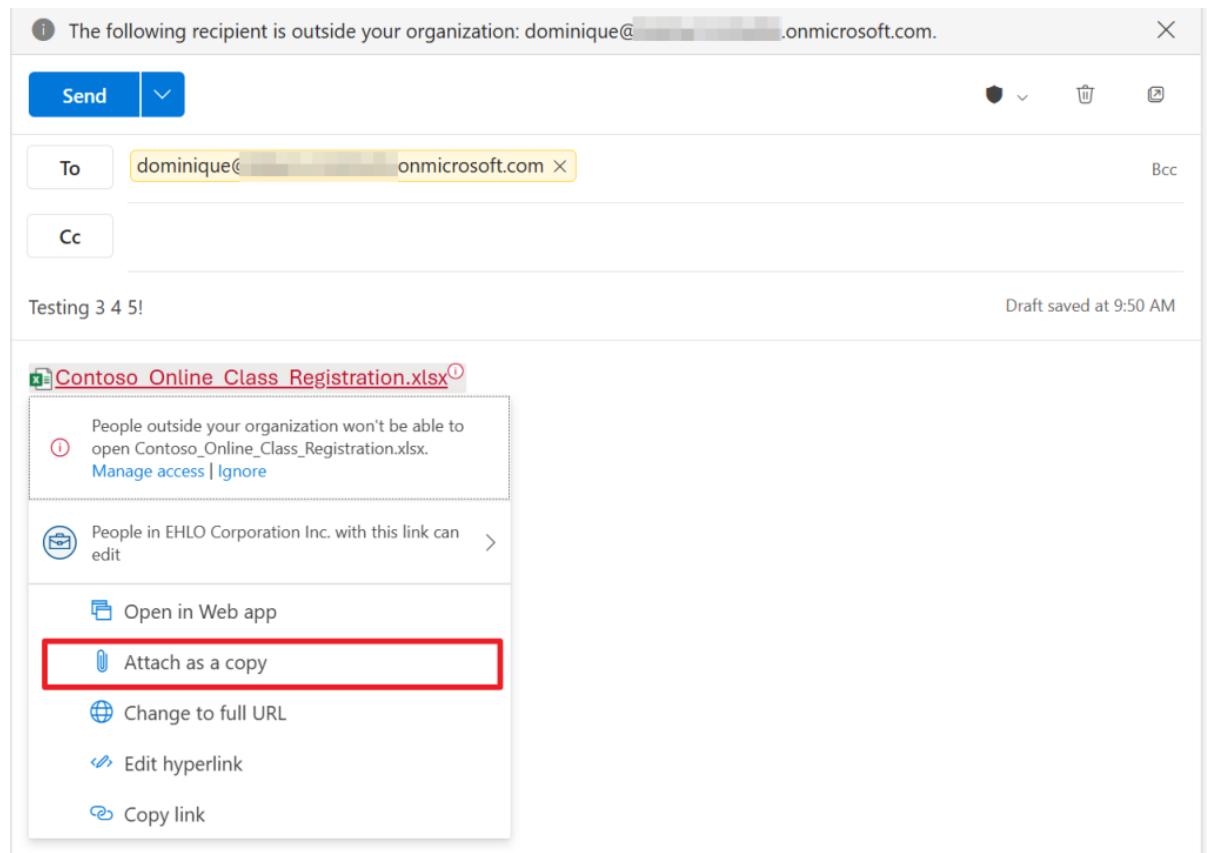
When sending the message anyway, another message pops up telling you that the recipients can't access the links in the document, after which you can send the message anyway (with an inaccessible link in it) or to don't send the message.



When the recipient receives the message, and tries to access the document:



This won't work because the tenant from which we shared the document did not create a guest account for the guest user. The simple reason for not creating the guest user is that the document could not be shared because of our DLP policy!



The following recipient is outside your organization: dominique@... .onmicrosoft.com.

Send **Send to** **Attachment** **File** **Help**

To: dominique@... .onmicrosoft.com **Bcc:**

Cc:

Testing 3 4 5! **Draft saved at 9:50 AM**

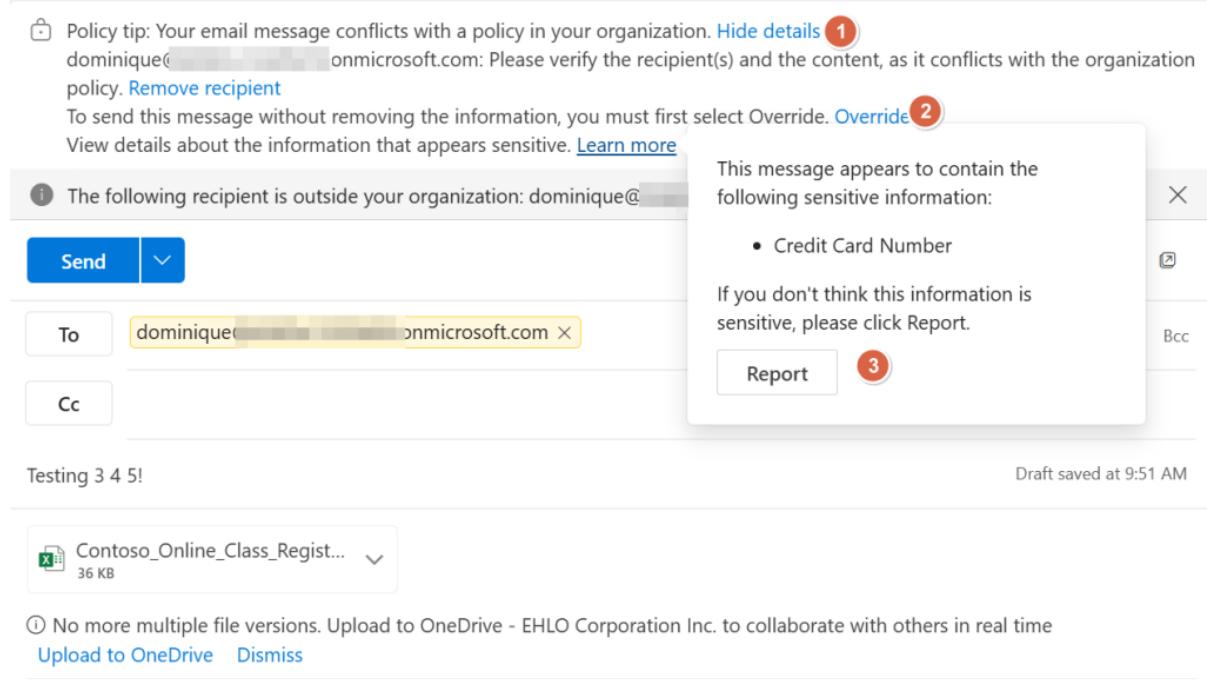
Contoso_Online_Class_Registration.xlsx ⓘ

ⓘ People outside your organization won't be able to open Contoso_Online_Class_Registration.xlsx. [Manage access](#) | [Ignore](#)

ⓘ People in EHLO Corporation Inc. with this link can edit [View](#)

[Open in Web app](#) [Attach as a copy](#) **Attach as a copy** [Change to full URL](#) [Edit hyperlink](#) [Copy link](#)

Let's try something else. Let's attach the file as a copy. Would this work?



ⓘ Policy tip: Your email message conflicts with a policy in your organization. [Hide details](#) 1

dominique@... .onmicrosoft.com: Please verify the recipient(s) and the content, as it conflicts with the organization policy. [Remove recipient](#)

To send this message without removing the information, you must first select Override. [Override](#) 2

View details about the information that appears sensitive. [Learn more](#)

ⓘ The following recipient is outside your organization: dominique@...

Send **Send to** **Attachment** **File** **Help**

To: dominique@... .onmicrosoft.com **Bcc:**

Cc:

Testing 3 4 5! **Draft saved at 9:51 AM**

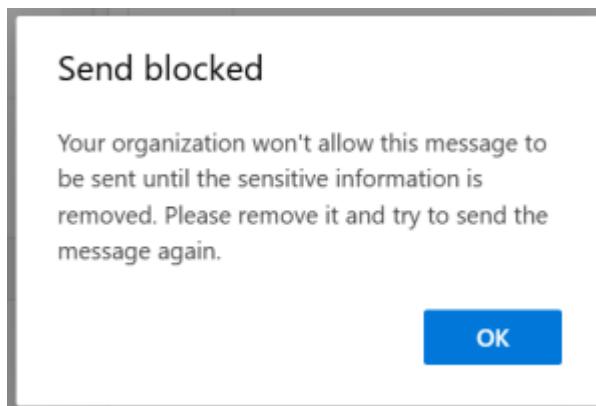
Contoso_Online_Class_Registration.xlsx ⓘ 36 KB

ⓘ No more multiple file versions. Upload to OneDrive - EHLO Corporation Inc. to collaborate with others in real time [Upload to OneDrive](#) [Dismiss](#)

As you can see, we are greeted with a couple of warnings:

1. A policy tip is shown that tells us that the email message conflicts with a (DLP) policy in our organization.

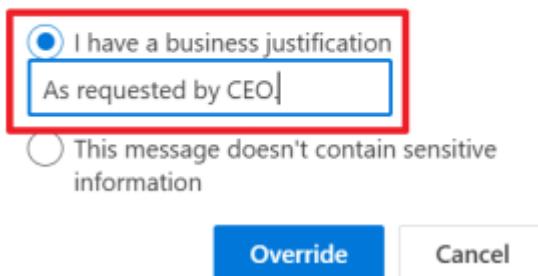
2. To send the message without removing the information (attachment in this case) we have to first select the 'override' button.
3. When selecting 'learn more', we can see that the document we attached contains Credit Card Numbers, which of course is found using the SIT we configured in our policy. We here have the option to report a false positive.



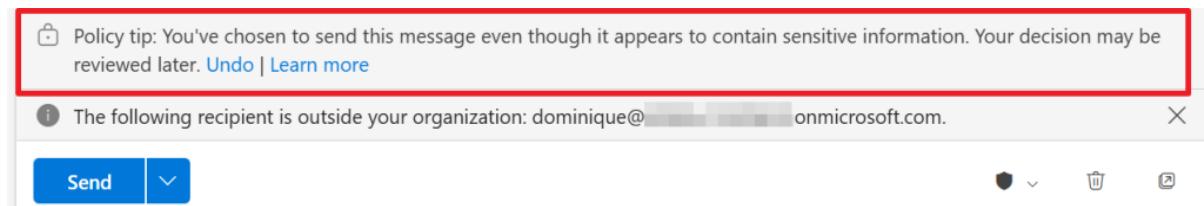
When trying to send the message without overriding, we see a message stating that we still won't be able to send the message without removing the attachment that contains sensitive information.

Explain why you want to send this message.

Your explanation may be subject to review later.



Now let's select 'override'. Now a business justification has to be entered or we have to select that we think the message doesn't contain sensitive information. Let's go with the first option and enter 'As requested by CEO'.



Another policy tip appears stating that your decision may be reviewed later as you've chosen to send this message even though it appears to contain sensitive information.

	A	B	C	D	E	F	G	H	I
1	Order Number	Order Total	Order Date	Order Time	Class Registration	Customer	Zip	Credit Card	Security Type
2	976890907	\$ 99.00	6/1/2015	5:46:31 PM	Yes	Berg, Karen	11716	3734-6623-7621-570	CID
3	976890831	\$ 198.00	6/1/2015	7:15:40 PM	Yes	Tolleven, Bjorn	32102	4539-6333-7119-3889	CVV2
4	976890396	\$ 149.00	6/2/2015	4:17:43 PM	Yes	Ceckova, Lucia	52643	6441-2087-8163-6199	CID
5	976890436	\$ 99.00	6/4/2015	1:32:48 PM	Yes	Saraiva, Jose	72441	4539-5907-4789-7710	CVV2
6	976890764	\$ 49.00	6/5/2015	1:11:51 PM	Yes	Balasubramanian, Dev	30447	4556-9525-3115-1167	CVV2
7	976890381	\$ 99.00	6/5/2015	4:18:17 PM	Yes	Andrews, Ben	34178	3447-9311-2642-660	CID
8	976890079	\$ 198.00	6/6/2015	11:40:49 AM	Yes	Halmaert, Franck	17617	3457-5582-5844-783	CID
9	976890423	\$ 149.00	6/6/2015	5:57:13 PM	Yes	Jain, Vikas	58234	5514-8818-2103-2514	CVC2
10	976890008	\$ 99.00	6/7/2015	7:44:13 AM	Yes	Jensen, Thomas Sonderg	77860	5493-1009-9211-6021	CVC2
11	976890868	\$ 49.00	6/7/2015	2:11:44 PM	Yes	Zighetti, Barbara	22808	4024-0071-9447-9662	CVV2

When the user receives the message, he/she is able to open the document containing the sensitive information. Let's understand why:

1. The user has overridden warnings from DLP stating that he/she is sending sensitive information to a party that should not have this information.
2. A business justification was added, that can be reviewed by a legal department of your organization for example.
3. Microsoft DLP does not encrypt files, so when an attachment is added a recipient can open the file as it's not encrypted. You could use encryption to counter this by using [sensitivity labels](#). These sensitivity labels can be teamed up with DLP for optimal protection!

In conclusion

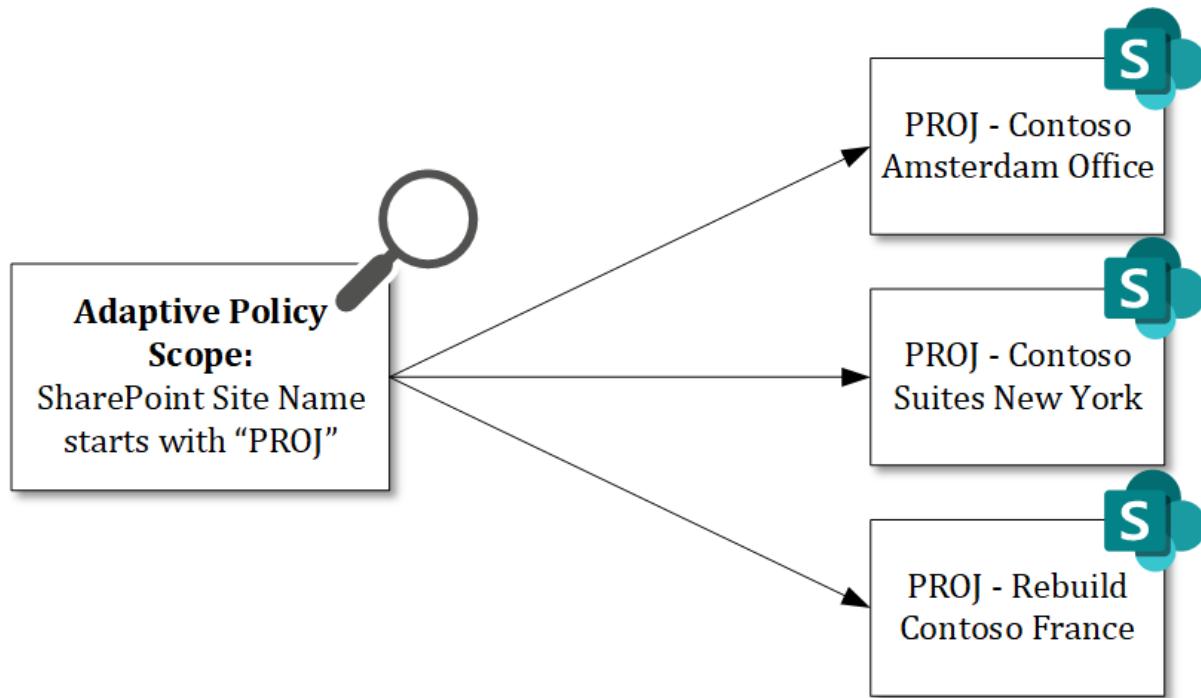
As we have seen in this chapter:

- Microsoft Data Loss Prevention (DLP) can be used to prevent users from oversharing information.
- It does this by applying actions to items that are considered to contain sensitive information.
- Sensitive information is found using Sensitive Information Types, also called SIT's.
- Policies can be ran in simulation mode first before enabling the policy in your environment.

Adaptive Scopes

In the chapter ‘Retention Policies (RM)’ I talked about retention policies and how you can apply them to your environment. To keep things simple, I used static scopes in that chapter. But since adaptive scopes are the recommended approach for retention policies I want to show you what they can do, and why Microsoft recommends them to be used instead of static scopes.

An adaptive scope is a scope that is dynamically filled based on a query. With that in the back of our mind, let’s start with the following scheme:



In this example, I’ve created 3 SharePoint sites. There’s also a naming convention that states that all project teamsites names should start with “PROJ”. Now we are going to create an adaptive policy scope that should find all of our sites.

The screenshot shows the Microsoft Purview interface with the 'Adaptive scopes' section selected. The 'Create scope' button is highlighted with a red box. The table below shows no data available.

	Name	Type	Created
No data available			

As adaptive scopes can be used in retention policies and communication compliance policies, they can be found under “Roles & Scopes”, “Adaptive Scopes” in the Microsoft Purview portal. Click “Create Scope” to create a new adaptive scope.

Name your adaptive policy scope

Name *

All Project SharePoint Sites

Description

This adaptive policy scope will contain all SharePoint sites starting with the name "PROJ".

Now, give the adaptive policy scope a logical name and description and click “Next”. In the following screen admin units can be selected. However, to keep things simple we’re going with “Full Directory”. If you would like to know more about admin units, take a look at the chapter about Data Lifecycle Management.

What type of scope do you want to create?

Each type of scope uses different attributes or properties to match the users, sites, or groups you want to detect in a policy.

Users

You'll select Microsoft Entra ID attributes used to define users (such as First name, Last name, and Department).

SharePoint sites

You'll select SharePoint properties used to define sites (such as site name, site URL, and refinable strings).

Microsoft 365 Groups

You'll select Microsoft Entra ID attributes used to define groups (such as Name, Description, and Email address).

As an adaptive policy scope “groups” objects together based on a query, we have to specify which objects we want to group. This can be users, SharePoint sites or Microsoft 365 groups. If you want to know what type of attributes or properties you can query, take a look at this [Learn article](https://learn.microsoft.com/en-us/purview/purview-adaptive-scopes#configure-adaptive-scopes) (<https://learn.microsoft.com/en-us/purview/purview-adaptive-scopes#configure-adaptive-scopes>). In this example, I choose “SharePoint sites”.

Create adaptive scope

The screenshot shows the 'Create adaptive scope' wizard. On the left, a navigation pane lists steps: Name, Admin unit, Scope type (selected), Site query, and Review. The main area is titled 'Create the query to define SharePoint sites'. It explains that the query consists of one or more SharePoint property/value combinations. A query builder interface is shown, with a 'Site properties' section containing a 'Site name' field with the operator 'starts with' and the value 'PROJ'. Below the query builder, a 'Query summary' box shows the resulting query: 'SiteTitle starts with PROJ;'.

Now for the query page. This is where the real magic happens. Here it is possible to enter the query or queries that contain the logic on which objects are selected to be part of the adaptive scope. As explained in the drawing above, we want to select all SharePoint sites that have a site name that starts with “PROJ”. Review the summary and finish the wizard.

Create adaptive scope

Your scope was created

It can take up to three days for the queries to fully populate, so the changes aren't immediate. Factor in this delay by waiting a few days before you add a newly-created scope to a policy.

Now that you've defined the scope, add it to a policy. If the attributes or properties that you specified change, the scope will automatically update to match them.

Now take note of the message that appears. It can take up to 3 days for the queries to fully populate, so keep that in mind when using your adaptive scope in a policy. When checking back a couple of days later, the following procedure can be followed to see whether our adaptive scope has queried our environment successfully:

Adaptive scopes > All Project SharePoint Sites

Scope

All Project SharePoint Sites

5 items Refresh

Filters: Filters

Display name	Location type	State
<input type="checkbox"/> https://[REDACTED].sharepoint.com/sites/M...	Site	Added
<input type="checkbox"/> https://[REDACTED].sharepoint.com/sites/PR...	Site	Added
<input type="checkbox"/> https://[REDACTED].sharepoint.com/sites/PR...	Site	Added
<input type="checkbox"/> https://[REDACTED].sharepoint.com/sites/PR...	Site	Added
<input type="checkbox"/> https://[REDACTED].sharepoint.com/sites/Pr...	Site	Added

Details

Name
All Project SharePoint Sites

Description
This adaptive policy scope will contain all SharePoint sites starting with the name "PROJ".

Admin unit
Full directory

Type
Site

Query summary
SiteTitle starts with PROJ

Last modified by
MOD Administrator

Last modified
Mar 26, 2024 8:13 PM

1. Navigate to “Roles & Scopes”, “Adaptive Scopes” in the Microsoft Purview portal
2. Click on the created adaptive scope.
3. Select “Scope Details” on the bottom of the screen.
4. You should see all objects that are now a “member” of your adaptive scope.

And now for the cool part, adaptive scopes are re-evaluated on a daily basis, so if you should create more SharePoint sites for projects like in this example (starting with “PROJ”), they are added to your retention policy within a day. The nice thing about this is that there's no manual action needed to add the new SharePoint site to your retention policy. This adds to the reliability of your retention policies. This is exactly why it's recommended to use adaptive scopes instead of static scopes. In the case of static scopes, the retention policy would have to be edited every time a new SharePoint site was added, for example.

Of course, this was just 1 example. With the flexibility that adaptive scopes provide, numerous options are possible!

Now you can use your adaptive scope in a retention policy. You can combine this chapter with information learned in the Data Lifecycle Management (DLM) chapter to create retention policies based on an adaptive scope!

eDiscovery (Premium)

Imagine you're a guy in a basement searching through millions and millions of boxes with documents (OK, that sentence may be a little exaggerated). You're responsible for processing requests for information for a government. These requests occur worldwide and are often enshrined in legislation. Examples include:

- The Netherlands: Wet Open Overheid (WOO)
- United States of America: Freedom of Information Act (FOIA)
- Canada: Access to Information Act
- United Kingdom: Freedom of Information Act

These requests often need to be handled within a specified timeframe. Not exactly doable if you have to comb through all these boxes like the guy in the picture. Good thing in reality you are not really that guy in the picture! You use modern methods and services like Microsoft 365 to store your information. And that's where the magic comes in that I will explain in this chapter. With Microsoft 365 you can use the eDiscovery feature to fulfill the requests for information. A really handy and quick tool that takes a lot of the manual work out of your hands. Let's see how this works by putting a request for information through its paces in eDiscovery Premium.

A quick note on licenses

Content Search	eDiscovery (Standard)	eDiscovery (Premium)
<ul style="list-style-type: none"> - Search for content - Keyword queries and search conditions - Export search results - Role-based permissions 	<ul style="list-style-type: none"> - Search and export - Case management - Legal hold 	<ul style="list-style-type: none"> - Custodian management - Legal hold notifications - Advanced indexing - Review set filtering - Tagging - Analytics - Predictive coding models And more...

Source: [Microsoft](#)

eDiscovery within Purview comes in 2 flavours: Standard and Premium. You can see the differences in the table above. There's also Content Search which can be used for the same purposes but has less features in terms of case management. In this chapter I will use eDiscovery Premium to demonstrate what features it has in store.

eDiscovery Roles

Role	Compliance Administrator	eDiscovery Manager & Administrator	Organization Management	Reviewer
Case Management	✓	✓	✓	
Communication		✓		
Compliance Search	✓	✓	✓	
Custodian		✓		
Export		✓		
Hold	✓	✓	✓	
Manage review set tags		✓		
Preview		✓		
Review		✓		✓
RMS Decrypt		✓		
Search And Purge			✓	

Source: [Microsoft](#)

To get your work done in eDiscovery you have to be assigned roles. The easiest way to get this done is to add the account that you will be using for eDiscovery to one of the available role groups: eDiscovery Manager, eDiscovery Administrator, Compliance administrator, Organization Management or Reviewer. As you can see in the table above, the eDiscovery manager and administrator roles give you the most extensive permissions available.

The screenshot shows the 'eDiscovery Manager' section of the Purview portal. On the left, a list of role groups is displayed, with 'eDiscovery Manager' selected and checked. On the right, details for the 'eDiscovery Manager' role group are shown, including its name, description, and assigned members. A red box highlights the 'eDiscovery Manager' role group details section.

Display name	Type	Admin units
There's no assigned member.		

Display name	Type	Admin units
There's no assigned member.		

Both are present in the roles & scopes -> permissions section of Purview. However, they're both available as a subgroup in parent role group 'eDiscovery Manager', as can be seen in the example above. Now, when navigating to eDiscovery, eDiscovery premium you can see the permissions that are assigned to you on the right hand of the screen.

eDiscovery Workflow

Advanced eDiscovery workflow

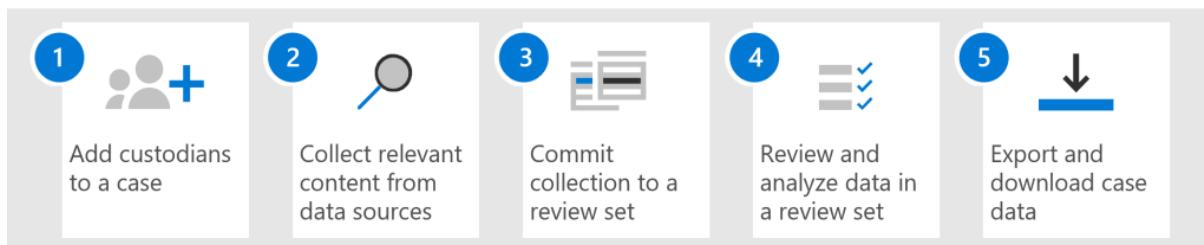


Image source: [Microsoft](#)

eDiscovery works with a workflow that helps you find the exact information that you need. Using this workflow, you will narrow down the volume of your data with each step, while making the results that remain at the end of the workflow as relevant as possible. Now that you are aware of the workflow we can dive in and take a look at what each step of the process looks like!

1. Create a case

Let's start by creating a case. Navigate to the Purview portal, eDiscovery, eDiscovery premium and select the 'cases' tab. Click 'Create a case'.

Advanced eDiscovery > New case

Name and description

Members and settings

Summary

Name your case

Provide basic information about this case and choose a case format

Name *

AV0001 - WOO Request Adele Vance

Description

WOO Request for Adele Vance, Customer Care reference number AV0001.

Number

AV0001

Case format

New (recommended)

- Improved performance and more durable data pipeline
- Increased collection, review set, and export limits
- Teams conversations collected as HTML transcript files (no support for .msg and .pdf formats)

[Learn more about the new case format](#)

Next **Cancel**

In the first screen, we have to create a name, description and Number for the case. The name and description speak for itself. The number can be any number that you can use as a reference to your customer management system for example. The new case format gives you more performance when having cases with a large number of items or a large size. Today, the old case format can not be used anymore. If you would like to have more information on the benefits of the new format, take a look at [this Microsoft learn page](https://learn.microsoft.com/en-us/purview/ediscovery-new-case-format) (<https://learn.microsoft.com/en-us/purview/ediscovery-new-case-format>).

Add team members and configure settings

Team members

Users
 X

Groups

Search and analytics

Reduce duplicates and link email threads

Similarity threshold (%) ^ v

Minimum word count ^ v

Maximum word count ^ v

Group items by theme

Create a saved query whenever analytics is performed

Text to ignore ()

Actual text or regular expression *

Apply to
 ▼

Case sensitive

+ Add

Optical character recognition (OCR)

Find text in images during advanced indexing

In the next screen, add your team members. These team members will investigate this case. Be sure to grant this members the correct eDiscovery permissions, as without them, they cannot access the correct parts of eDiscovery. Team members can be added as individuals or as groups.

We can also configure the search and analytics features here like reducing duplicates in the information that is found or linking email threads. The parameters for these features can be altered here, or when the case is created. Also, items can be grouped by theme and your search query can be saved.

If you have certain items with text you want to ignore, you can add them here. This actual text or regex can be applied to near-duplicates, email threads or themes and can be configured as case sensitive if you prefer.

Lastly, you can enable optical character recognition to find the text you specify in images.

Note that the actual search query is not defined in this screens. For this, we will configure a collection later.

2. Add Custodian(s)

eDiscovery (Premium) > Cases > AV0001 - WOO Request Adele Vance

Overview Data sources Collections Review sets Communications Hold Processing Exports Jobs Settings

Start your case by setting up people as custodians to quickly identify and preserve data sources with which they are associated. [Learn more](#)

Add data source Refresh

Add new custodians

0 items Search Filter Group Customize columns

Source type	Locations	Source status	Indexing status	Hold status	Index date

Add new custodians
Import custodians
Data Locations
Add data locations



Now that our case is created we have to add data sources. Data sources can be of one of two types:

- Custodial data sources: a custodian is someone who has access to a piece of information that might be relevant to this case. This piece of information can be an email, document, Teams group, SharePoint site or other type of information.
- Data Locations: A data location is a location that has data that isn't tied to a person but can be relevant to your case. This is also called a non-custodial data source.

Both data types above are indexed by a process called advanced indexing. This process optimizes searching for it in the next step.

New custodian

Identify custodian

Hold settings

Review

Select custodian

Identify new custodians from your organization's active directory

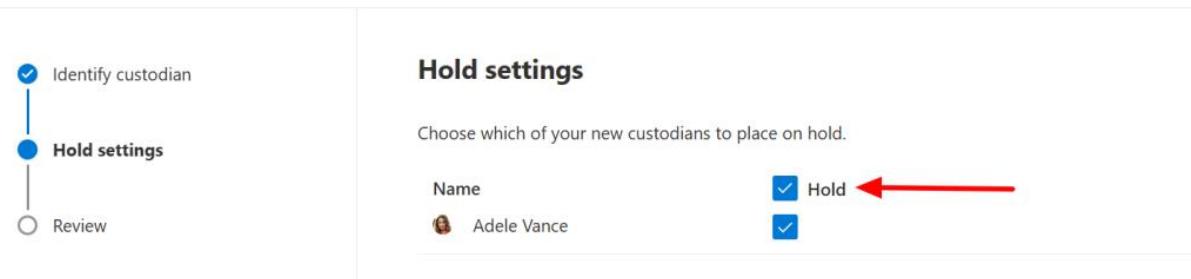
Adele Vance Please type minimum 3 characters to get the mailbox list.

Expand each custodian to view and add locations.

Custodian	Count	Clear	Edit
Adele Vance	5		
Mailboxes	1/1 (Default)	Clear	Edit
OneDrives	1/1 (Default)	Clear	Edit
Exchange	0		Edit
SharePoint	5		Edit
Teams	4		Edit
Yammer	0		Edit

In this demonstration, I add Adele Vance as a custodian and map her mailbox, OneDrive and various SharePoint sites and Teams groups to be linked to her as a custodian.

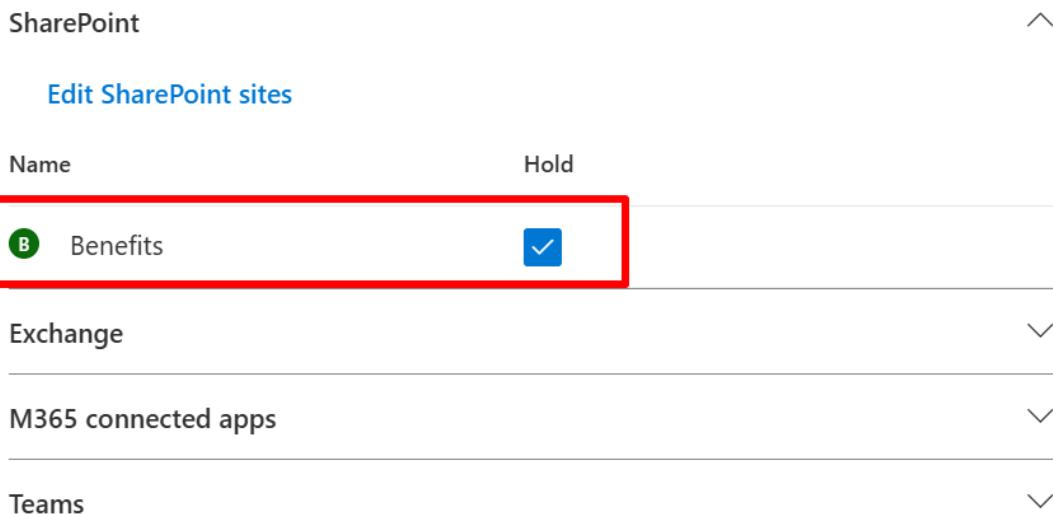
New custodian



The screenshot shows the 'New custodian' wizard with three steps: 'Identify custodian', 'Hold settings' (which is selected and highlighted with a blue circle), and 'Review'. The 'Hold settings' step is titled 'Hold settings' and contains the sub-instruction 'Choose which of your new custodians to place on hold'. It lists 'Name' and 'Adele Vance' with a 'Hold' checkbox. The 'Hold' checkbox for Adele Vance is checked and highlighted with a red arrow.

In the next step, I can choose to place this data on hold. This preserves the data and makes sure it can't be deleted by the user or anyone else because we need it in our case.

New non-custodial data locations



The screenshot shows the 'Edit SharePoint sites' page. It lists 'SharePoint' locations: 'Benefits' (selected and highlighted with a red box), 'Exchange', 'M365 connected apps', and 'Teams'. The 'Benefits' site has a 'Hold' checkbox checked, indicated by a red arrow.

For completeness I also added the 'Benefits' SharePoint site as a data source that isn't tied to a custodian and also place this location on hold.



The screenshot shows the 'Data sources' page. It lists data sources: 'Benefits' (selected and highlighted with a red box) and 'Adele Vance'. The 'Benefits' data source has a 'Hold' checkbox checked, indicated by a red arrow. The page includes a table with columns: 'Source name', 'Source type', 'Locations', 'Source status', 'Indexing status', 'Hold status', and 'Index date'. The 'Benefits' entry shows 'Data location' as the source type, 1 location, 'Active' status, 'Indexing' indexing status, 'Applying' hold status, and 'Sep 27, 2024 11:06 AM' index date. The 'Adele Vance' entry shows 'Custodian' as the source type, 10 locations, 'Active' status, 'Fully indexed' indexing status, 'Applying' hold status, and 'Sep 27, 2024 11:05 AM' index date.

When this is done the system goes through its paces to index the sources and apply the holds.

3. Create a collection

Now indexing is complete and holds have been applied, let's move to the 'collections' tab. The creation of a collection allows you to build search queries to fetch the information you need from the data sources you added to the case. As with most wizards, let's start by specifying a name and description for the collection. Next, we have to select the custodial and non-custodial data sources that we want to use in this collection. As this is just a simple case, we select each custodial and non-custodial data source that's created and include all services to be included. In a real-world scenario, you would use collections to narrow down your search result.

Additional locations

Choose additional locations to search. An additional location is a data source that isn't associated with the custodians you selected on the previous page. Note: these additional locations will not include advanced indexing unless they are added as data sources in the case, [learn more in this article](#).

ⓘ Location limit: Partially indexed items cannot be collected from additional locations. Additional locations don't include advanced indexing either.

Status	Location	Included	Excluded
Off	ⓘ Exchange mailboxes Microsoft 365 Groups Teams Yammer user messages		
Off	ⓘ SharePoint sites OneDrive sites Microsoft 365 Group Sites Team Sites Yammer Networks		
Off	ⓘ Exchange public folders		

Additional search options

ⓘ These options have been set as default by your eDiscovery admin. However, adjustments per collection are enabled. Modify the defaults to meet the requirements of your collection by selecting preferred options.

Locations

Search additional locations during tenant-wide searches. Note: including these additional locations can cause searches to take longer to complete.

<input type="checkbox"/> Guest mailboxes
<input type="checkbox"/> Shared Teams channels
<input type="checkbox"/> Include departed users (inactive mailboxes) in search scope
<input type="checkbox"/> Include group mailboxes in search scope

On the next page, you can specify additional locations to be included in your collection. It's good to be aware of the fact that the locations or search options selected here will not be indexed by the advanced index process!

Define your search query

Use the query builder or editor to define your search. [Learn more about queries](#)

Query language-country/region: None 

Use new query builder

Query builder

KQL editor

Filters  Clear all



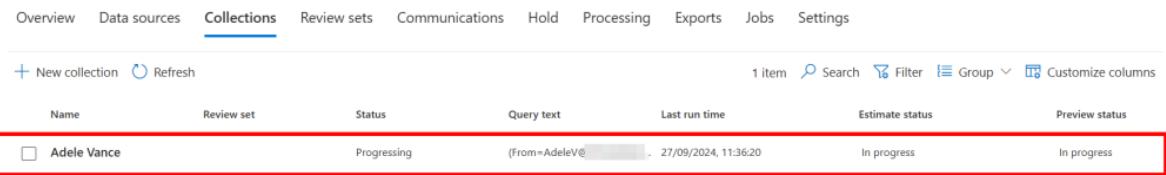
OR

Sender Equals any of AdeleV@...

Keywords Equal adele vance (cs) adele (cs...)

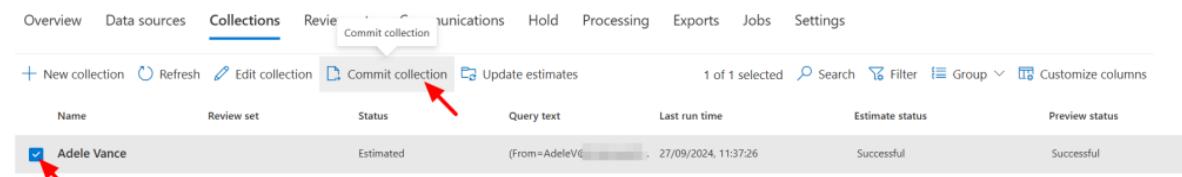
+ Add filter 

Now we arrive on the page where we can construct our query that actually will fetch the data we are searching for. You'll have the option to do this with the query builder or by using the KQL editor. I opted to go for a simple query I drafted up with the query builder. It will search for Adele as the sender OR adele vance, adele or vance as keywords in items. In a production environment, this is where serious thought goes into creating your query to be able to fetch the right information.



Name	Review set	Status	Query text	Last run time	Estimate status	Preview status
<input type="checkbox"/> Adele Vance		Progressing	(From=AdeleV@...	27/09/2024, 11:36:20	In progress	In progress

Once again, the system starts running to fetch the information you requested from the data sources that were specified earlier! When the search is complete you can click the collection and take a look at the statistics of the data found. Now you can edit your query to better match your needs or when you're happy with the end result, you can commit the collection to a review set by selecting the collection and pressing 'commit collection'.



Name	Review set	Status	Query text	Last run time	Estimate status	Preview status
<input checked="" type="checkbox"/> Adele Vance		Estimated	(From=AdeleV@...	27/09/2024, 11:37:26	Successful	Successful

When you commit a collection, the data that is found using the query is copied to a secure Azure Storage location. When the data is arrived at that destination, it is indexed again so you are provided with fast search results when looking at the items in your review set.

Commit items to a review setGather items and process them into a review set. [Learn more about committing items](#) Add to new review set

Review set name

AV0001-RS01

 Add to existing review set

Select review set

Retrieval

Identify additional items to collect.

 Teams and Yammer conversations
Collect up to 12 hours of related conversations when a message matches a search. Cloud attachments
Collect items from links to SharePoint or OneDrive. All document versions
Collect all versions of SharePoint documents. If not selected, only current versions are collected. Partially indexed items
Collect unsearchable items that might be relevant.**Collection ingestion scale** Add all of collection to review set Add only collection sample to review set. [Edit sample parameters](#)

When committing the items to a review set, you have to specify a name for the set and choose if you would like the retrieve additional items like Teams and Yammer messages, cloud attachments by following links to SharePoint or OneDrive or by including all document versions or partially indexed items.

Lastly, you can choose to add all of your collection to the review set or use a collection sample. This sample can consist of documents based on a confidence level or a random sample size which you can specify.

4. Review your review set

When your data is committed to the review set, you can select the 'review sets' tab to see it. Select your review set and select 'open review set'. Let's see what we can do with our review set.

Overview

eDiscovery (Premium) > Cases > AV0001 - WOO Request Adele Vance > AV0001-RS01

Saved filter queries ▾

Filters

AND

75 items

<input type="checkbox"/>	#	Subject/Title	Status	Tag Status	Date (UTC +02:00)	Sender/Aut
<input type="checkbox"/>	1	You have late tasks	<input type="radio"/> Ready	No Tag	Aug 31, 2024 2:24:...	Microsoft o...
<input type="checkbox"/>	2	You have late tasks	<input type="radio"/> Ready	No Tag	Sep 1, 2024 5:28:10...	Microsoft o...
<input type="checkbox"/>	3	You have late tasks	<input type="radio"/> Ready	No Tag	Aug 20, 2024 8:38:...	Microsoft o...
<input type="checkbox"/>	4	You have late tasks	<input type="radio"/> Ready	No Tag	Jul 2, 2024 9:13:47 ...	Microsoft o...
<input type="checkbox"/>	5	You have late tasks	<input type="radio"/> Ready	No Tag	Sep 2, 2024 6:26:42...	Microsoft o...
<input type="checkbox"/>	6	V6E4G2	<input type="radio"/> Ready	No Tag	Jun 24, 2024 1:16:0...	MOD Adm...

Subject line

3

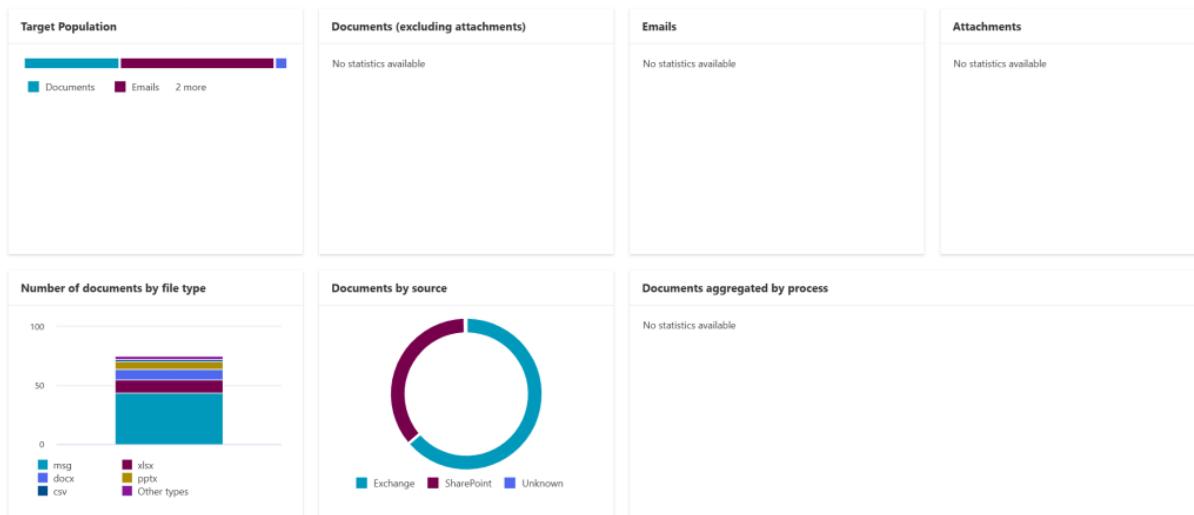
Select an item from the list to preview its content

In the overview pane, we can set (and save) filters to further narrow down the list of documents in this review set (1). We can also see the list of all files in this review set (2) and even take a look at the document content in the preview pane (3).

The first button provides you with an overview of the current review set, which includes the total extracted documents, total processed items and total number of failed extracted documents. It also shows you how many load sets are in this review set. A load set is a given number of documents that are added to a review set at a single time.

Analytics

Analytics



The second button allows you to run document and email analytics which in turn provide you with some nice analytics as can be seen in the image above.

Document Review

The figure is a screenshot of the Document Review interface. It shows a preview of a document titled "Contoso Electronics to Open Miami Store.docx". The preview area includes:

- Source:** Plain text, Annotate, Metadata.
- Show pinned metadata:** A checkbox with the label "Show pinned metadata".
- Word:** A preview of the Microsoft Word document content.
- Image:** A preview of an image showing a keyboard.
- Text:** The document text, including a section titled "FOR IMMEDIATE RELEASE" and a paragraph about Contoso's plans to open a store in Miami.
- Navigation:** A set of navigation icons for navigating through the document.
- Page Number:** "Page 1 of 2" at the bottom left.
- Page Control:** A blue button labeled "Tag" at the bottom left.
- Header:** "Contoso Electronics to Open Miami Store".
- Footer:** Accessibility Mode, Print, Find, Immersive Reader, and a "Give Feedback to Microsoft" link.

Now when we select a document there's a couple of things we can do that help us in the review process. Let's start with the button on the lower left. This provides us with a tagging option to provide the document with a tag. This tag can be part of a tag group that you can use within other documents as well and you can use them in search queries. We can also add notes, download the original file or pdf. It's even possible to take a look at the metadata or plain text version of the files!

Take a look at the video above to see the annotation possibilities in action, really cool stuff if you ask me. Select actions, commit redactions to pdf to save them when you are done. You can also use the same menu to add documents to another review set. Another powerful tool in this menu is the ability to group documents by conversations/related items or by families. If you want to know more about this, see the following [Microsoft Learn article](https://learn.microsoft.com/en-us/purview/ediscovery-view-documents-in-review-set#grouping) (<https://learn.microsoft.com/en-us/purview/ediscovery-view-documents-in-review-set#grouping>).

Number	File Type	Subject Line	Status	Tag Stat
63	Word	Contoso Electronics to Open Miami S...	Ready	No Tag
64	Excel	Q4 Marketing Analysis.xlsx	Ready	No Tag
65	Word	International Marketing Campaigns.d...	Ready	No Tag
66	Word	Proposed_agenda_topics.docx	Ready	No Tag
67	Excel	Inventory.xlsx	Ready	No Tag
68	Word	Home	Ready	No Tag

When you are done with reviewing all the files, there are several ways to export the information for further review by third parties. You can do this by downloading the documents in source or pdf (containing annotations) format, but it's also possible to export a selection or all documents in bulk. This contains the metadata, native files, text files, and redacted documents that have been saved to a PDF file. You can do this by selecting the files and pressing 'export'.

Export options

Export name *

AV0001 - Export 01

Description

Description for the export

Export these documents *

- Selected documents only
- All filtered documents
- All documents in the review set

Expand selection

None

**Output options ***

- Report only
- Loose files and PSTs (email is added to PSTs when possible)
- Condensed directory structure
- Condensed directory structure exported to your Azure Storage account

Container URL

Container URL

SAS token

SAS token for the Azure Storage account

Include

- Tags
- Text files
- Replace redacted natives with converted PDFs

Export**Cancel**

To conclude this chapter, let's take a look at the options we have here. We can provide a name and description for the export and select which documents we want. (Selected, all filtered, all documents in the review set). This selection can be expanded by using family groups or conversation groups if you have these created.

The output options at our disposal are really nice to, we can only export a report, loose files and PST's or a condensed directory structure which you can download directly to your pc or to your Azure Storage Account. I also opted to include Tags, text files and to replace redacted natives with converted PDF's (not shown in image).

The screenshot shows the Microsoft 365 Compliance Center interface. At the top, there is a file list with the following entries:

Name	Date modified	Type	Size
AV0001__Export_01_1of1	27/09/2024 13:45	Compressed (zipped)...	2,092 KB
Summary	27/09/2024 13:44	Microsoft Excel Com...	1 KB

Below this, a detailed view of the export structure is shown:

- Today** (selected):
 - Export_Loadfile_1of1 (Microsoft Excel Com... 15 KB)
 - Warnings_Errors_1of1 (Microsoft Excel Com... 1 KB)
 - NativeFiles (File folder)
 - Extracted_text_files (File folder)

Red arrows highlight the 'Today' folder and its sub-items. A larger red arrow points from the 'Today' folder down to the detailed view of the extracted text files.

On the right, a preview pane shows a document page with the following text:

FOR IMMEDIATE RELEASE

Redacted - Contoso announced today that it plans to open a Miami, Florida retail store early next year. The company's recent product line has seen a dramatic rise in both online and retail store sales, with a posted quarterly revenue of \$28.27 million in the just concluded quarter. **Redacted** or diluted earnings per share for the quarter came to revenue of \$21.65 million and net quarterly profit of \$4.45 million, or **Redacted** per diluted share, in the year-ago quarter. Gross margin was 42.3 percent compared to **Redacted** in the year-ago quarter. International sales accounted for 27 percent of the quarter's revenue.

"We are thrilled to bring our products to the people of Miami. With our very strong

Redacted text is circled in the preview pane.

At the bottom, another file list shows the following entries:

Name	Date modified	Type	Size
a9c534d639f1fc604a924405da7bf182fee15...	27/09/2024 13:51	Text Document	13 KB
a42be22258c6403fc6375935fd7cbb5130497f677...	27/09/2024 13:51	Text Document	2 KB
c3a478417621f7d0fd530a80a1e4c44d009e2c942d01...	27/09/2024 13:51	Text Document	1 KB
c38aaa66b7b24b68e0ac8944c212e553e121b4d5f50ce3...	27/09/2024 13:51	Text Document	1 KB
c232c26560db8088ab439b9d7e66e84b9f2139147f267f...	27/09/2024 13:51	Text Document	10 KB

The output is displayed above.

Communication Compliance

Communication Compliance in Microsoft Purview can detect messages in your organization that are considered to be inappropriate. Besides detection it can also capture and take action on the messages that it finds. Microsoft Purview is equipped with several out-of-the-box policies and gives you the possibility to create your own. Communication compliance policies can be used to check for inappropriate messages in internal and external communications that take place in email (Exchange), Meeting/IM (Teams chat, channel messages, meeting transcripts with recordings), Viva Engage and interactions with Microsoft 365 Copilot.

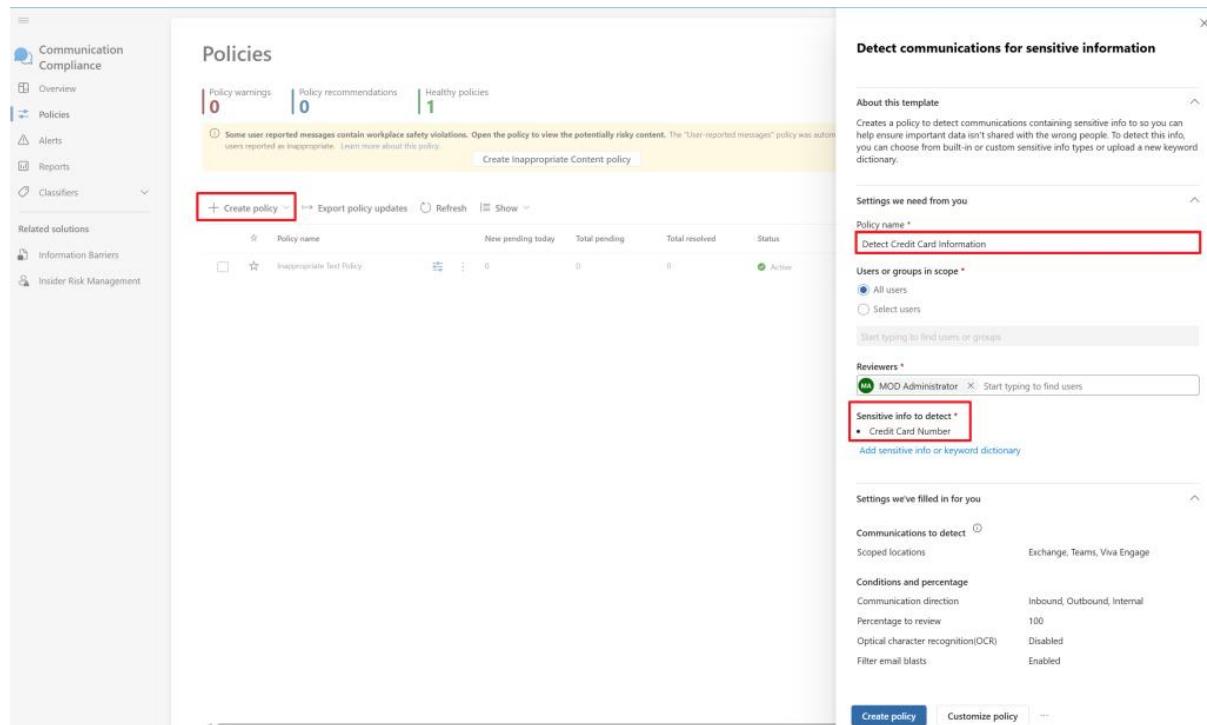
You can think of the following messages being inappropriate in your environment:

- Messages that contain sensitive content.
- Messages that contain inappropriate content, text or images.
- Messages that contain conflict of interest.
- Messages that contain information that is against laws or compliance policies.
- And so on!

In this chapter I want to show you how to create a communication compliance policy, what it looks like for the user that sends messages being inappropriate, how these messages are captured and how you can take action. Are you ready? Let's go!

Communication Compliance Policy Configuration

In this example we're going to set up a Communication Compliance policy that detects credit card numbers and we'll use Microsoft Teams to demonstrate what this looks like for the users involved. Before we begin though, make sure the audit log is enabled in your environment.



The screenshot shows the Microsoft Purview Communication Compliance interface. On the left, the navigation menu includes 'Communication Compliance', 'Overview', 'Policies' (which is selected and highlighted in red), 'Alerts', 'Reports', 'Classifiers', 'Related solutions' (with 'Information Barriers' and 'Insider Risk Management' listed), and 'Audit logs' (with 'Audit logs' and 'Audit logs history' sub-options). The main 'Policies' page displays a summary: 0 Policy warnings, 0 Policy recommendations, and 1 Healthy policies. A message states: 'Some user reported messages contain workplace safety violations. Open the policy to view the potentially risky content. The "User-reported messages" policy was automatically created for you.' Below this is a button to 'Create Inappropriate Content policy'. The table below shows one policy: 'Inappropriate Test Policy' (Status: Active). The right side of the screen shows the configuration pane for the 'Detect communications for sensitive information' template. It includes sections for 'About this template', 'Settings we need from you' (with 'Policy name' set to 'Detect Credit Card Information'), 'Users or groups in scope' (set to 'All users'), 'Reviewers' (set to 'MOD Administrator'), 'Sensitive info to detect' (set to 'Credit Card Number'), 'Settings we've filled in for you' (with 'Communications to detect' set to 'Inbound, Outbound, Internal'), and 'Conditions and percentage' (with 'Communication direction' set to 'Inbound, Outbound, Internal', 'Percentage to review' set to '100', 'Optical character recognition(OCR)' set to 'Disabled', and 'Filter email blasts' set to 'Enabled'). At the bottom are buttons for 'Create policy' and 'Customize policy'.

Now, let's start by navigating to the Microsoft Purview portal, select the Communication Compliance solution and navigate to the policies window. There, select 'Create policy' to create a new policy. Here we'll need to configure a few fields of information:

1. Give our policy a name. I'll go with 'Detect Credit Card Information'.
2. I select all users to be in scope of this policy, which means that this policy will apply to all users in my tenant.
3. The reviewer will be 'MOD Administrator'. As this is just for demo purposes, selecting the administrator is fine, however in a real world scenario I would advise you to delegate the task of reviewing policy matches to a dedicated resource.
4. For sensitive info to detect we'll select 'Credit Card Number'. Do note that you can select any Sensitive Info Type (SIT) or keyword dictionary to be detected in communications.
5. The scope of where this policy will be applied is Exchange, Teams and Viva Engage and this will be applied for inbound, outbound and internal communication.
6. We will leave 'percentage to review' at 100%, so we'll have to review all policy matches.
7. Optical Character Recognition is turned off at this screen, but is enabled later when selecting 'customize policy'. This feature can even detect forbidden communications in images that are being sent, really cool!
8. The 'Filter email blasts' feature is enabled, so messages from bulk email senders will not generate alerts.

When done, we'll click 'customize policy', so we have a chance to:

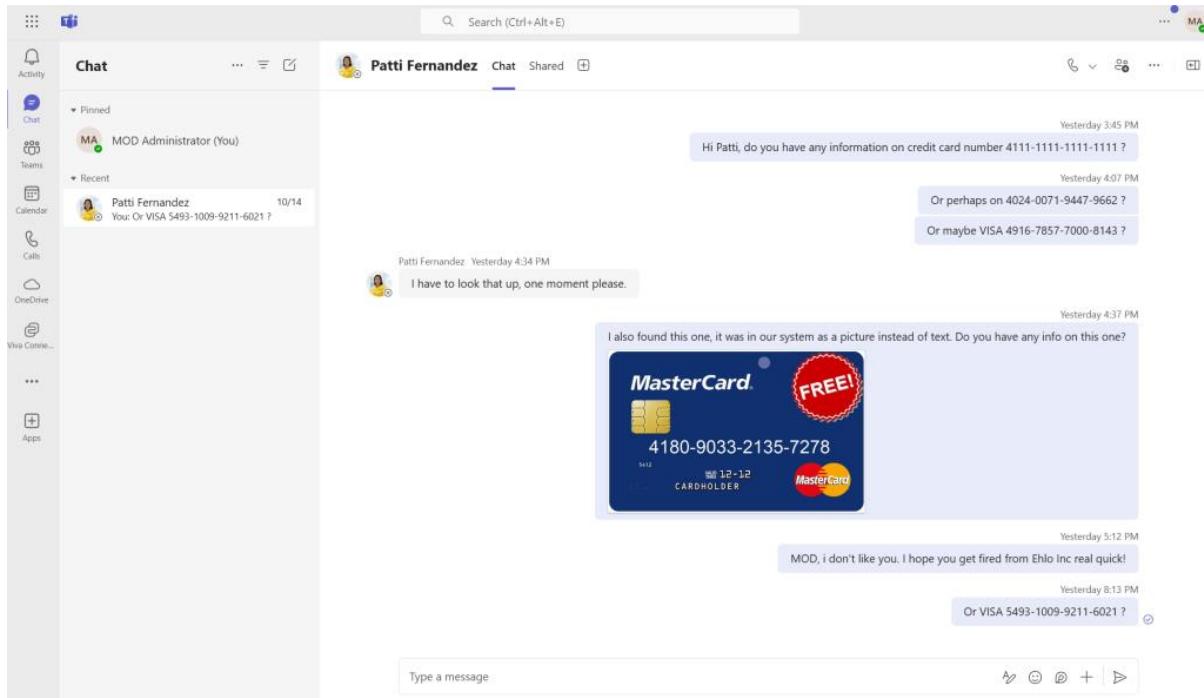
- Exclude users and groups from the policy.
- Select additional Microsoft 365 locations (like Microsoft 365 Copilot) or non-Microsoft apps.
- Edit the SIT.



Policy name	New pending today	Total pending	Total resolved	Status	Policy health	Last modified (UTC)
Detect Credit Card Information	0	0	0	Activating	Healthy	Oct 14, 2024 12:46 PM

When done you will see that the status of your new policy is now 'Activating'. After a while the policy will have a status of 'Active' and policy health of 'Healthy'. When this is the case, we can start our testing.

Testing, 1, 2, 3



To demonstrate what is going to happen, I set up a chat between MOD Administrator and Patti Fernandez as you can see in the image above. The Administrator has sent various messages that contain (fictional) credit card numbers and even sends 1 (fictional) credit card number as a picture.

Content type	Time to detection
Email body content	1 hour
Teams body content	1 hour
Viva Engage body content	1 hour
Viva Engage attachment	Up to 24 hours
Microsoft 365 Copilot and Microsoft Copilot body content (prompts and responses)	1 hour
Email OCR	24 hours
Teams OCR	24 hours
Email attachment	24 hours
Team attachment	24 hours
Teams modern attachment	24 hours
Teams metadata	1 hour
Email metadata	1 hour
Teams shared channels	24 hours
Teams transcripts	1 hour

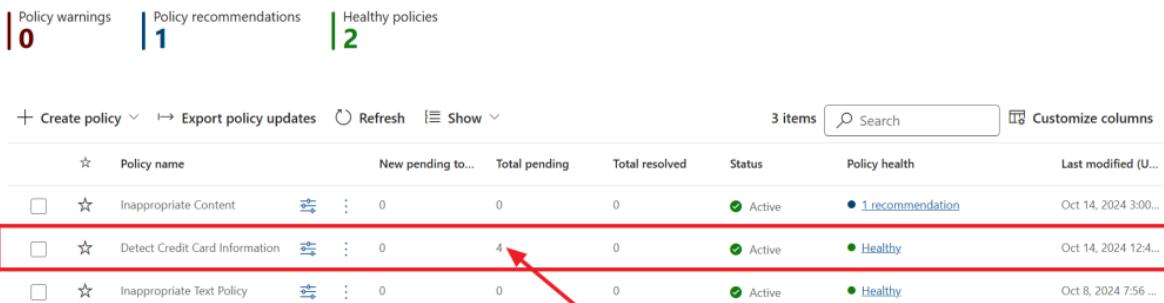
Image Source: [Microsoft](#)

Now, the Communication Compliance service starts detecting and processing these messages after which they pop up in the Communication Compliance policy screen. In the screenshot above you can see how long it takes for the service to pick up the various types of content.

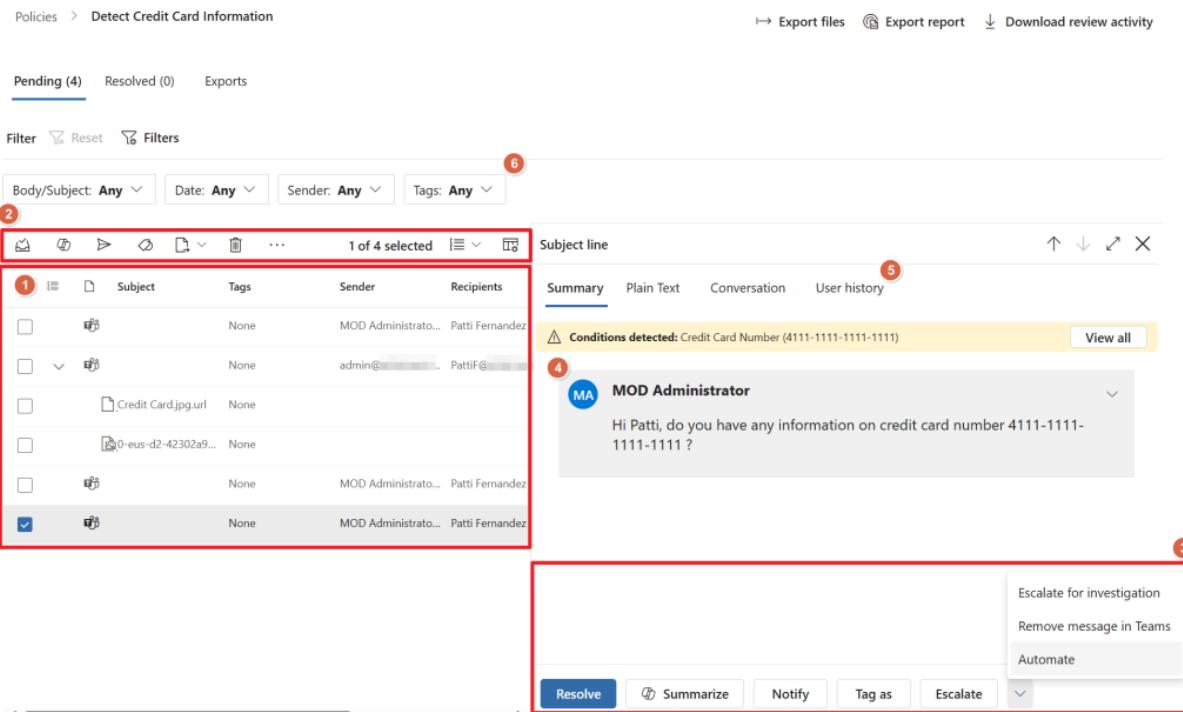
Communication Policy Match Overview

Policies

 Recommended actions  Learn  What?



Back at the Purview Communication Compliance policies screen, you can see that our ‘Detect Credit Card Information’ Policy has picked up 4 messages that match our policy.



When clicking the policy, we get some insights on what's exactly going on. We can (numbers below match numbers in the picture above):

1. See an overview of all messages that match with our communication policy and included SIT.
2. Apply various actions to one or multiple messages:
 - o Resolve: Move the message from the pending queue to the resolved queue. No further action is taken.
 - o Summarize: Let Copilot summarize a message (requires separate license)

- Notify: Notify the sender of the message using a notice template.
- Tag as: Tag the message with one of the predefined filters or with a custom one that can both be used to filter messages.
- Escalate: Let other people in your organization review the message.
- Escalate for investigation: Create an eDiscovery (premium) case for further review.
- Remove message in Teams: Removes the message for sender and recipient in Teams and shows a help tip why the message was blocked.
- Automate: Use a Power Automate action on the selected message.

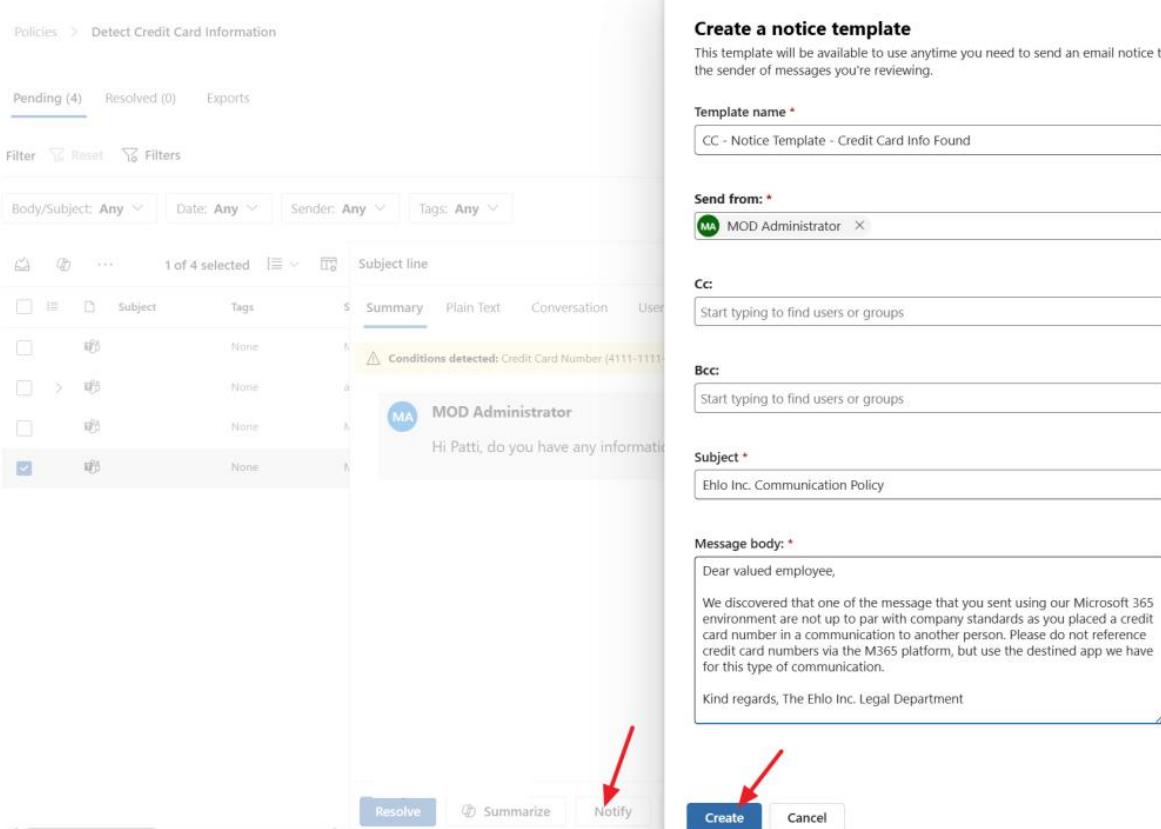
3. Same actions as for number 2, however in this case the scope is only the highlighted message.

4. Preview of the message selected.

5. Show the message summary, in plain text, a snippet of the conversation or the user history.

6. Filter the message list.

Notify the user



The screenshot shows the Microsoft 365 Compliance Center with the following details:

- Policies > Detect Credit Card Information**
- Pending (4) Resolved (0) Exports**
- Filter Reset Filters**
- Body/Subject: Any Date: Any Sender: Any Tags: Any**
- 1 of 4 selected**
- Subject line**
- Conditions detected: Credit Card Number (4111-1111-1111-1111)**
- MOD Administrator** (highlighted in blue)
- Hi Patti, do you have any information**
- Create a notice template** (right-hand pane)
- Template name ***: CC - Notice Template - Credit Card Info Found
- Send from: ***: MOD Administrator
- CC:** (input field: Start typing to find users or groups)
- Bcc:** (input field: Start typing to find users or groups)
- Subject ***: Ehlo Inc. Communication Policy
- Message body: *** (rich text area):

Dear valued employee,

We discovered that one of the message that you sent using our Microsoft 365 environment are not up to par with company standards as you placed a credit card number in a communication to another person. Please do not reference credit card numbers via the M365 platform, but use the destined app we have for this type of communication.

Kind regards, The Ehlo Inc. Legal Department
- Buttons at the bottom:** Resolve, Summarize, Notify, Create, Cancel. The 'Create' button is highlighted with a red arrow.

By selecting one or multiple messages and selecting 'Notify', we can notify the sender of a message. When no notice templates are configured yet, here you also have the opportunity to create one. I choose to create one that informs our user that a message he or she has sent isn't

up to par to the company's standards. After selecting your template you have the opportunity to change it to your liking.

MOD Administrator
To: MOD Administrator

Dear valued employee, We discovered that one of the message that you sent using our Microsoft 365 environment are not up to par with company standards as you placed a credit card number in a communication to another person. Please do not reference credit card numbers via the M365 platform, but use the destined app we have for this type of communication. Kind regards, The Ehlo Inc. Legal Department

Reply Forward

See above what the message looks like in the mailbox of the user. Unfortunately the original message isn't present so the user can't easily match the notice with a message he might have sent.

Message Removal

1 of 4 selected

Subject line

Summary Plain Text Conversation User history

Conditions detected: Credit Card Number (4916-7857-7000-8143) View all

MOD Administrator
Or maybe VISA 4916-7857-7000-8143 ?

Escalate for investigation
Remove message in Teams
Automate

Resolve Summarize Notify Tag as Escalate

When selecting 'remove message in Teams' the message is... removed in Teams! This is a quite powerful feature to dispose of messages that do not comply with company standards.

Patti Fernandez Chat Shared +

Yesterday 3:45 PM

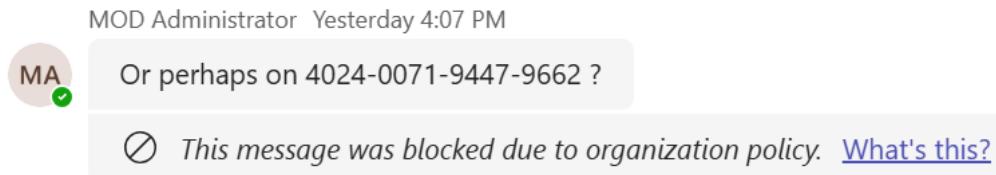
Hi Patti, do you have any information on credit card number 4111-1111-1111-1111 ?

Yesterday 4:07 PM

Or perhaps on 4024-0071-9447-9662 ?

This message was blocked. What can I do?
Or maybe VISA 4916-7857-7000-8143 ?

From the sender's point of view, the message is still displayed, but with an added note that the message is actually blocked. The message contains a help link to inform the user of the process that took place to actually block the message.



On the recipient side, the message is completely removed and is replaced by a message stating that the message was blocked due to organization policy. Ain't that nice?

Resolve a message

By pressing the resolve button, one or multiple messages are being moved from the pending queue to the resolved queue. This indicates that no further action has to take place on the message.

Resolve

Selected message will be moved to the "Resolved" tab. Add a comment to let others know why they were resolved.

- Because cross-policy resolution is turned on, selected messages will be automatically resolved in all other policies that detected them. You can turn this off from Communication Compliance settings.

Comment (optional)

False positive. This user was allowed to send this message.

A comment can be placed that will be stored with the message in the Communication Compliance Center. Notice that by default, the message will be resolved in all other policies that detected it. This can be turned off if it not to your liking.

Optical Character Recognition (OCR)

Conditions detected: Credit Card Number (4180-9033-2135-7278)

I also found this one, it was in our system as a picture instead of text. Do you have any info on this one?

This has to be one of the coolest parts of this feature. In the picture above a message is linked with an image. Take a look at the credit card number that the OCR feature found.

It's actually found in a picture! So your SIT's will also be detected when hidden in a picture! Also the picture can be removed from the chat in this case!

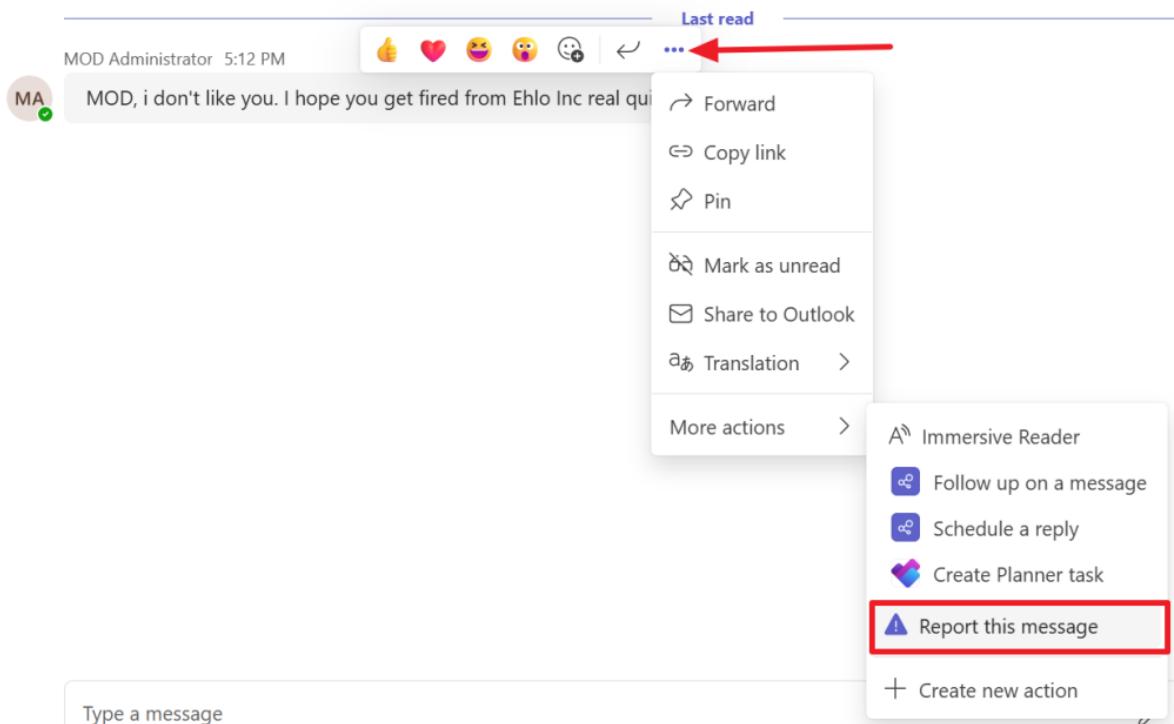
Bonus: Quick Setup: User-reported messages & Inappropriate Content

I hope the above explanation was a good example of the inner workings of communication compliance. However, when you don't want to go through the hassle of setting up policies yourself, Communication Compliance provides you with 2 out of the box options. The first one is 'user-reported messages'.

Policy name	New pending to...	Total pending	Total resolved	Status	Policy health	Last modified (U..)
Detect Credit Card Information	2	2	0	Active	Healthy	Oct 14, 2024 12:4...
Inappropriate Text Policy	0	0	0	Active	Healthy	Oct 8, 2024 7:56 ...

When you first enter the Communication Compliance console, a yellow bar appears stating 2 things:

1. The ‘user-reported messages’ policy was automatically created by Microsoft to detect Teams and Viva Engage messages that *users reported* as inappropriate.
2. The option to create an ‘Inappropriate Content Policy’.



The first option creates a policy that will kick in whenever a user reports a message in Teams or Viva Engage via the ‘Report this message’ button. After the button is clicked, the user can mark the message as a security risk or as inappropriate. The message is then shown in the Communication Compliance console for further investigation.

The second recommended policy (which can be created with the button in the yellow bar) creates an ‘Inappropriate Content Policy’ that monitors communications for content that matches one of the following trainable classifiers:

- Sexual
- Violence
- Hate
- Self-Harm

This is a pretty powerful feature that starts monitoring your messages on the topics above with little configuration and is great to start your Communication Compliance policy with!

Insider Risk Management (IRM)

Let's talk about the following scenario. You have an employee that is leaving your company. Because the leaving employee thinks that he has the rights to all of the companies documents he starts downloading them for later use and sends the documents to his private email account using Dropbox. Wouldn't you want to be notified of such a scenario?

Enter Microsoft Purview Insider Risk Management (IRM). A solution that collects information from all kinds of different sources like Microsoft 365 and perhaps other services like your HR-system. In IRM, you can create various policies that let you monitor all sorts of policy violations. A few examples are:

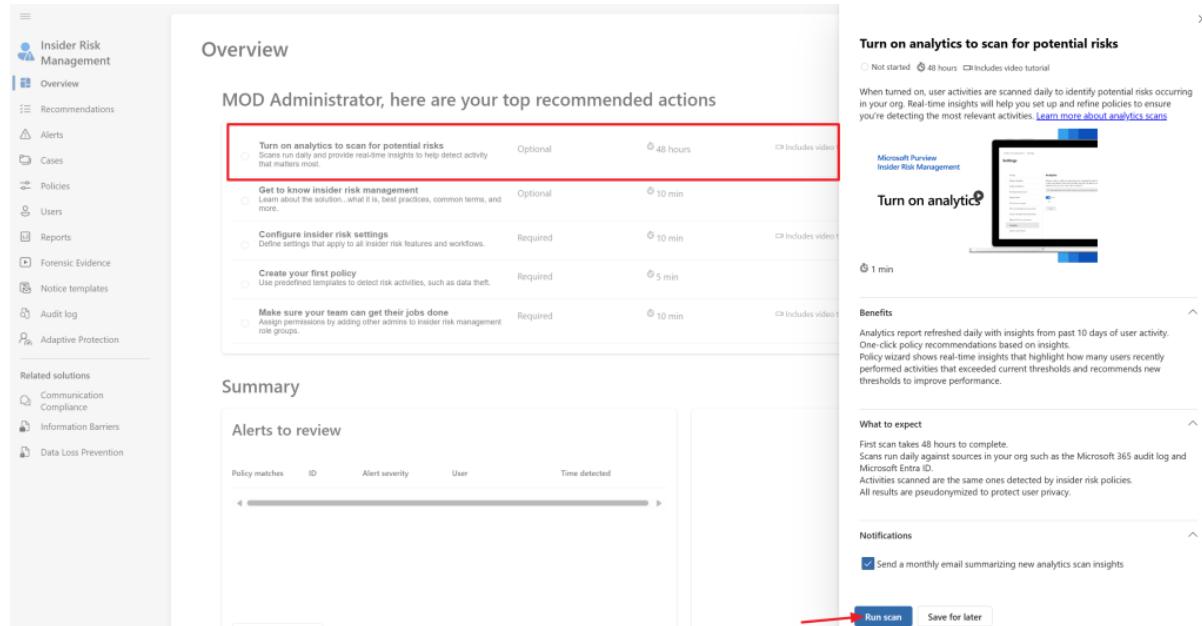
- Data theft by departing users
- Various kinds of data leaks
- Various kinds of policy violations
- Health record misuse
- Risky browser usage

IRM provides you with workflows to help your organization detect the above potential risks, manage them by cases and take various actions on the risks that are found in your environment. Ready to find out more by using an example? Let's go!

Prerequisites

When you first navigate to the IRM console, you are greeted with a few recommendations to get you started.

Turn on analytics



The screenshot shows the IRM console's 'Overview' page. On the left, a sidebar lists 'Insider Risk Management' under 'Overview', along with 'Recommendations', 'Alerts', 'Cases', 'Policies', 'Users', 'Reports', 'Forensic Evidence', 'Notice templates', 'Audit log', and 'Adaptive Protection'. Under 'Related solutions', it lists 'Communication', 'Compliance', 'Information Barriers', and 'Data Loss Prevention'. The main area is titled 'Overview' and shows 'MOD Administrator, here are your top recommended actions'. The first recommendation is 'Turn on analytics to scan for potential risks', which is highlighted with a red box. It includes a description: 'Scans run daily and provide real-time insights to help detect activity that matters most.', a duration of 'Optional 48 hours', and a video link 'Includes video'. Below this are three more recommendations: 'Get to know insider risk management' (duration 10 min), 'Configure insider risk settings' (duration 10 min), and 'Create your first policy' (duration 5 min). To the right of the recommendations, there is a 'Turn on analytics' section with a video thumbnail and a duration of '1 min'. Further down, there are sections for 'Benefits', 'What to expect', and 'Notifications', with a 'Run scan' button at the bottom right highlighted by a red arrow.

The first one is shown in the picture above and let's you turn on analytics to scan for potential risks. When you enable this, user activities are scanned on a daily basis to identify potential

risks occurring in your environment. The first scan can take 48 hours to complete as it scans your [audit log](#) and Microsoft Entra ID. When it's done, it provides you with an email when there are insights and IRM policy recommendations to check up on. Note that if you don't have audit logging enabled for your tenant (which is enabled by default nowadays) now is a good time to do so.

Read the Manual

When you have taken the first hurdle, the second recommendation is to read about the solution on Microsoft Learn so be sure to take a night out of your schedule to read up on the latest and greatest of IRM.

Configure global IRM settings

After this, we can get to work by configuring the global IRM settings.

Configure insider risk settings

Recommended settings to get up and running

Privacy

When users perform activities matching your insider risk policies, decide whether admins can see their actual names or pseudonymized versions. We recommend showing pseudonymized versions to mask their identities (this is on by default).

[Manage the privacy setting](#)

Policy indicators

Indicators are essentially the activities you want to detect from your users (like downloading content from SharePoint or copying files to a USB). Indicators you choose will be available when you create a policy.

At a minimum, we recommend selecting:

- All Office indicators.
- Indicators relevant to your org.
- Risk score boosters.

[Choose indicators](#)

Admin email notifications

Stay up-to-date on alerts across your org

Send me the following emails

When a new policy generates its first alert

When new high severity alerts are generated

Weekly email summarizing policies that have unresolved warnings

[Manage all IRM email notifications](#)

[Save email preferences](#)

[Mark as complete](#)

Here there are various settings to think about and take care of. Let's start with the **privacy setting**. Here, you can choose whether you want the portal to show actual usernames or pseudonymized versions of the usernames. Do note that pseudonymizing usernames disables data sharing between IRM and other solutions within Purview.

Insider Risk Management settings

Policy indicators

[Built-In Indicators](#) [Custom Indicators](#)

Insider risk policy templates define the type of risk activities you want to detect and investigate. Each template is based on indicators that trigger alerts when users perform related activities. Choose one or more indicators to include in your policy templates. [Learn more](#)

[+ New indicator variant](#)

Office indicators

- Select all
- Sharing SharePoint files with people outside the organization
- Sharing SharePoint folders with people outside the organization
- Sharing SharePoint sites with people outside the organization
- Downloading content from SharePoint
- Syncing content from SharePoint
- Downloading content from OneDrive
- Syncing content from OneDrive
- Adding people outside organization to priority SharePoint sites
- Adding people inside organization to priority SharePoint sites
- Downgrading sensitivity labels applied to SharePoint files

Like the help text tells us, **policy indicators** are the actual activities you want to detect in your environment. These can be Office indicators like in the screenshot above. For example, downloading content from SharePoint, Syncing content from OneDrive and so on. There are also indicators from other categories, like device indicators which let us monitor activities on devices. Examples here are creating or copying files to USB, printing files and using a browser to download content from a third-party site. In this case, I select all available office indicators for our demo. If the default indicators don't match your needs, you can always create your custom indicators.

The last thing to set up is which **email notifications** your admin account should receive. I leave the defaults to only receive emails when a new policy receives its first alert.

When this is taken care of the last steps let you [set up permissions for IRM](#) so your team can do their job and lastly **create your first policy**. Let's zoom in on creating a policy and what its effects are in your environment.

Setting up a policy

Let's navigate to the Microsoft Purview portal and select the Insider Risk Management solution. Next, click policies and select 'Create policy'.

The screenshot shows the Microsoft Defender Policy Wizard. On the left, a navigation pane lists steps: 'Policy template', 'Name and description', 'Users and groups', 'Content to prioritize', 'Triggering event', 'Indicators', and 'Finish'. The main area is titled 'Policy templates specify the conditions and indicators that define the risk activities you want to be alerted to.' It shows a list of policy types: 'Data theft', 'Data leaks', 'Security policy violations (preview)', 'Health record misuse (preview)', and 'Risky browser usage (preview)'. The 'Data leaks' section is selected. It includes a 'Prerequisites' section with a checked 'DLP policy' (optional) and three optional sections: 'Devices onboarded' (optional), 'Physical badging connector' (optional), and 'Connect cloud applications' (optional). The 'Triggering event' section lists two options: 'User performs selected exfiltration activities that exceed specific thresholds.' and 'User performs an activity matching specified DLP policy.' The 'Activities detected include' section lists three items: 'Downloading files from SharePoint', 'Printing files', and 'Copying data to personal cloud storage services'. A red box highlights the 'Data leaks' section, and a red arrow points to the 'Next' button at the bottom left.

Immediately we are shown with a dialog asking us what kind of policy that we want to create. Let's select 'Data leaks'. Remember those data sources we talked about? Here it's time to review which data sources are mandatory or optional for this policy type. As I already have a Data Loss Prevention (DLP) policy in my environment, it is found by the wizard. As can be seen it's also possible to connect a physical badging connector to see to which physical locations a user had access and use this in the policy or you can choose to connect cloud applications like Dropbox and Google drive, so you can monitor them for data being uploaded or downloaded to or from the services.

To trigger this policy, a user has to perform an exfiltration activity that exceeds specific thresholds or a user performs an activity matching specified DLP policy. Activities that are detected in this policy are downloading files from SharePoint, printing files or copying data to personal cloud storage services (which is why one would want to connect Dropbox and the likes).

Next, name your policy and give it a nice description.

Choose users, groups, & adaptive scopes

Choose users, groups, and adaptive scopes within your organization who this policy will apply to.

- All users, groups, and adaptive scopes
- Specific users, groups, and adaptive scopes

Users

Users		
Name	Email	Remove
Patti Fernandez	PattiF@...OnMicrosoft.com	X

Groups

Groups		
Groups		
Name	Email	Remove

No groups added yet.

Adaptive scopes

Adaptive scopes		
Adaptive scopes		
Name	Scope type	Remove

No adaptive scopes added yet.

Next we can use an adaptive scope or select users and groups to which we want to apply this policy. To make things simple for this example, I've chosen to apply the policy to a user named 'Patti Fernandez' only. In the next screen, we can choose to prioritize content on a specific SharePoint site, or content that contains sensitivity labels or certain sensitive info types.

Sensitive info types to prioritize

Any activity associated with content that contains this sensitive info will be assigned a higher risk score.

Add or edit sensitive info types		7 items
Info type ▾		
Australia Bank Account Number		X
Canada Bank Account Number		X
International Banking Account Number (IBAN)		X
Israel Bank Account Number		X
Japan Bank Account Number		X
New Zealand bank account number		X
U.S. Bank Account Number		X

In this example, I have chosen to prioritize content in the 'Sales and Marketing' SharePoint site, content that holds the 'Confidential – Finance' sensitivity label or one of the sensitive info types that are shown above. If you want, you can also choose to get only alerts for activities that include priority content, or get an alert for all activity that is seen in your environment that matches this policy.

- Policy template
- Name and description
- Users and groups
- Content to prioritize
- Triggering event
- Indicators
- Finish

Choose triggering event for this policy

Choose one or more triggering events to determine when a policy will begin assigning risk scores to a user's activity. [Learn more](#)

User matches a data loss prevention (DLP) policy
Policy will start assigning risk scores when a user performs an activity matching the DLP policy you select. The DLP policy must be configured to generate 'High' severity incident reports. [Learn more about DLP policy requirements](#).

Select a DLP policy

User performs an exfiltration activity
Policy will start assigning risk scores when specific thresholds are detected for activity relating to the following indicators:

Select which activities will trigger this policy

① Unable to select some indicators? This is because they're currently turned off in your organization. To make them available to select, you can turn them on now.

Downloading content from SharePoint

Sending email with attachments to recipients outside the organization

Printing files

Creating or copying files to USB

Using a browser to upload files to the web

Sharing SharePoint files with people outside the organization

File copied to remote desktop session

Now for the part that we (or at least I) have been waiting for. How is this policy triggered? We can do so by choosing a DLP policy that has to be matched OR select an exfiltration activity that takes place. I've chosen the latter option. In this policy, a user (Patti Fernandez in this case) has to send an email with attachments to recipients outside the organization or share SharePoint files with people outside the organization to trigger this policy.

Select which sequences will trigger this policy

i Some sequences require specific indicators to be turned on in 'Settings' before they can be selected below.

[Turn on indicators](#)

Download from Microsoft 365 location then exfiltrate

Download from Microsoft 365 location, obfuscate, then exfiltrate

Download from Microsoft 365 location, exfiltrate, then delete

Download from Microsoft 365 location, obfuscate, exfiltrate, then delete

Archive then exfiltrate

Archive, obfuscate, then exfiltrate

Archive, exfiltrate, then delete

Archive, obfuscate, exfiltrate, then delete

Downgrade or remove label then exfiltrate

Downgrade or remove label, download, then exfiltrate

Next, we have to choose which sequence will trigger the policy. In my case I've selected the ones above. Exfiltration is the fact that data is stolen from your environment.

Choose thresholds for triggering events

The policy will start assigning risk scores to activity only when specific thresholds are met for the exfiltration activities you selected as the triggering event. Thresholds are based on the number of events recorded for an activity per day. You can use recommended thresholds or specify your own.

Apply built-in thresholds RECOMMENDED

Choose your own thresholds

In the next screen, we have to choose which thresholds we want for triggering events. For example, here you can choose the number of times that a document is downloaded from a Microsoft 365 location and then exfiltrated.

Indicators

The following indicators are used to generate alerts for the activity detected by the policy template you selected. [Learn more](#)

Total indicators selected: 32/79

① Unable to select some indicators? This is because they're currently turned off in your organization. To make them available to select, you can turn them on now.

[Choose indicators](#)

Office indicators (28/28 selected)

- Select all
- Sharing SharePoint files with people outside the organization
- Sharing SharePoint folders with people outside the organization
- Sharing SharePoint sites with people outside the organization
- Downloading content from OneDrive
- Syncing content from OneDrive
- Downloading content from SharePoint
- Syncing content from SharePoint
- Adding people outside organization to priority SharePoint sites
- Downgrading sensitivity labels applied to SharePoint files
- Removing sensitivity labels from SharePoint files

[Back](#) [Next](#) [Cancel](#)

Next we have to select indicators (yes, again) which will generate an alert for the activity detected by the policy template we selected (Data leaks in our case). Let's choose to select all. The same alert generation options have to be chosen for detected sequences.

Cumulative exfiltration detection

Detects when the number of exfiltration activities that a user performs over a certain time exceeds the normal amount performed by users in your org over the past 30 days (for example, if a user shared more files than most users over the past month). [Learn more about cumulative exfiltration detection](#)

① Your settings are configured to detect when a user's exfiltration activities exceed organization norms and peer group norms. You can adjust your configuration in settings. [Change Cumulative exfiltration detection preferences in settings](#)

- Select all
- Detect when a user's exfiltration activities exceed norms

Risk score boosters

Risk scores for policy alerts might be increased for these circumstances. For example, a medium-severity alert might be increased to high.

- Select all
- Activity is above user's usual activity for that day

In the 'cumulative exfiltration detection' screen, we can configure whether the service should detect when the exfiltration activities that a user performs over a certain time exceeds the normal amount for users in your org over the past 30 days. This can then be used to boost the risk score in your policy.

Choose threshold type for indicators

Each indicator you selected uses thresholds to influence the activity's risk score, which in turn determines whether an alert's severity is low, medium, or high. Each threshold is based on the number of events recorded for an activity per day.

- Analytics has been turned on. The first scan takes 48 hours to complete. We will have personalized threshold recommendations ready for you once the scan is complete. Return afterwards to update your selection. X
- Apply thresholds provided by Microsoft
Built-in thresholds will be applied to all indicators you selected.
Built-in thresholds will be applied to all indicators you selected.
- Apply thresholds specific to your users' activity RECOMMENDED
Thresholds based on your users' recent activity patterns will be applied to all built-in indicators you selected.
Thresholds based on your users' recent activity patterns will be applied to all built-in indicators you selected.
- Choose your own thresholds
Customize thresholds that are prepopulated with built-in values from Microsoft.
Customize thresholds that are prepopulated with built-in values from Microsoft.

The indicators we selected before (downloading content from SharePoint, Syncing content from OneDrive and so on) uses thresholds that influence the activity's risk score, which in turn will determine what the severity for an alert is. For these thresholds we can configure to use the Microsoft defaults or use our own thresholds. However, as can be seen in the screenshot above, after 48 hours it's possible to let the service use personalized threshold recommendations that are based on your users' recent activity patterns!

Lastly, as always, review the summary and submit the new policy. Take notice of the screen that tells you that it might take up to 24 hours before policy matches will start showing up on the Alerts tab.

Good user gone bad

The screenshot shows a SharePoint document library interface. A file named "Cross Cultural Marketing Campaigns.pptx" is selected, indicated by a checkmark icon. A sharing dialog box is overlaid on the screen, asking "Share outside of your organization?" with "Continue" and "Cancel" buttons. The background shows a list of other files including "General", "Monthly Reports", "Branding Elements.pptx", "Cross Cultural Marketing Cam...", "DG-1000 Product Overview.pptx", "DG-2000 Product Overview.docx", "DG-2000 Product Pitch.pptx", "DG-2000 Product Specification.docx", "International Marketing Campaigns.docx", and "Marketing To Your Audience.docx".

Now let's enter Patti Fernandez as our 'good user gone bad'. Patti starts to share a couple of documents from SharePoint to an external user. She sees the notification above but chooses to ignore it and share the data anyway. She shares a couple of files in the same way.

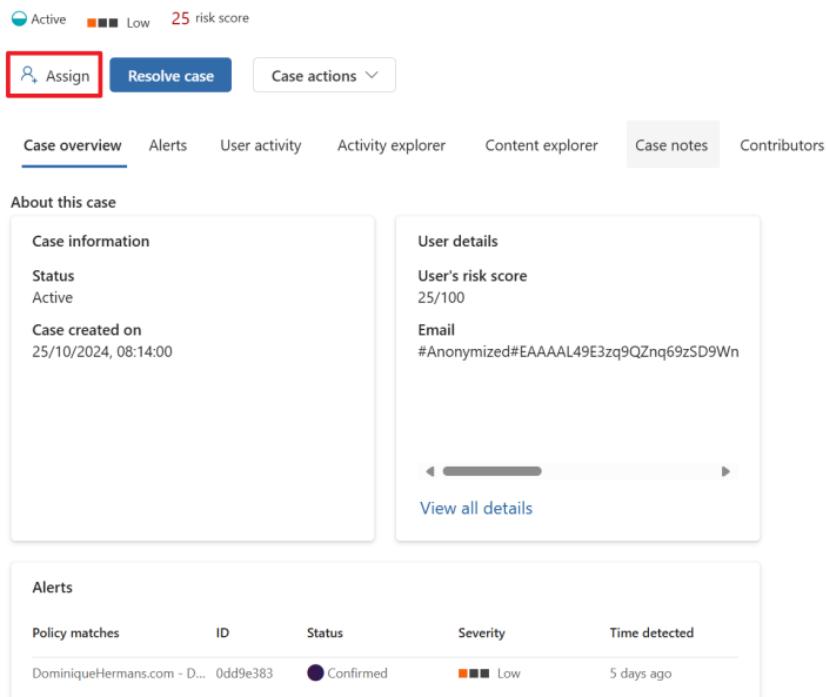
Red alert!

Now let's return to our admin user and select the 'Alerts' screen in IRM. We can see that an alert is generated for a user (which cannot be seen because we choose to pseudonymize usernames) and the status for the alert is set to 'Needs review'. We can also see that the alert is not yet assigned to a user tasked with taking care of alerts and no case has been generated yet.

When the alert is opened, we are provided with a whole world of details. What was the triggering event (which I might have slightly manipulated for the policy to trigger 😊), the top exfiltration activities, the severity of the alert and some other details. Let's confirm the alert and create a case by clicking 'confirm all alerts & create case'.

Files shared from OneDrive

(371ed16a) Files shared from OneDrive



The screenshot shows the IRM portal's Case overview screen for a case titled 'Files shared from OneDrive'. The top navigation bar includes 'Case overview', 'Alerts', 'User activity', 'Activity explorer', 'Content explorer', 'Case notes' (which is the active tab), and 'Contributors'. Below the navigation is a section titled 'About this case' with two panels: 'Case information' (Status: Active, Case created on: 25/10/2024, 08:14:00) and 'User details' (User's risk score: 25/100, Email: #Anonymized#EAAAAL49E3zq9QZnq69zSD9Wn). A 'View all details' button is at the bottom of this section. The 'Alerts' section shows one policy match: 'DominiqueHermans.com - D...' with ID '0dd9e383', Status 'Confirmed', Severity 'Low', and Time detected '5 days ago'. A red box highlights the 'Assign' button in the top right corner of the main content area.

After moving to the cases screen in the IRM portal, select the case we just created and let's assign it to a fellow IRM employee to start reviewing the case. Let's go through each tab of this interface to see what information is provided to us via the case. Note that a case can have multiple alerts attached to it, in this demo it is only 1 alert, which can be seen on the 'alerts' tab.

Case overview Alerts **User activity** Activity explorer Content explorer Case notes Contributors

Filter: Risk category: Any Activity Type: Any Reset all

Sort by: Date occurred

- User added to policy: Testing** ...
Oct 24, 2024 (UTC)
Policy added to: DominiqueHermans.com - Data Leaks Policy
- Exfiltration: Emails with attachments sent outside the organization** ...
Oct 24, 2024 (UTC) | Risk score: 5/100
5 emails: sent to 1 recipient outside the organization (Explore content)
5 emails: containing sensitive info, including: All Full Names (Explore content)
- Access: Sensitive SharePoint files accessed** ...
Oct 24, 2024 (UTC) | Risk score: 5/100
12 events: Sensitive files accessed from 1 SharePoint site (Explore content)
7 events: Sites that have labels applied, including: (Explore content)
- Access: People outside the organization added to sensitive or priority SharePoint sites** ...
Oct 24, 2024 (UTC) | Risk score: 5/100
1 event: 1 Person outside the organization added to 1 SharePoint site (Explore content) (Explore content)
- Collection: Files downloaded from SharePoint** ...
Oct 24, 2024 (UTC) | Risk score: 5/100
47 events: Files downloaded from 2 SharePoint sites (Explore content)

User activity scatter plot 6 Months 3 Months **1 Month**

User added to policy: Testing

Exfiltration: SharePoint files shared

Oct 24, 2024 (UTC) | Risk score: 25/100
2 events: Files shared (Explore content)
(Explore content) 0

RISK SCORE

Oct 1, 2024

Legend: Access (blue), Deletion (pink), Collection (yellow), Exfiltration (teal), Infiltration (purple), Obfuscation (dark brown), Security (light blue), Custom Indicator (orange), Defense Evasion (red), Privilege Escalation (dark red), Communication Risk (green), AI Usage (light green). Sequence (arrow), Cumulative Exfiltration (arrow).

The user activity tab shows us a list of all activities that lead up to this alert and it accumulates this for us in a so-called ‘user activity scatter plot’ in which you can see all user activities that happened over time. Now, this is just an example that was created here, but imagine that your HR system noted that an employee is going to leave the company in a few weeks (which would be a great trigger event), and after this the system shows all kinds of data being exfiltrated to the users personal Dropbox.

This shows the real power of the IRM solution as it uses all these events to create 1 easy to read chart that shows exactly which risk the user actions pose to your company!

Files shared from OneDrive

(371ed16a) Files shared from OneDrive

Active Low 25 risk score

MOD Administrator Resolve case Case actions

Case overview Alerts User activity **Activity explorer** Content explorer Case notes Contributors

Filter: Risk factor: Any Sort by: **User added to policy: Testing** (Oct 24, 2024 (UTC))

Export 76 items Reset columns Customize columns Save this view Views

Filter Reset Filters

Activity: Any Date (UTC): 3/29/2024-10/29/2024

Date (UTC)	Activity	File name	Object ID	Workload
Oct 24, 2024 2:04 PM	Email sent to external recip...	AttachedImage.At...	<Share-48105da1-70de-60...	IrmHygiene
Oct 24, 2024 2:03 PM	Email sent to external recip...	AttachedImage.At...	<Share-3d105da1-d0a3-60...	IrmHygiene
Oct 24, 2024 2:03 PM	File accessed on SPO	DG-2000 Product...	https://[REDACTED].sh...	SharePoint
Oct 24, 2024 2:03 PM	File accessed on SPO	DG-2000 Product...	https://[REDACTED].sh...	SharePoint
Oct 24, 2024 2:02 PM	Email sent to external recip...	AttachedImage.At...	<Share-2a105da1-f0ae-60...	IrmHygiene
Oct 24, 2024 1:58 PM	Email sent to external recip...	AttachedImage.At...	<Share-f60f5da1-6007-60...	IrmHygiene

When we select the 'activity explorer' tab we can dive a little deeper by taking a look at all the activities the service linked to this case. This really shows a nice detail of what is going on. We can even take a look at the content included in the case from the 'content explorer' tab. Really nice to have all this features in 1 GUI!

Last but not least, we can use the 'case notes' tab to write down important information about the case and use the 'contributors' tab to add other colleagues that may manage the case.

Notice templates

Insider Risk Management

Overview Recommendations Alerts Cases Policies Users Reports Forensic Evidence Notice templates Audit log Adaptive Protection

Related solutions

Communication Compliance Information Barriers

Notice templates

Create notice template

Notice templates

No data available

Create a new notice template

This template will be available to use anytime you need to send an email notice to a user included in an insider risk case.

Template name *

Breach of Ehlo Information Protection Guidelines

Send from *

MOD Administrator

CC

Bcc

Subject *

Breach of Ehlo Information Protection Guidelines

Message body *

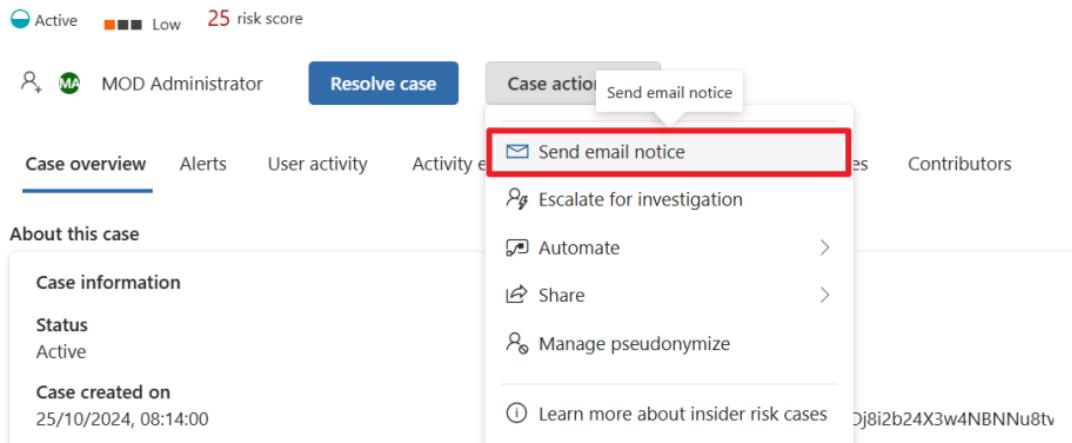
Hi,
You have breached the Ehlo Information Protection Guidelines. A meeting has been scheduled to give the opportunity to give feedback on this matter.
Kind regards,
The Legal Department

Now, with our case created let's create a notice template by navigating to IRM, Notice templates. These templates can be used as a starting point to communicate about a case with a user involved in this case. I drafted up the one that is shown above.

Case actions

Files shared from OneDrive

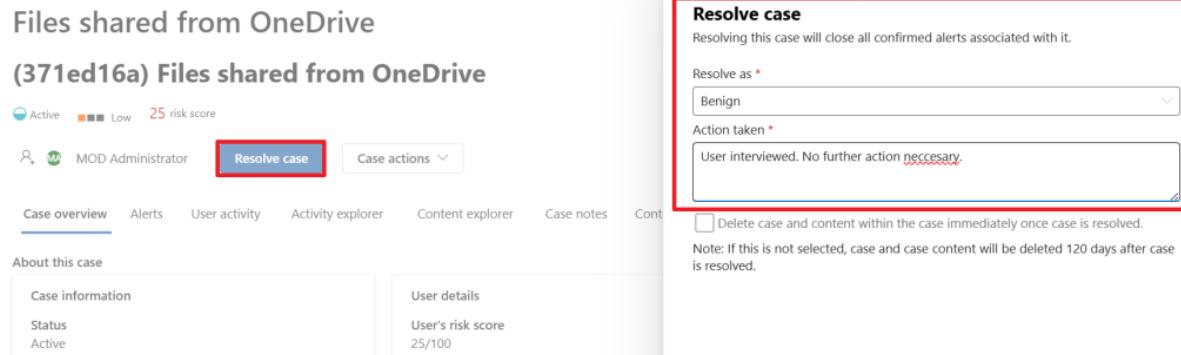
(371ed16a) Files shared from OneDrive



The screenshot shows the Microsoft 365 IRM portal interface. At the top, there are status indicators: 'Active' (green), 'Low' (orange), and a 'risk score' of '25'. Below this, the user 'MOD Administrator' is logged in. A 'Resolve case' button is visible. A 'Case actions' dropdown menu is open, showing several options: 'Send email notice' (highlighted with a red box), 'Escalate for investigation', 'Automate', 'Share', 'Manage pseudonymize', and 'Learn more about insider risk cases'. The 'Case overview' tab is selected in the main navigation bar.

Let's return to the case we've created and select 'case actions'. It provides us with a few options like escalate the case for investigation, which actually opens an [eDiscovery \(premium\)](#) case in your Microsoft 365 environment that allows the IRM case to be further reviewed, for example for legal purposes. We can also apply a Power Automate action or share the case. Lastly we can manage the pseudonymize options for this case.

Let's go for 'send email notice'. We can now select the email notice we created earlier and send the message to the user. We might not be able to identify the user because the data in the IRM portal is pseudonymized, however the system is still able to track down the user.



The screenshot shows the Microsoft 365 IRM portal interface. The 'Resolve case' dialog box is open on the right side. It contains fields for 'Resolve as' (set to 'Benign') and 'Action taken' (containing the text 'User interviewed. No further action necessary.'), both of which are highlighted with a red box. Below the dialog, there is a note about deleting the case and content. The main interface shows the 'Case overview' tab selected, with 'Case information' and 'User details' sections visible.

Lastly, depending on the input of our user, the case would be resolved as 'benign' which closes the case as no further action is necessary. It's also possible to resolve the case as a 'Confirmed policy violation' in which case you have to take appropriate actions depending on the severity of the case.

Mastering the (Unified) Audit Log

The Microsoft Purview (Unified) Audit Log. Not the first component of Purview you think of when there's Data Loss Prevention, Data Lifecycle Management and other cool features. However, the most basic feature like the Audit Log can be quite interesting. So in this chapter, I want to take you through the basics of the Unified Audit Log.

The Audit log is the place where a lot of user and administrator interactions with the various Microsoft 365 services are stored so they are accessible and searchable for security ops, IT admins, insider risk teams, and compliance and legal investigators in your organization. That's exactly why it's called the **unified** audit log. It collects almost everything from the various services in your Microsoft 365 subscription.

Let's talk basics

The Purview Audit Solution comes in 2 flavors: Audit Standard and Audit Premium. Both are enabled by default in newer M365 tenants. If you have an older tenant, you'll need to use the following command in Exchange Online PowerShell to verify whether the Audit Solution is enabled in your tenant:

```
Get-AdminAuditLogConfig | Format-List UnifiedAuditLogIngestionEnabled
```

While both flavors come packed with the Audit Search tool in Purview and Compliance portals and Search-UnifiedAuditLog PowerShell cmdlet to search through audit events, exportable audit records to CSV's and access to audit logs via the Office 365 Management Activity API (albeit that Audit Premium includes higher bandwidth access to the API), differences are noticeable in the following area's:

Audit Standard has a 180-day audit log retention where Audit Premium can be configured for up to 10-year audit log retention. Also, Audit Premium includes Audit Log retention policies and Intelligent Insights. Intelligent insights provides access to important events to conduct forensic and compliance investigations. Take a look at the following [Learn Article](https://learn.microsoft.com/en-us/purview/audit-solutions-overview#audit-premium-activity-properties) (<https://learn.microsoft.com/en-us/purview/audit-solutions-overview#audit-premium-activity-properties>) to see the activity properties that are supported.

Besides the above, you have to be assigned the 'Audit Logs' or 'View-Only Audit Logs' roles in the Microsoft Purview (compliance) portal to be able to search the audit log. These roles are part of the 'Audit Manager' and 'Audit Reader' role groups.

Another thing to note is that Microsoft doesn't guarantee the time it takes for an audit log to be available (or searchable) after an auditable event has happened. It does note however that for the general services (Exchange, SharePoint, OneDrive and Teams) availability is typically 60-90 minutes after an event occurs.

What information is stored in the Unified Audit Log?

The answer to this question depends on the type of license that is assigned to specific users. The license type defines which user or admin activity generates an audit record and for how long this record is stored (and thus is searchable) in the audit log.

- When a user has a Office 365 E5, Microsoft 365 E5 or Microsoft 365 E5 Compliance or Microsoft 365 E5 eDiscovery and Audit add-on license, audit records for the following services are retained for 1 year by default:
 - Microsoft Entra ID
 - Exchange
 - SharePoint (and thus OneDrive)
- Organizations with the licenses mentioned above can create audit log retention policies to retain audit records for activities in other services for longer retention periods than the default of 180 days for these other services.
- Users that are assigned any other license than the ones mentioned in the first bullet, audit records are retained for 180 days.

Using Audit Log Search

Let's dive into the portal and take a look at how Audit Log Search works.

The screenshot shows the Microsoft Purview Audit Log search interface. The left sidebar has a 'Audit' solution selected (1). The main area has a search filter section (2) with fields for date range, activities, users, and keyword search. Below it is a table of previous search jobs (3).

Search name	Job status	Progress	Search time	Total results	Creation time	Search performed by
Jun 13 - Aug 13 dlprulematch	Completed	100%	7m, 8s	21	Aug 13, 2024 8:27 ...	admin@contoso.onmicrosoft.com

When you navigate to the Microsoft Purview Console and select the Audit Solution on the left hand side (1), you'll enter the main Audit Console which consists of 2 parts:

- The filter section to narrow down your search of all items in the Audit Log (2). The more filters you configure for your search, the more narrowed down your search result will be.
- The previous search jobs section that shows previous search jobs and their status.

Let's go and run through an example to get an idea of how this works.

1 Keyword Search

2 Activities - friendly names

3 Users

4 Workloads

For our audit search, let's configure the search filter so that we're searching for all accessed files in SharePoint by Grady Archie in the last week:

- Date and Time range set to last week (1).
- Activity is 'accessed file' (notice the nice drop down list on 'activities – friendly names') (2).
- User is Grady Archie (3).
- Workload is 'SharePoint' (also selectable by a drop down list).

Search name	Job status	Prog...	Search...	Total results	Creation tim...	Search performed by
Aug 7 - Aug 14 GradyA fileaccessed SharePoint	Completed	100%	2m, 12s	23	Aug 14, 2024 9:09 ...	admin@ onmicrosoft.com

When pressing Search a new line appears in the 'previous search jobs' section. Now we wait until the 'Job Status' changes to 'completed' and click it so we can see the results. Notice that the search name is automatically filled in for us (since we lacked to do so).

1 Aug 13, 2024 12:33 PM

2 Details

3 Details

4 Date (UTC)

5 IP Address

6 Users

7 Activity

8 Item

9 Admin Units

10 AppAccessContext

11 CreationTime

12 Id

13 Operation

14 Close

The results pane presents us with all audit log entries that match our search filter (1). When selecting one of the entries, the details pane (2) shows us more information. Ranging from CreationTime to UserID to the fact if the device that accessed the file is a managed device or not.

If you want to take a look at all the available properties take a look at this [Microsoft Learn Article](https://learn.microsoft.com/en-us/purview/audit-log-detailed-properties) (<https://learn.microsoft.com/en-us/purview/audit-log-detailed-properties>). If you would like to see some scenarios where the Audit Log could help you troubleshoot common support issues, take a look at [this Learn Article](https://learn.microsoft.com/en-us/purview/audit-troubleshooting-scenarios) (<https://learn.microsoft.com/en-us/purview/audit-troubleshooting-scenarios>).

Hope you learned something in this chapter!

Utilizing the Content and Activity Explorer

In this chapter I want to take you along the likes of the Content Explorer and Activity Explorer. According to the Purview documentation, we can leverage the Content Explorer to explore email and documents in your environment that contains sensitive information or items that have labels applied. Activity Explorer can be used to take a look at all the actions that took place with sensitive info or items that have labels applied.

Let's dive right in and start with the possibilities of Content Explorer

Content Explorer

Let's start taking a look at Data classification, overview. Here you can see a glance at the sensitive information and labels used in your environment. A lot of this data is actually coming from the Content Explorer.

Content explorer

Explore the email and docs in your organization that contain sensitive info or have labels applied. You drill down further by reviewing the source content that's currently stored in Exchange, SharePoint, and OneDrive. Support for more locations is coming soon. [Learn more](#)

Sensitive info types	Count
All Full Names	287
Credit Card Number	61
New Zealand Social Welfare Number	57
Hungarian Social Security Number (TAJ)	50
Portugal Tax Identification Number	50
All Medical Terms And Conditions	48
EU Debit Card Number	48
U.S. Bank Account Number	47

Name	Count	Location
Exchange	28	>
OneDrive	24	>
SharePoint	9	>
Teams	0	>

Now, when clicking through to Content Explorer on the left hand side, we see a current snapshot of all items in your environment that have:

- a sensitivity label
- a retention label
- have been classified as a sensitive information type you defined or are available by default.

Let's talk about permissions next, which are managed on different levels.

First, to get access to the Content Explorer, you'll have to have the Global administrator, Compliance administrator, Security administrator or Compliance data administrator role.

Second, because the Content Explorer shows files (and it's content) that contain sensitive information, you'll have to have more permissions to see the file title or contents:

- Content Explorer List viewer (and specifically it's data classification list viewer role) gives you permission to see the item and it's location.
- Content Explorer Content viewer (and specifically it's data classification content viewer role) gives you permission to view the contents of each item in the list.

Content explorer

Explore the email and docs in your organization that contain sensitive info or have labels applied. You drill down further by reviewing the source content that's currently stored in Exchange, SharePoint, and OneDrive. Support for more locations is coming soon. [Learn more](#)

	Name	SSN	Bank Account number	Card Type	Credit card
1	Name	SSN	Bank Account number	Card Type	Credit card
2	Alex Darrow	533414056	980441377-8026967	Visa	448578659
3	Allie Bellew	534558859	364684619-7519006	Visa	448565116
4	Anne Wallace	532258026	231951785-0182689	Visa	453908742
5	Bonnie Kearney	533414573	948490562-6112913	MasterCard	522464339
6	Garth Fort	531337513	804785034-5007576	MasterCard	519873946
7	David Longmuir	169482480	980441377-8026945	MasterCard	554566598
8	Denis Dethenne	454814186	364684633-7519006	MasterCard	533931753
9	Doreen Paschke	540453187	231951785-2382689	MasterCard	543726953
10	Garret Vargas	246575457	948490562-6112913	MasterCard	51125047
11	Janet Schorr	600434056	804785034-5004276	MasterCard	528227467
12	Julian Isla	284076634	948491662-6112913	Visa	402400715
13	Katie Jordan	324077404	804785034-5049576	Visa	492953228
14	Robin Counts	511239308	980441377-8033545	Visa	453260213

When diving in Content Explorer, let's check out what's in store here. I've clicked the 'Credit Card Number' SIT, where content explorer shows me that there are actually 4 items in my environment that match the SIT. When drilling down further into the SharePoint site in this case, it shows me the actual files and even contents of these files, because I have the appropriate permission. The possibilities in this view are as follows (numbers below match numbers on the screenshot above):

1. **Filter** on label name, SIT or categories. For instance, if you don't want to scroll all the way through the list, you can just type in 'credit' and it shows you only the Credit Card Number info type.
2. **Export**. Provides you with a CSV file with all the content that is currently on screen.
3. **Search**. Search for items in the current view. In the screenshot above, it allows me to search for items with the Credit Card Number SIT in the SharePoint site 'newemployeeonboarding'.
4. **Detailed view**. Shows the contents of your items (if you have the correct permission).
5. **Provide Feedback** on a matched SIT or trainable classifier. This feedback can be used to further optimize your SIT's / classifiers.

In conclusion, a nifty explorer to find out all about the sensitive content in your environment, and where it resides.

Activity Explorer

Now let's turn our attention to the Activity Explorer. In short, the activity explorer gives you an historical view of activities on your labeled content. This information is actually sourced from the Microsoft 365 unified audit logs and made available in the activity explorer UI. This data is available for 30 days and can be filtered using over 30 different pre-configured filters.

Just as with the content explorer, let's talk permissions first.

Microsoft Purview Roles	Microsoft Purview Role Groups	Microsoft 365 Roles	Microsoft 365 Role Groups
Information Protection Admin	Information Protection	Global Admins	Compliance Administrator
Information Protection Analyst	Information Protection Admins	Compliance Admins	Security Administrator
Information Protection Investigator	Information Protection Investigators	Security Admins	Security Reader
Information Protection Reader	Information Protection Analysts	Compliance Data Admins	
	Information Protection Readers		

Image source: [Microsoft](#)

To get access to the activity explorer, you need one of the roles above or be part of one of the role groups described here.

Activity explorer



Let's take a look at what's going on here. I've filtered my view so it provides me with a certain time range. In that time range a Data Loss Prevention rule was matched in my environment. Options we have here are (numbers below match with numbers on the screenshot above):

1. Built-in filters that you can choose from are for example Endpoint DLP activities, Egress activities or activities for a certain target domain for example. This actually doesn't stop there. With the filters button in the same row, you can define your own filters using a vast numbers of properties to choose from.
2. When you change your filter, the properties that you can choose from on screen change. This way you can filter on properties that belong to a pre-defined filter.
3. A graph is shown of when an activity happened and how many items in your current view took place in that timespan.
4. The ability to export the current list of activities to a CSV file.
5. The current activities in your view.

In conclusion, the activity explorer provides you with a nice activity overview of activities in your environment which allows you to see if the controls that you have in place in your environment are effective or whether you need to finetune your configuration!

Configuring Alert Policies for High Risk Activities

In some cases you may want to be informed immediately when certain actions are being performed in your Microsoft 365 environment by your users. Examples are documents being shared with external parties that should not have access to the documents, or maybe you have a certain user that you want keep tabs on. Of course there are many ways to achieve this in Microsoft Purview, and the configuration of alert policies for these high risk activities is one of them.

A word on RBAC Permissions

To start off with the necessities, the required RBAC permissions to view alerts can be found on this [Microsoft Learn page](https://learn.microsoft.com/en-us/purview/alert-policies#rbac-permissions-required-to-view-alerts) (<https://learn.microsoft.com/en-us/purview/alert-policies#rbac-permissions-required-to-view-alerts>). However this isn't one simple permission that grants a user or administrator the permissions to view all alerts. As alerts are categorized, the user or admin tasked with viewing alerts has to have permissions to view alerts in the specific category.

Alert Policies Overview

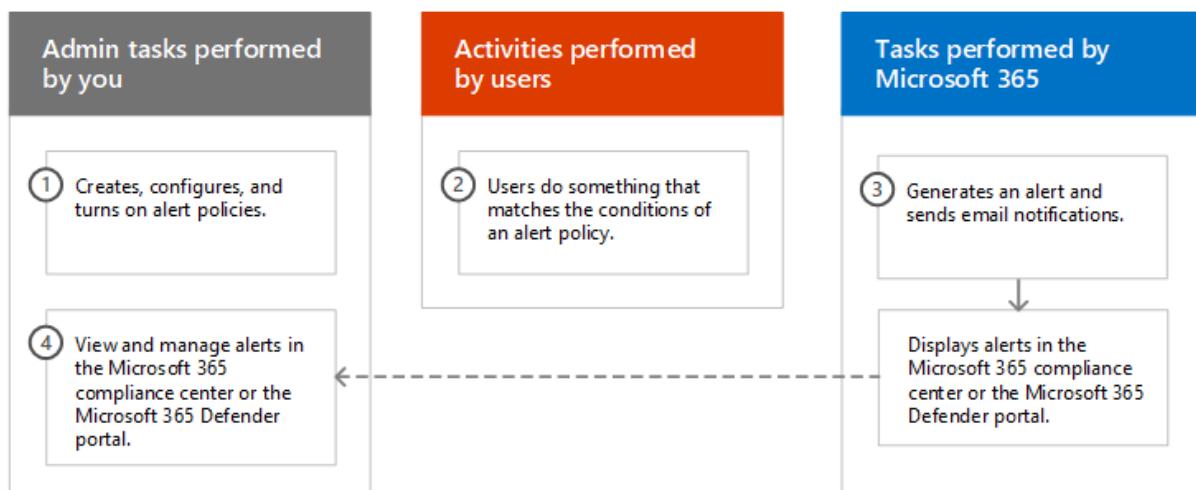


Image credit: Microsoft

As you can see in the image above, alert policies work as follows:

1. An administrator creates, configures and turns on alert policies.
2. A user performs an action that matches the conditions of an alert policy.
3. Microsoft 365 generates an alert, sends email notifications and displays alerts in the Microsoft 365 compliance center and the Microsoft 365 Defender portal.
4. An administrator views and manages alerts in the Microsoft 365 compliance center or Defender Portal.

Configuring Alert Policies

Let's take a look at how we can configure a policy and how it works when a policy is triggered by the user.

Microsoft Purview

Home

Compliance Manager

Data classification

Data connectors

AI Hub (preview)

Alerts

Policies

Roles & scopes

Trials

Solutions

Catalog

Audit

Content search

Communication compliance

Policies

Set up policies to govern data, manage devices, and receive alerts about various activities in your org. [Learn more about these policies](#)

Policy type

Alert (3)

Alert

Defender for Cloud Apps

Alert policies

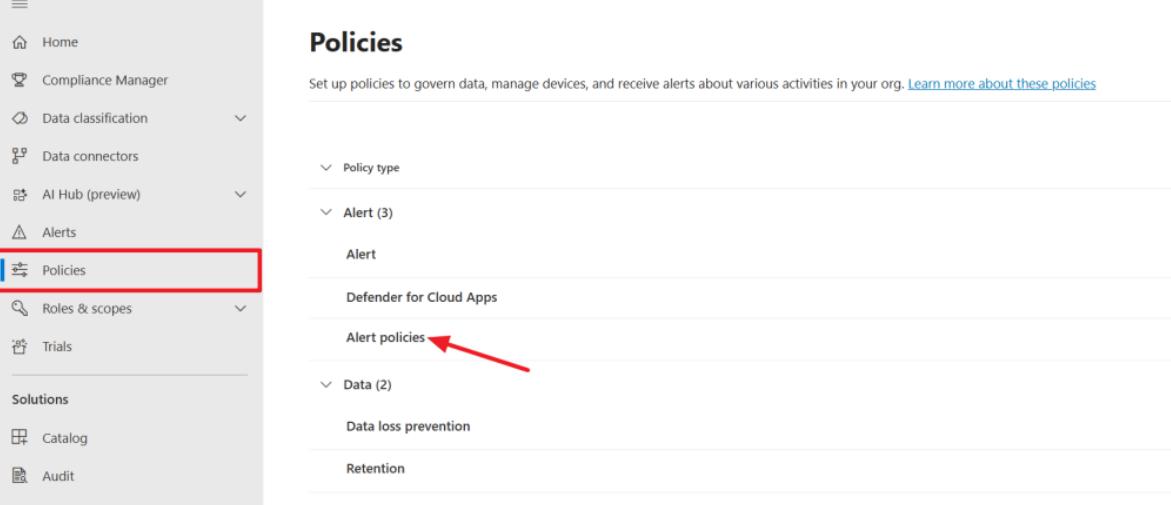
Data (2)

Data loss prevention

Retention

Access (1)

Terms of use



Navigate to the Microsoft Purview Compliance Center and click on Policies, Alert Policies.

As you can see we are immediately redirected to the alert policies in the Microsoft Defender console (which can be found manually by navigating to ‘Email and Collaboration’, ‘Policies and Rules’, ‘Alert Policy’). Notice that there are a lot of policies already present out of the box. Let’s click on ‘New Alert Policy’ to start configuring a new alert policy to get a peak inside the inner workings of alert policies.

New Alert Policy

Name your alert

Create alert settings

Set your recipients

Review your settings

Name your alert, categorize it, and choose a severity.

Assign a category and severity level to help you manage the policy and any alerts it triggers. You'll be able to filter on these settings from both the 'Alert policies' and 'View alerts' pages.

Name *

Grady A - Sharing Files Externally

Description

Alert that is triggered when Grady A shares files with an external party.

Severity * ⓘ

High

Category *

Information governance

In the first screen of the wizard we can name and describe the policy accordingly and assign it a severity and category. Note that you may want to bundle certain alerts under a certain severity. Do also note that configuring a category requires a user to have a specific RBAC role belonging to the category before he or she can view and manage the alert in this category.

New Alert Policy

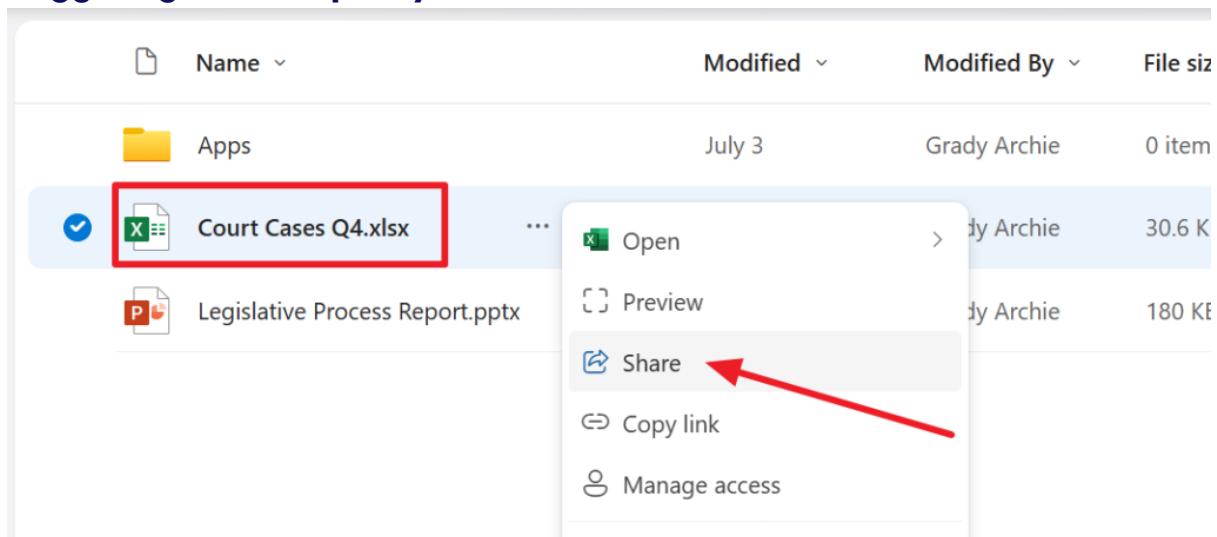
Next we have arrived at the activity and conditions page, where we can configure both activities and conditions that trigger the alert policy. This page consists of a massive library of activities that you can generate an alert on, which can be combined with other conditions that are specific to the activity you selected. In this example I configure the activity to be 'Shared file externally' AND user equals 'Grady Archie'. I also want the alert to trigger every time Grady shares a file with an external party. You may also choose to only fire the alert when a certain threshold is reached, or 'when the volume of matched activities becomes unusual'.

New Alert Policy

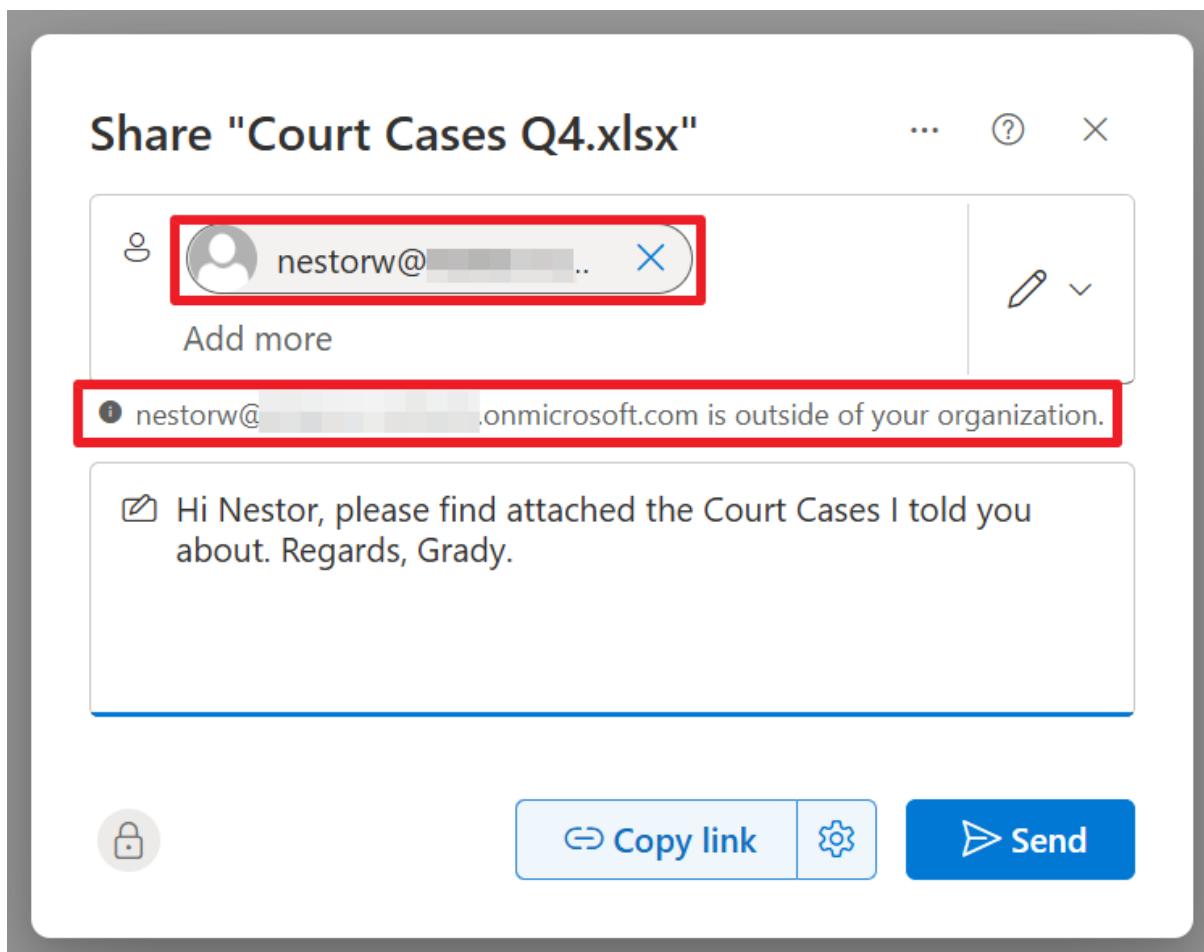
On the next page you can decide if you want to notify people when the alert is triggered. I can imagine that for alerts with a higher sensitivity level, emails will be sent. For alerts with lower sensitivity levels, an alert in the dashboard will suffice. I find the 'Opt-In for Email Notifications' text a bit misleading here, as the option only lets you specify email recipients here together with a daily notification limit. So when turning it off, the option to configure email recipients disappears. It doesn't give users the option to opt-in for email notifications later from the portal for example. I choose to send my admin an email when this alert is triggered.

In the last screen, doublecheck your settings and turn on the policy. At this point, it can take up to 24 hours for the alert policy to take effect as the policies have to be synced to the alert detection engine.

Triggering the alert policy



Now let's log in as Grady Archie and share a file with someone from an external company.



In this case, I share the file with my good friend “Nestor” whom I’ve told about this list of Court Cases I’ve been involved in.

Examining the results

Note that Nestor receives the email with no problem, as we don't block anything by setting an alert policy. However, the following does happen:

The screenshot shows the Microsoft Defender Console interface. The left sidebar has a red arrow pointing to the 'Alerts' link under 'Incidents & alerts'. The main area is titled 'Alerts' and shows a list of alerts. The first alert in the list is highlighted with a red box. The alert details are: Name: 'Grady A - Sharing Files Externally', Severity: High (indicated by three red squares), Status: Queued (blue dot), Category: Information governance, Detection source: MDO, and Impacted assets: gradya, Grady Archie.

After a while, an alert shows up in the Alerts panel in the Microsoft Defender Console.

The screenshot shows the Microsoft Defender Console interface with the 'Alerts' link highlighted in the left sidebar. The main area shows the details of the 'Grady A - Sharing Files Externally' alert. The 'Activity list' section is highlighted with a red box and shows a single entry: 'SharingInvitationCreated' on Aug 20, 2024 7:50 PM by user gradya@... with item https://... and IP address ...'. The right side of the screen displays various alert details such as Classification (Not Set), Alert state (Assigned to Unassigned), and Alert details (Category: Information governance, Detection source: MDO, etc.).

When opening the alert, it shows us all kind of information:

- The activity that triggered the alert.
- Category of the alert.
- Detection Source.
- Date and time of the activity.
- The user that triggered the alert.
- Activity being performed by the user.
- The item on which the activity is being performed.
- IP address where the action took place from.

We can perform additional actions on the alert:

- Set the classification of the alert.
- Assign the alert to someone for further investigation.

- Link the alert to another incident to find a correlation.

The screenshot shows the Microsoft Purview Alerts interface. On the left, a navigation menu includes Home, Compliance Manager, Data classification, Data connectors, AI Hub (preview), Alerts (selected), Policies, Roles & scopes, Trials, Solutions, Catalog, Audit, Content search, Communication compliance, Data loss prevention, eDiscovery, Data lifecycle management, Information protection, Information barriers, Insider risk management, Records management, Privacy risk management, and Subject rights requests. The main area is titled 'Alerts' and shows a list of alerts. One alert is highlighted with a red box: 'Grady A - Sharing Files Externally' (Severity: High, Status: Active, Investigation state: Investigating). To the right of the alert list is a detailed view of the selected alert. The 'Alert details' section shows the alert name, severity (High), date (Aug 20, 2024 7:51 PM), category (Information governance), activity count (1), user (gradya@[REDACTED].onmicrosoft.com), and activity (ExternalFileSharing). The 'Alert status' section shows the status (Active) and comment (New alert). The 'Policy details' section shows the alert policy (Grady A - Sharing Files Externally) and notification sent to (admin@[REDACTED].onmicrosoft.com). A red arrow points to the 'Resolve' button in the top right of the alert details view.

When taking a look at the Alerts section in the Purview portal, notice that the alert is also registered here. Here also various details about the event that triggered the alert are available, albeit with less details than in the Microsoft Defender Console.

A high-severity alert has been triggered

⚠️ Grady A - Sharing Files Externally

Severity: ● High

Time: 8/20/2024 5:51:00 PM (UTC)

Activity: ExternalFileSharing

User: gradya@[REDACTED].onmicrosoft.com

Details: Alert that is triggered when Grady A shares files with an external party.

See details in the Microsoft 365 security center

[View alert details](#)

Thank you,
The Office 365 Team

As expected, an email was also sent to my administrator providing details of the alert and the possibility to view alert details.

Alerts

[Export](#) 1 Week

Filter set:

[Add filter](#)

Alert name

Tags

Severity

Investigation state

Status

Grady A - Sharing Files Externally

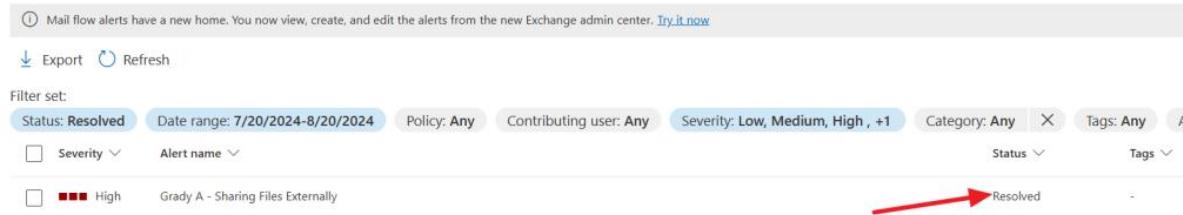
■■■ High

Queued

Resolved

Let's resolve the incident by pressing 'resolve' in the Purview console. After specifying a comment why the issue was resolved the alert will not be present in the default view anymore, but can be seen when changing the filter to include resolved alerts.

Alerts



The screenshot shows the Microsoft Purview Alerts interface. At the top, a message says "Mail flow alerts have a new home. You now view, create, and edit the alerts from the new Exchange admin center. Try it now." Below this are "Export" and "Refresh" buttons. The filter bar includes "Status: Resolved", "Date range: 7/20/2024-8/20/2024", "Policy: Any", "Contributing user: Any", "Severity: Low, Medium, High, +1", "Category: Any", "Tags: Any", and dropdowns for "Severity", "Alert name", "Status", and "Tags". A search bar is also present. A red arrow points to the "Resolved" status filter. Below the filter bar, a table lists one alert: "High" severity, "Grady A - Sharing Files Externally".

When moving back to the Microsoft Defender console, note that the alert was resolved here as well!

Summary

- Alert Policies can be used to notify you when a certain action is taken in your Microsoft 365 environment.
- Ideally you configure alert policies for high risk activities that you want to monitor.
- Alerts can be seen in the Purview and Defender console where the latter gives you more detail.
- Alerts can be emailed to recipients when configured, ideally you'll only do this for high risk activities.

That's it! Hope you learned something new in this chapter!

Information Barriers

Information Barriers in Microsoft 365 can be used to block (or allow) communication between groups of users and will apply to Teams, SharePoint and OneDrive. An example for a use case is a R&D department which may not communicate with the marketing department. The example used in this chapter is a multi-national company where users in the Netherlands are not allowed to communicate with users in Belgium because of high regulatory laws and requirements.

In this chapter, I will introduce you to Information Barriers in Microsoft Purview by walking you through setting up the backend and show you the impact it has on your users. Let's go!

Setting the scene

For this demonstration, we are going to use the following employees in our fictional company:

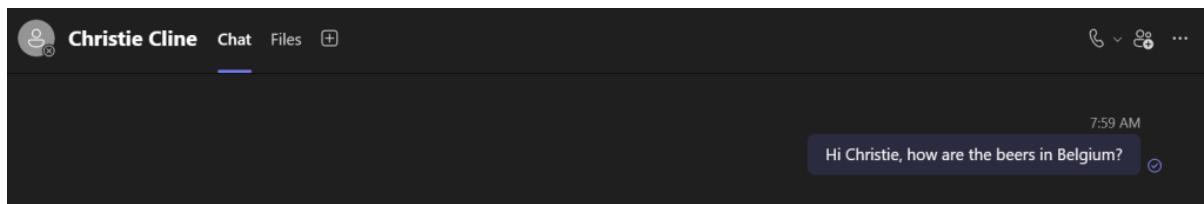
1. Allan Deyoung and Adele Vance, working in the Netherlands Office.
2. Christie Cline and Megan Bowen, working in the Belgium Office.

When you want to implement Information Barriers in your environment, make sure to take your time planning the solution. First become familiar with all the possibilities, next make a design and finish by implementing it. Configuring Information Barriers may seem simple, but a lot of actions take time to take effect in your environment and behavior can be a bit flaky when you don't wait for 24 hours after implementing your solution.

Information Barriers versioning

Information Barriers are available for quite some time now. However, it is only recently that Microsoft introduced some new features into Information Barriers and named this version "Information Barriers v2". v2 was launched on March 6th 2023. So, If your tenant is created after this date, you have the option to use the v2 features. If your tenant is created before this date, you will have to wait for the option to migrate to v2, which is currently scheduled for July 2024. All information in this chapter is based on v1.

Pre-Information Barriers behavior



1 on 1 chat between Allan Deyoung and Christie Cline

PNW Coffee kick-off

Hi, PNW Coffee Social Campaign team, Awesome job all around getting things organized with our client this morning. I really love the participation, questions and collaboration coming from this group. The client expressed how excited they are to work with us. I created this channel to keep us on task. I've created folders for this channel so we can organize and collect assets here.

Megan Bowen 2/12 11:25 AM
Thank you Christie Cline for coordinating!

Johanna Lorenz 2/12 11:25 AM
Agreed! Christie, we should probably add Diego and Miriam from the comms team to this chat.

Christie Cline 2/12 11:25 AM
Agreed.

Irvin Sayers 2/12 11:25 AM
I look forward to our brainstorm session this afternoon. Chat soon!

Allan Deyoung 8:00 AM
This is an awesome reply to an existing message of Christie.

Reply by Allan Deyoung on a Teams-post by Christie Cline

When there are no information barriers in place, users may communicate freely with other users in your Microsoft 365 tenant. In the example above, you can see that user Allan Deyoung can communicate with Christie Cline using Teams chat and reply to a post made by Christie in the “PNW Coffee kick-off” Team.

Setting up Information Barrier Basics – Segments

First, we are going to create segments. Segments are used to create a group of users based on a common property. In this example, I will use the property “Usage Location”. Other properties that can be used are for example Department, Street etc.

In the Microsoft Purview portal, navigate to Information Barriers, Segments and click “New Segment”.

Contoso Electronics Microsoft Purview

Information barrier > Segment > Create segment

Name
Employees in the Netherlands

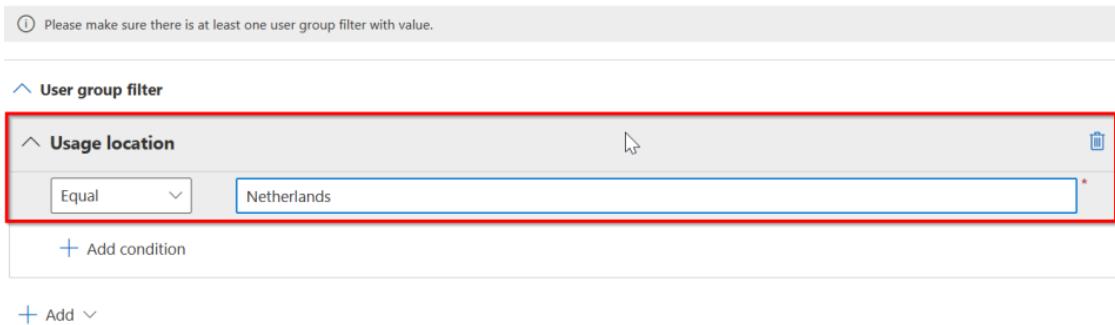
User group filter

Review your settings

Provide a segment name

First provide a segment name, for which I'll use “Employees in the Netherlands”

Add user group filter



Please make sure there is at least one user group filter with value.

User group filter

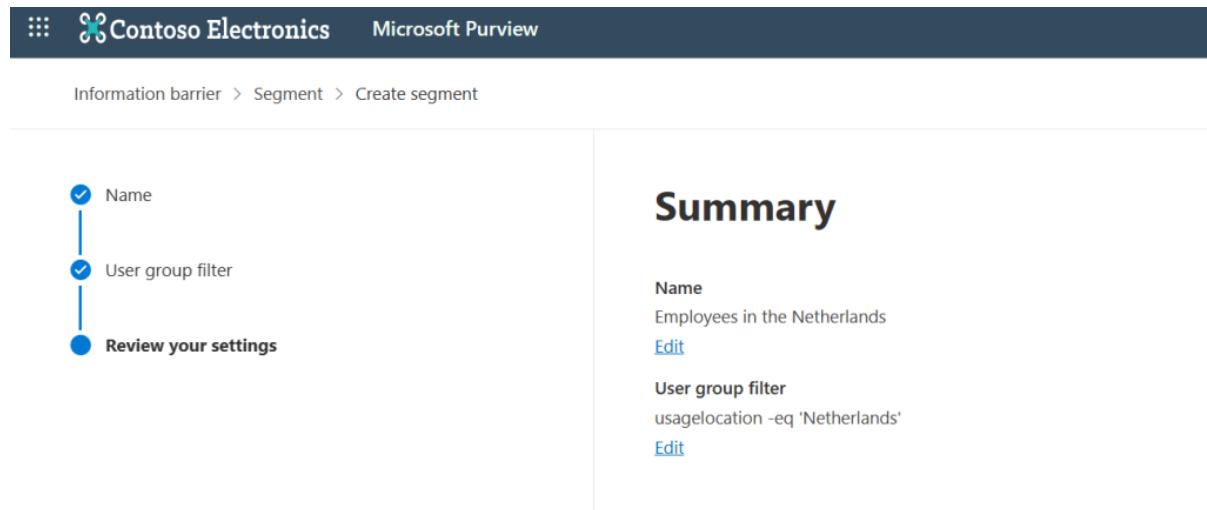
Usage location

Equal Netherlands

Add condition

Add

Next, create the group filter. Here I'll use "Usage location" Equals "Netherlands". This will form a group of all users in my tenant which have property usage location set to Netherlands.



Contoso Electronics Microsoft Purview

Information barrier > Segment > Create segment

Name: Employees in the Netherlands

User group filter: usagelocation -eq 'Netherlands'

Review your settings

Summary

Name: Employees in the Netherlands

User group filter: usagelocation -eq 'Netherlands'

Review the summary and finish the wizard.

Segments

In addition to your initial list of policies, make a list of segments for your organization.

Segments		
	Last modified by	Last modified
<input type="checkbox"/> Name		
<input type="checkbox"/> Employees in Belgium	MOD Administrator	Mar 4, 2024 7:19 AM
<input type="checkbox"/> Employees in the Netherlands	MOD Administrator	Mar 4, 2024 7:18 AM

Create another segment using the same steps as before, however here I've changed the usage location to "Belgium". Now I have 2 segments configured.

Note that information barrier segments have to be designed and configured in a way that all users are a member of only 1 segment. This limitation is removed in Information Barriers v2, but if you're still using v1 and you're adding users to multiple segments, applying your policies will fail.

Setting up Information Barrier Basics – Policies

Next, let's configure policies. Policies are the glue between segments as you can choose to block or allow communication between segments with them. Under the information barriers section, click "Policies" and select "Create Policy".

First, create a name for the policy. I chose to show the names of the segments that I'm going to use and the action (block).

Next, assign the first segment, which is "Employees in the Netherlands" in this case.

Configure communication and collaboration details

Note: Communication over Teams and collaboration on SharePoint & OneDrive would be restricted based on this policy.

On the next page, select the action "Blocked" and choose the second segment which is "Employees in Belgium" in my case. Set the policy status to Active in the next step.

Information barrier > Policies > Create policy

Name

Assigned segment

Communication and collaboration

Policy status

Review your settings

Summary

Name
Usage Location Netherlands - Block - Usage Location Belgium
[Edit](#)

Assigned segment
Employees in the Netherlands
[Edit](#)

Blocked segments
Employees in Belgium
[Edit](#)

Policy status
Active
[Edit](#)

Review the summary and finish the wizard. Now the first policy is created. Now here comes the first catch. To be able to apply this policy, you should make the same policy in reverse. What I mean by this is that you also have to create a policy where users with usage location Belgium should block users with usage location Netherlands. This has to be done because 1-way communication blocks are not supported.

Information barrier > Policies > Create policy

Name

Assigned segment

Communication and collaboration

Policy status

Review your settings

Summary

Name
Usage Location Belgium - Block - Usage Location Netherlands
[Edit](#)

Assigned segment
Employees in Belgium
[Edit](#)

Blocked segments
Employees in the Netherlands
[Edit](#)

Policy status
Active
[Edit](#)

Here you see the second policy I've created.

Policies

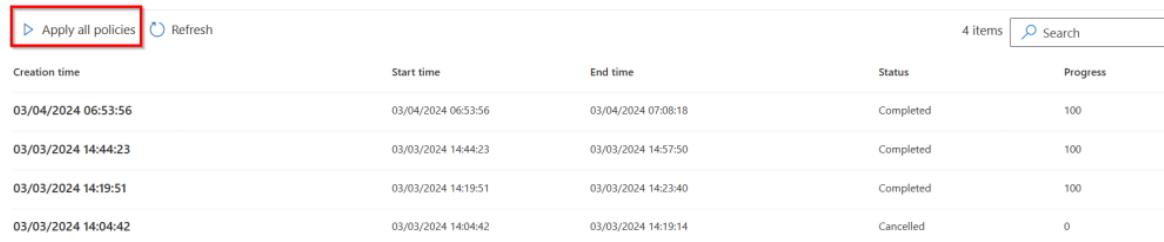
<input type="checkbox"/> Name	Last modified by	Last modified	Status
<input type="checkbox"/> Usage Location Belgium - Block - Usage Location Netherlands	MOD Administrator	Mar 4, 2024 7:24 AM	Active
<input type="checkbox"/> Usage Location Netherlands - Block - Usage Location Belgium	MOD Administrator	Mar 4, 2024 7:23 AM	Active

And there we are, policy creation finished! ✓

Setting Up Information Barrier Basics – Applying Policies

Now for the last step, we have to apply the policies we made.

Policy application



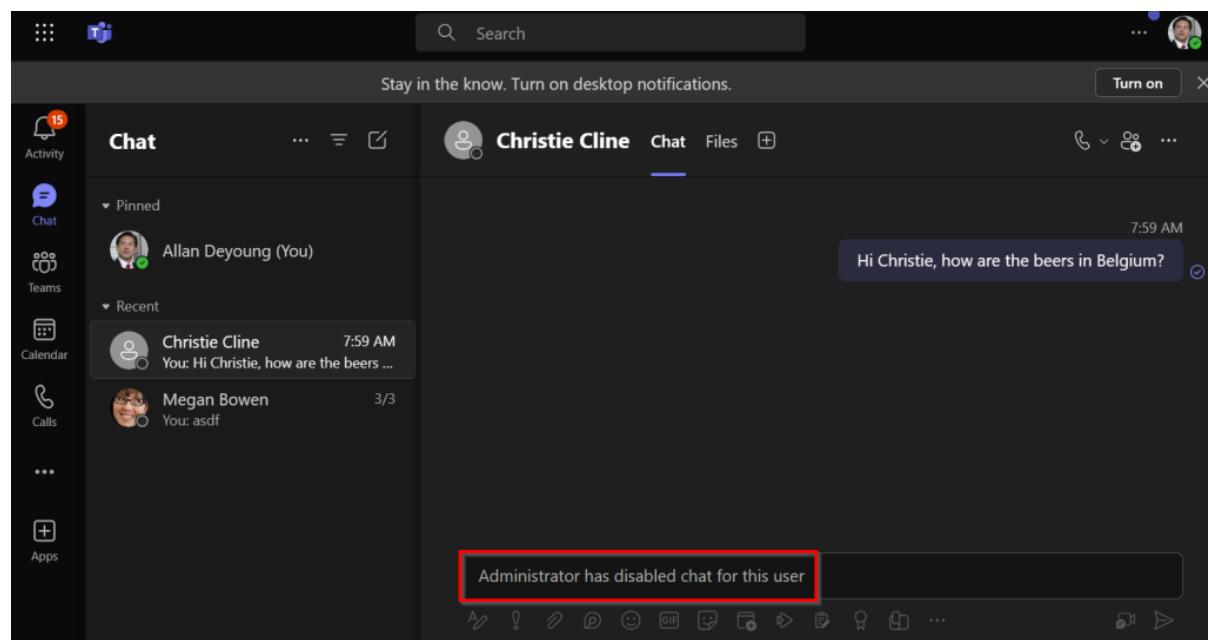
Creation time	Start time	End time	Status	Progress
03/04/2024 06:53:56	03/04/2024 06:53:56	03/04/2024 07:08:18	Completed	100
03/03/2024 14:44:23	03/03/2024 14:44:23	03/03/2024 14:57:50	Completed	100
03/03/2024 14:19:51	03/03/2024 14:19:51	03/03/2024 14:23:40	Completed	100
03/03/2024 14:04:42	03/03/2024 14:04:42	03/03/2024 14:19:14	Cancelled	0

Navigate to “Policy application” under the Information Barriers section and press “Apply all policies”. The status will cycle through the phases NotStarted, ApplyInProgress, PendingCompletion and Completed. I would advise you to take your time and wait at least 24 hours after applying your policies to make sure they are distributed through your entire environment as otherwise your experience will not be consistent through the entire environment.

A quick note on editing and changing segments and policies. To put it simple, this cannot be done. When you want to remove a policy, keep in mind that you have to edit it first and set “Active” to off. If you want to remove a segment, first edit it’s filter so that it doesn’t include any users. If this doesn’t work, apply the policies again or wait a while longer.

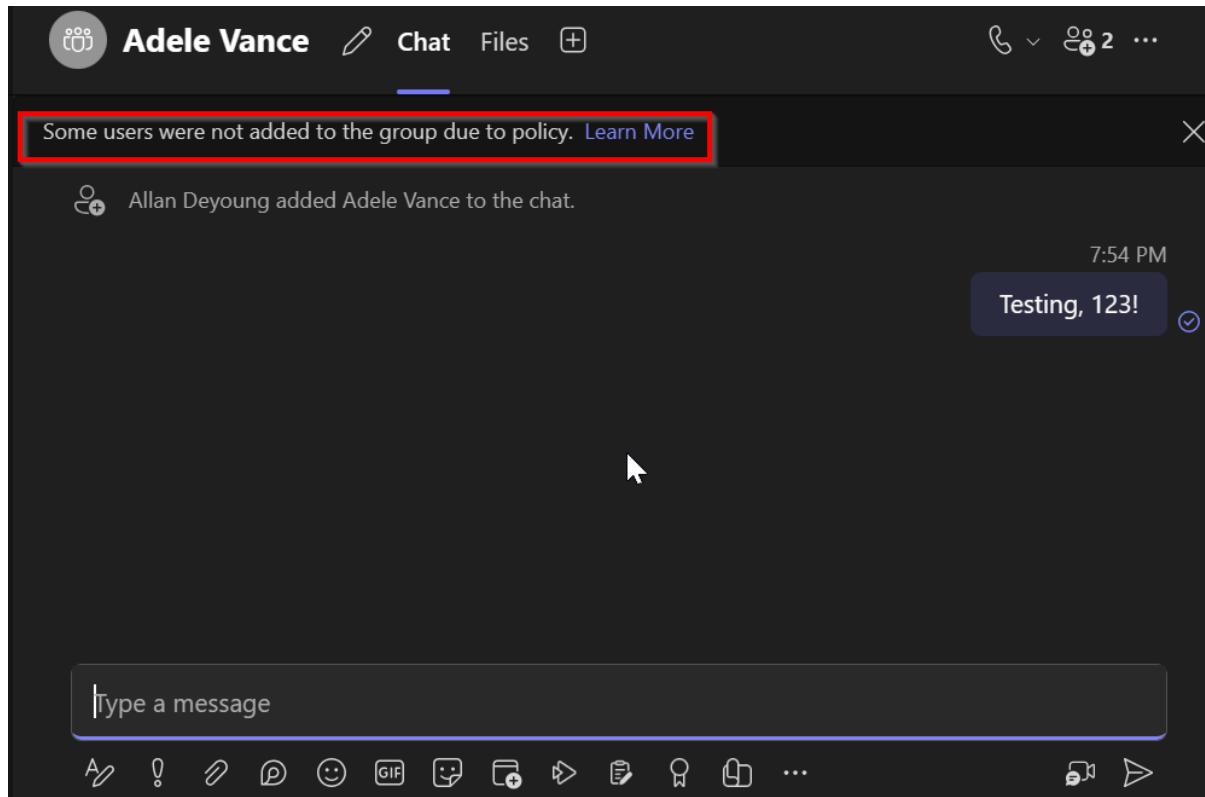
The user Experience – Teams

When taking a look at the user experience for teams I’ve created some examples. Let’s see what happens when Allen from the previous example wants to continue his 1 on 1 chat with Christie.



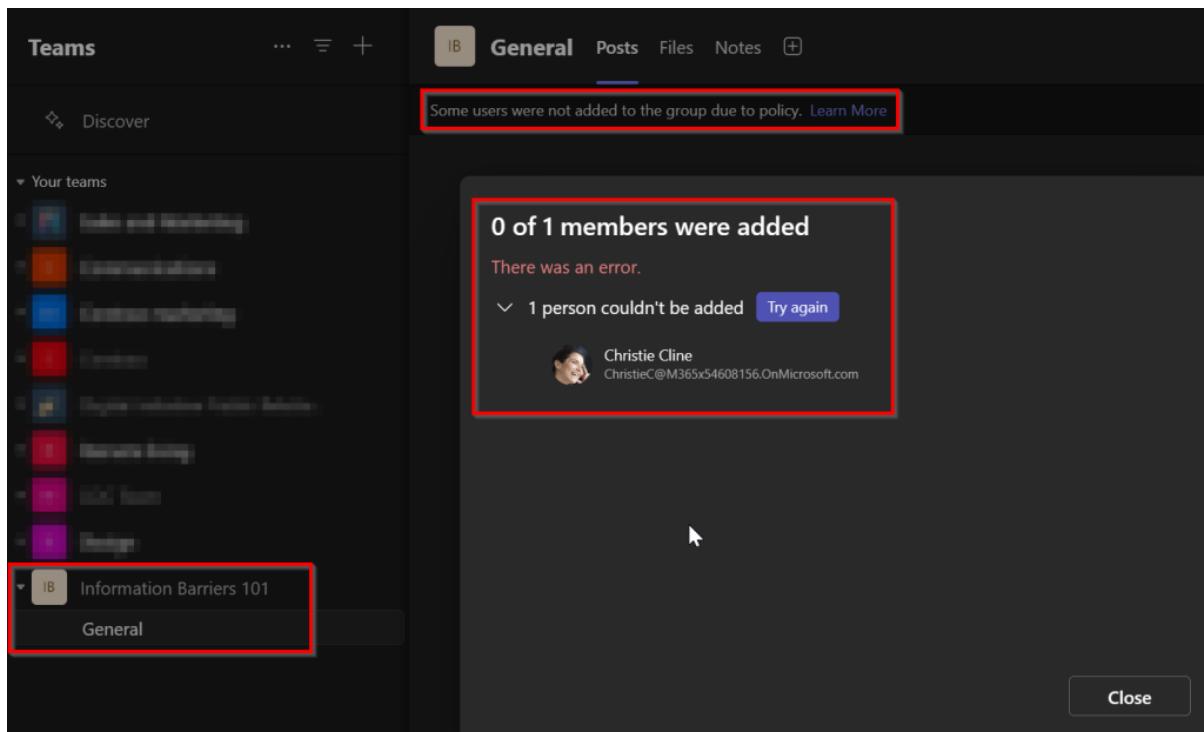
We receive a message that an “Administrator has disabled chat for this user”. This is ofcourse because both users are on different segments and we’ve blocked communication between these segments. Of course this works both ways, Christie is also not able to communicate with Allan anymore and Teams is displaying the same message in her Teams client.

Let’s take a look at group chats. In this example Allan Deyoung creates a group chat with Adele Vance and Christie Cline.



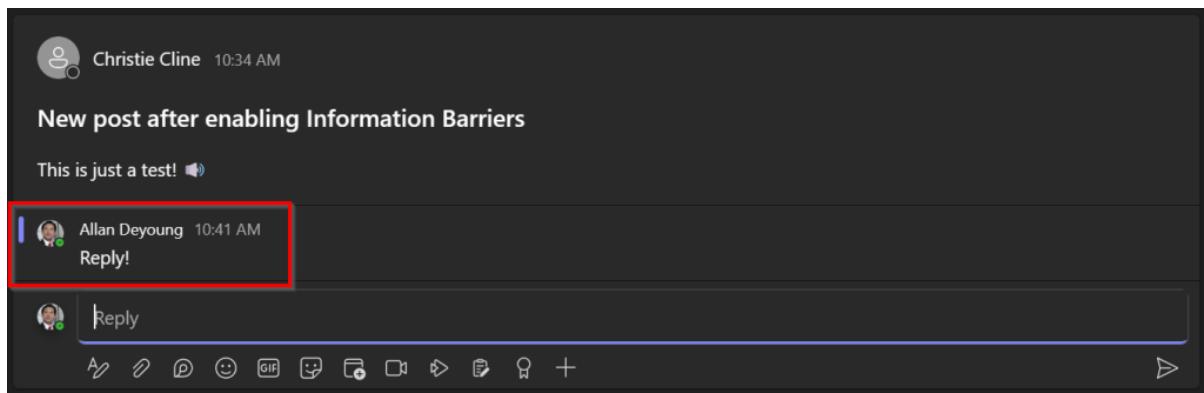
Allan is immediately notified that some users were not added to the group due to policy. In this case this would be Christie, since she's in another segment to where communication is blocked.

Now onto Teams. Specifically the creation of a team. In this example Allen Deyoung create a public team and adds Christie to team.



Can you guess the outcome? Correct, Christie could not be added!

However, the behavior is different when the team is already created.



Even when Christie creates a new post in an existing team, Allan can still post a comment!

These are just some examples of what Information Barriers can do for you in Teams. Take a look at this [Microsoft Learn](https://learn.microsoft.com/en-us/purview/information-barriers-teams) (<https://learn.microsoft.com/en-us/purview/information-barriers-teams>) article for more!

Setting up Information Barriers – Sharepoint Configuration

Microsoft Teams setup was reasonably quick and easy. SharePoint however needs some more work to get Information Barriers working.

For SharePoint (and also OneDrive) first install the SharePoint Online Management Shell by starting PowerShell 5 (included with Windows 11, unfortunately PowerShell 7 is not supported by the module at time of writing) and enter the following command:

```
Install-Module -Name Microsoft.Online.SharePoint.PowerShell
```

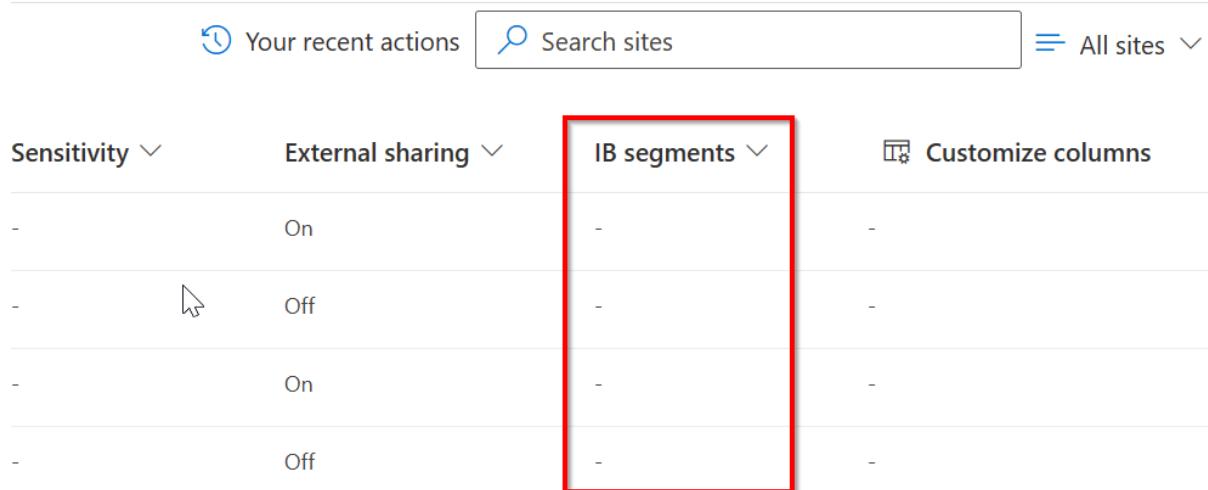
When this is done load the module and connect to your SharePoint environment using the following cmdlets:

```
Import-Module Microsoft.Online.SharePoint.PowerShell
```

```
Connect-SPOService -Url https://YourSharePointURL-admin.sharepoint.com
```

Next, run the following cmdlet to lift the suspension from Information Barriers in SharePoint which enables it:

```
Set-SPOTenant -InformationBarriersSuspension $false
```



Sensitivity	External sharing	IB segments	Customize columns
-	On	-	-
-	Off	-	-
-	On	-	-
-	Off	-	-

Now, Navigate to the SharePoint Online Admin Portal and click Active Sites. You should see a column named “IB Segments” as in the screenshot above. Information Barrier Segments are now enabled for your SharePoint sites.

Before configuring the site, you should know that a SharePoint site can be configured in 1 of 4 Information Barrier Modes:

Mode	Description	Examples
Open	When a SharePoint site doesn't have segments, the site's IB mode is automatically set as <i>Open</i> . See this section for details on managing segments with the <i>Open</i> mode configuration.	A Team site created for picnic event for your organization.
Owner Moderated	When a SharePoint site is created for collaboration between incompatible segments moderated by the site owner, the site's IB mode should be set as <i>Owner Moderated</i> . See this section for details on managing <i>Owner Moderated</i> site.	A site is created for collaboration between VP of Sales and Research in the presence of VP of HR (site owner).
Implicit	When a site is provisioned by Microsoft Teams, the site's IB mode is set as <i>Implicit</i> by default. A SharePoint Administrator or Global Administrator can't manage segments with the <i>Implicit</i> mode configuration.	A Team is created for all Sales segment users to collaborate with each other.
Explicit	When segment is added to a SharePoint site either via end-user site creation experience or by a SharePoint Administrator adding segment to a site, the site's IB mode is set as <i>Explicit</i> . See this section for details on managing segments with the <i>Explicit</i> mode configuration.	A research site is created for Research segment users.

Source: [Microsoft](#)

In this example, our site is configured with the IB Mode “Explicit”. You can always check the current mode of your site using the following cmdlet:

```
Get-SPOSite -Identity YourSiteURL | Select InformationBarriersMode
```

Now on to the configuration of our SharePoint site:

Benefits

Communication site

View site Delete

General Activity Membership **Settings**

External file sharing

Only people in your organization

Sensitivity label

None

[More sharing settings](#)

Information barriers

Employees in Belgium

[Edit](#)

Navigate to one of your sites and click on the “Settings” tab. Add one of your created segments to your SharePoint site in the section shown above. By adding the site to a Information Barriers segment it can now be used by your Information Barrier policies. Let’s take a look at an example what this means for a user.

Access Denied

Due to organizational policies, you can't access this resource.

Here are a few ideas:

Please contact your organization.

If this problem persists, contact your support team and include these technical details:

Correlation ID: [REDACTED]

Date and Time: 3/4/2024 11:48:02 AM

User: [REDACTED]

Issue Type: User has encountered a policy issue.

In the configuration above, I’ve added the SharePoint site “Benefits” to the “Employees in Belgium” segment. Now I can still add Allan to the site members, but since he hits the information barrier, he can’t access the site as can be seen in the screenshot above, since he’s not in the Belgium usage location.

Again, this is just 1 example. Take a look at this [Microsoft Learn article](#) (<https://learn.microsoft.com/en-us/purview/information-barriers-sharepoint>) to show all usecases.

Setting up Information Barriers – OneDrive Configuration

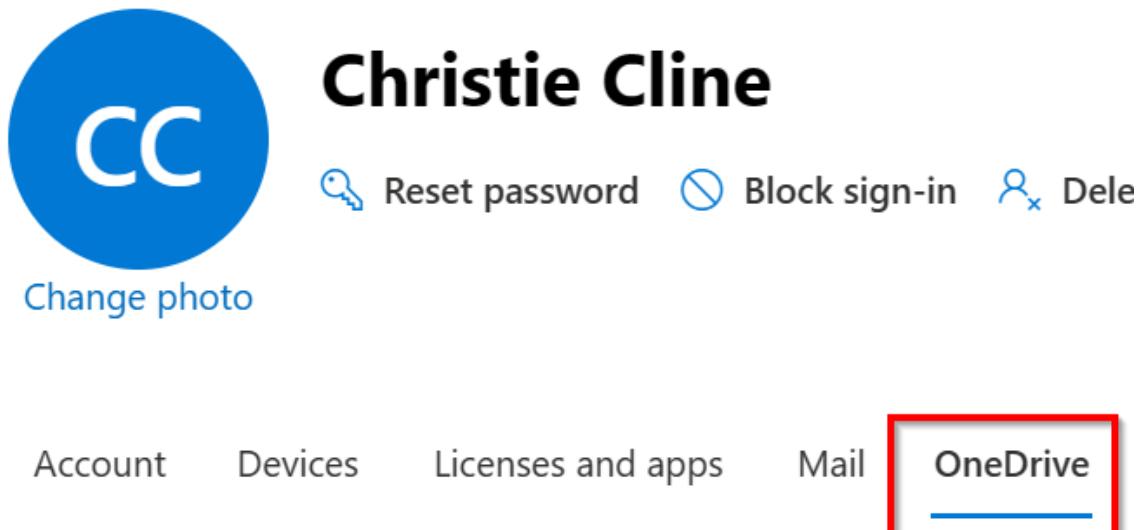
When we enabled Information Barriers for our SharePoint environment, we also enabled it for our OneDrive environment. So that's a nice advantage. As with SharePoint, OneDrive can also be configured with different Information Barrier Modes:

Mode	Description
Open	When a non-segmented user provisions their OneDrive, the site's IB mode is set as Open, by default. There are no segments associated with the site.
Owner Moderated	When a OneDrive is used for collaboration with incompatible users in the presence of the site owner/moderator, the OneDrive's IB mode can be set as Owner Moderated. See this section for details on Owner Moderated site.
Explicit	When a segmented user provisions their OneDrive within 24 hours of enablement, the site's IB mode is set as <i>Explicit</i> by default. The user's segment and other segments that are compatible with the user's segment and with each other get associated with the user's OneDrive.
Mixed	When a segmented user's OneDrive is allowed to be shared with unsegmented users, the site's IB mode can be set as <i>Mixed</i> . This is an opt-in mode that the SharePoint admin can set on OneDrive of a segmented user.

Source: [Microsoft](#)

The difference with OneDrive is however, that when a user is added to an Information Barrier segment, this segment is stamped on it's OneDrive and the Information Barrier mode of the users OneDrive is set to "Explicit", both within 24 hours of setting the segment.

The mode can be seen using the "Microsoft.Online.SharePoint.PowerShell" module that's used for the SharePoint configuration above. You can return to the same PowerShell window to execute the cmdlets. However, first you have to find the URL of the users OneDrive.



The screenshot shows the Microsoft 365 Admin Portal with the 'Users' section selected. Under 'Active users', a user named 'Christie Cline' is listed. The 'OneDrive' tab is highlighted with a red box. The navigation bar at the top includes 'Account', 'Devices', 'Licenses and apps', 'Mail', and 'OneDrive'.

Get access to files

Create a link to view and edit Christie Cline's OneDrive files.

[Create link to files](#)

This can be found by navigating to the Microsoft 365 Admin Portal, Users, Active Users. Then select your user and click on the Onedrive tab. Then click “create link to files” and copy the link that is shown.

To see the mode, run the following cmdlet in the PowerShell window:

```
Get-SPOSite -Identity https://TenantName-
my.sharepoint.com/personal/christiec_TenantName_onmicrosoft_com | Select
InformationBarriersMode
```

```
PS C:\Users\DominiqueHermans> Connect-SPOSERVICE -Url https://[REDACTED]-admin.sharepoint.com
PS C:\Users\DominiqueHermans> Get-SPOSITE -Identity https://[REDACTED]-my.sharepoint.com/personal/christiec_[REDACTED]_onmicrosoft_com | Select InformationBarriersMode
InformationBarriersMode
-----
Explicit
```

As can be seen in the screenshot above, the mode is indeed set to “Explicit”. Explicit means the following according to [Microsoft Learn](https://learn.microsoft.com/en-us/purview/information-barriers-onedrive#explicit) (<https://learn.microsoft.com/en-us/purview/information-barriers-onedrive#explicit>):

Explicit

When a OneDrive has information barriers segments and the mode is set to *Explicit*:

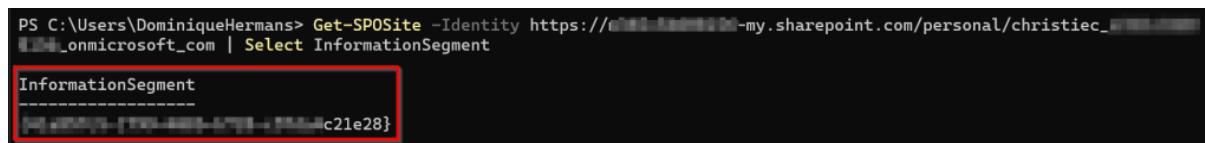
- The option to share with *Anyone with the link* is disabled.
- The option to share with *Company-wide link* is disabled.
- Files and folders can be shared only with users whose segment matches that of the OneDrive.

To change the mode, use:

```
Set-SPOSite -Identity https://TenantName-my.sharepoint.com/personal/christiec_TenantName_onmicrosoft_com -InformationBarriersMode OwnerModerated
```

or to show which segment was assigned, use:

```
Get-SPOSite -Identity https://TenantName-my.sharepoint.com/personal/christiec_TenantName_onmicrosoft_com | Select InformationSegment
```



```
PS C:\Users\DominiqueHermans> Get-SPOSite -Identity https://... | Select InformationSegment
InformationSegment
-----
c21e28
```

As shown in the screenshot above, the result returned is not the name of a segment, but a GUID. This GUID should be matched with the GUID that's a result of the "Get-OrganizationSegment" cmdlet.

But this is in another module. And to be precise the "Exchange Online PowerShell module". Weird name? Yes. But it gets the job done. And this module is supported by Powershell 7.0.3 and later. Install with command

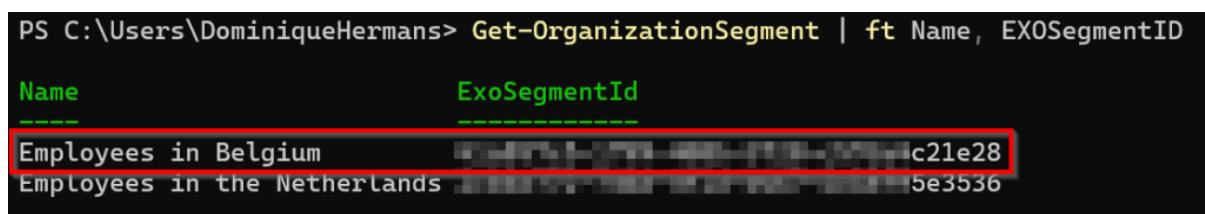
```
Install-Module -Name ExchangeOnlineManagement
```

Take a look at it's [Microsoft Learn page \(https://learn.microsoft.com/en-us/powershell/exchange/exchange-online-powershell-v2?view=exchange-ps#install-and-maintain-the-exchange-online-powershell-module\)](https://learn.microsoft.com/en-us/powershell/exchange/exchange-online-powershell-v2?view=exchange-ps#install-and-maintain-the-exchange-online-powershell-module) for all the prerequisites.

Run the following commands to match the GUID above with a meaningful name:

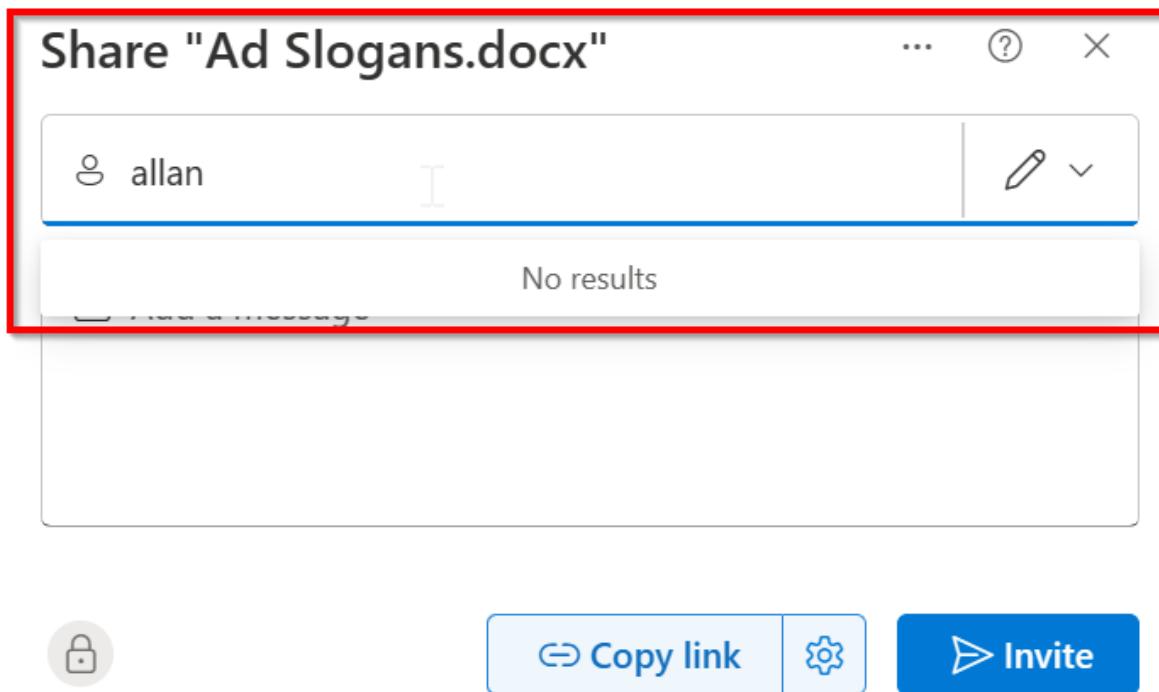
```
Import-Module ExchangeOnlineManagement
Connect-IPSSession -UserPrincipalName admin@TenantName.onmicrosoft.com
Get-OrganizationSegment | ft Name, EXOSegmentID
```

The result is as follows:

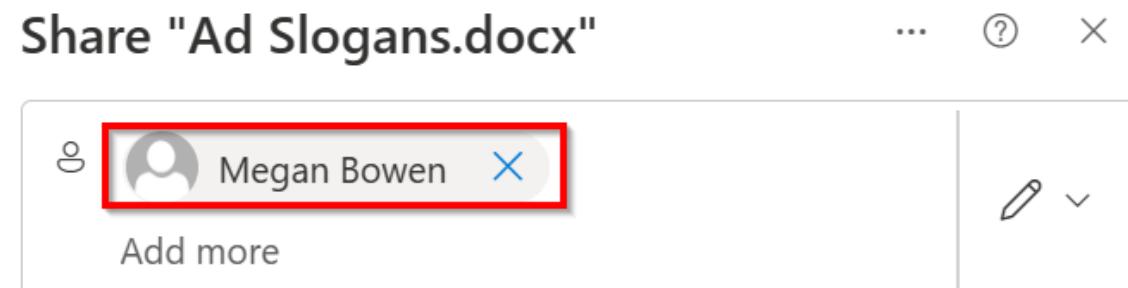


Name	EXOSegmentID
Employees in Belgium	c21e28
Employees in the Netherlands	5e3536

And this checks out, since it's Christie's OneDrive and Christie is an employee from Belgium! To wrap this up, let's take a look at the user experience.



When Christie wants to share a file in her OneDrive with Allan (which is in another segment) this isn't possible. Allan doesn't even show up in the list of available users!



However, when sharing the same document with Megan Bowen, who is also a user from Belgium, this is possible and Megan can be found in the recipients list.

As with the Teams and SharePoint examples, this OneDrive scenario was just 1 example. Take a look at [Microsoft Learn](https://learn.microsoft.com/en-us/purview/information-barriers-onedrive) (<https://learn.microsoft.com/en-us/purview/information-barriers-onedrive>) to see other possibilities!

Compliance Manager

Microsoft Purview has a pretty sweet feature called “Compliance Manager”. It can be used to assess your Microsoft 365 (and other non-Microsoft) environments based on various regulations like “ISO 27001” or you can create your own custom assessment that’s not based on a regulation. (Do note that at the time of writing this chapter, the creation of custom assessments [is disabled](#) due to an update of the process by Microsoft.)

Basics first, as always

Before we dive into the world of assessments and regulations, let’s start with the basic components of compliance manager:

- A **control** can be a technical setting in your Microsoft 365 environment or a procedure that has to be followed. Examples are “turn on MFA” or “create a document with rules that employees have to accept before they can access their new workspace”.
- An **assessment** is a group of controls.
- Assessments can be based on a **regulation**, which groups all the controls that are in the scope of a regulation in 1 assessment. When you comply to this assessment, it can be stated that you are compliant to the regulation it is based on.

These 3 components form the basics of your compliance solution. Other components include:

- A **solution** is a service within the Microsoft 365 or Azure ecosystem that can be checked based on controls. For example “Exchange Online”.
- Assessments can be **grouped** together. A group of assessments can share the same improvement actions.
- An **improvement action** can be a change you make to improve your system, like “turn on MFA”.

Compliance Manager, here we come!

Now that’s out of the way let’s get our hands dirty and dive into Compliance Manager!

When you navigate to the [Microsoft 365 Compliance center](#) and make your way to the compliance manager, you are immediately greeted with the following page:

Compliance Manager



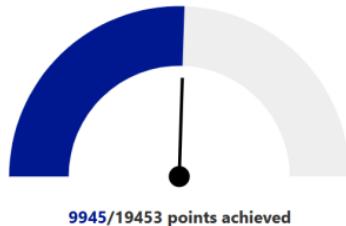
Overview Improvement actions Solutions Assessments Regulations Alerts Alert policies

Compliance Manager measures your progress in completing actions that help reduce risks around data protection and regulatory standards. [Find guidance and documentation](#)

Filter

Overall compliance score

Your compliance score: 51%



Key improvement actions

Not completed	Completed	Out of scope
485	4	0

Improvement action	Impact	Test status	Group	Action type
Use boundary protection to limit access to...	+27 points	None	Default Group	Technical
Configure VPN integration	+27 points	None	Default Group	Technical
Digitally encrypt secure channel data	+27 points	None	Default Group	Technical
Automate notifications and configure act...	+27 points	None	Default Group	Technical
Block outdated ActiveX controls	+27 points	None	Default Group	Technical

It shows a cool gauge that tells you your current compliance score. But wait. Compliance score? Based on what assessment? We didn't even create an assessment! That's right, but when we move forward to the "Assessments" tab, we can see that it is set up with a predefined assessment:

Assessment	Status	Progress	Your improvement ...	Microsoft actions	Group	Service	Regulation
<input type="checkbox"/> Data Protection Baseline for Micr...	Incomplete	51%	4 of 489 completed	591 of 606 completed	Default Group	Microsoft 365	Data Protection Baseline

It also shows that it is based on the "Data Protection Baseline" regulation and that's currently scanning our environment (progress at 51% in the screenshot above) to check all the controls in the assessment. This assessment checks your environment on a set of controls for key regulations and standards for data protection and general data governance. You should check out [this Microsoft Learn article](https://learn.microsoft.com/en-us/purview/compliance-manager-assessments#data-protection-baseline-default-assessment) (<https://learn.microsoft.com/en-us/purview/compliance-manager-assessments#data-protection-baseline-default-assessment>) if you would like to know more about the Data Protection Baseline assessment.

How to create an assessment based on a regulation

If you want, start by glancing at the "Regulations" tab on the top of the page, it will show you a dazzling number (376 at the time of writing) of regulations that you can base your assessment on. Now let's move forward by creating an assessment of our own.

Add assessment	Export actions	Update actions	1 item	Search	Group
Filters:					
Assessment	Status	Progress	Your improvement ...	Microsoft actions	Group
<input type="checkbox"/> Data Protection Baseline for Micr...	Incomplete	51%	5 of 489 completed	591 of 606 completed	Default Group
					Microsoft 365
					Data Protection Baseline Admin

1. Click the "Assessments" tab and click the "Add Assessment" button next. You will be presented by a nice wizard. Click "Select Regulation" to select a regulation to base your assessment on.

Base your assessment on a regulation

[Change selection *](#)

Regulation name
ISO/IEC 27001:2022 

Availability
Premium

Activation
Inactive

Created by
Microsoft

Last modified
10/18/2023

Date created
7/11/2023

- I chose the “ISO 27001 version 2022” regulation since that’s a well known regulation in the Netherlands. Click next.

Add name and group

Create a name for your assessment and assign it to an existing or new group. The assessment name must be unique within the group. Group names must be unique within your organization. [Learn more about groups](#)

Assessment name *



Assessment group *

Use existing group 

Create new group 



- Give your assessment a name. Remember that assessments can be grouped together? Here the wizard provides you with this possibility. I choose to go with the default group for this demonstration. Click next.

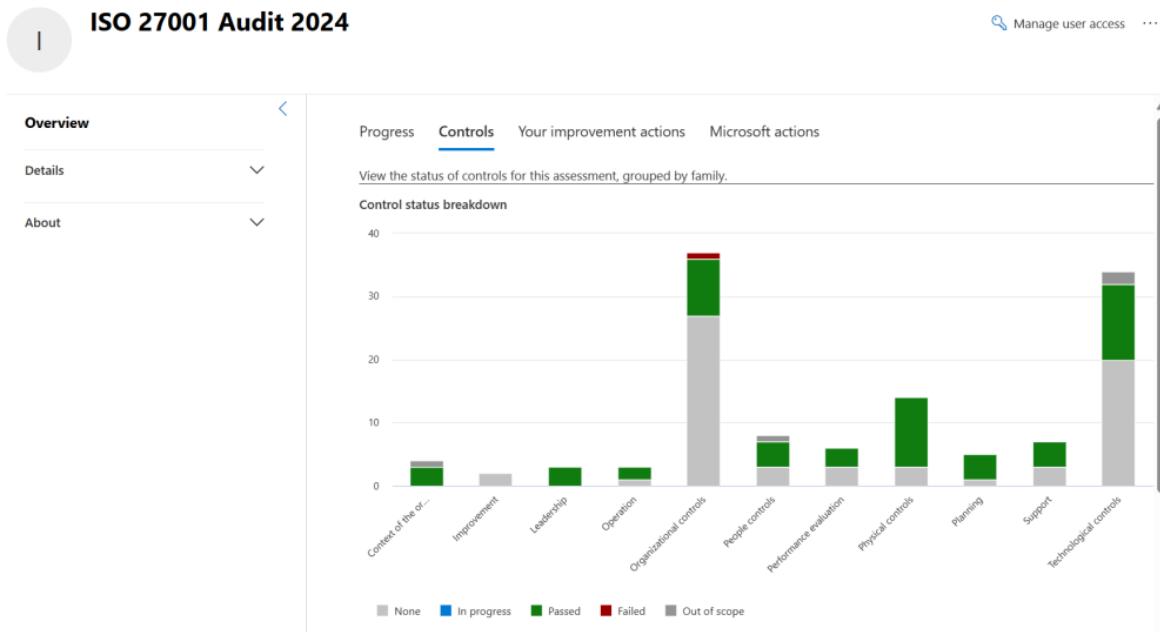
Select services

Select the services this assessment will apply to. [Learn more about services](#)

 Select services  Add new service 1 item

Service
Microsoft 365 

- By default, you can use your assessment to assess your Microsoft 365 environment. However, it is possible to extend this to Zoom or Salesforce if you want. Click next and review your selections. Click “Create Assessment” when you’re done.



When your assessment is saved it starts checking your environment based on the controls I mentioned earlier. When you take a look at the “Controls” tab, you can see exactly what controls are being checked and what controls don’t match (or fail if you will) with the properties of the regulation you selected.

Progress Controls **Your improvement actions** Microsoft actions

Review improvement actions managed by your organization. Select an improvement action to edit its status and view implementation guidance

Improvement action status

Passed Failed Low Risk Failed Medium Risk Failed High Risk Not Assessed Partially Tested 6 more

✓ Accept all updates Assign to user 110 items Search

Filter Reset Filters

Service: Any Control family: Any Status: Any Service Instances: Any Testing type: Any +1 more

Improvement action	Service	Test status	Impact	Points achieved	Regulat
Add an on-premises application ...	Microsoft 365	None	+9 points	0/9	Data Pr...
Add or update users profile infor...	Microsoft 365	None	+27 points	0/27	ISO/IEC

Remember that grouped assessments can share the same improvement actions? Also remember that I choose to put my ISO 27001 in the default group? Well, as you can see in the screenshot above, the Data Protection Baseline is also in this group and so the improvement actions are also shared in one view. So when you click the “Your Improvement Actions” tab, it’ll show you the controls you have to change to comply with the regulation you selected.

Review the status of improvement actions managed by Microsoft. Select an improvement action to view details, including implementation and testing notes.

155 items

Filter

Microsoft action	Service	Test status	Points achieved	Regulations
Access Control for Transmission M...	Microsoft 365	Passed	3/3	Data Protection Baseline, IS...
Access Enforcement	Microsoft 365	Passed	27/27	Data Protection Baseline, IS...
Access Restrictions for Change	Microsoft 365	Passed	9/9	Data Protection Baseline, IS...
Account Management - Authorizat...	Microsoft 365	Passed	27/27	Data Protection Baseline, IS...

Since your Microsoft 365 environment is hosted by Microsoft there are also a lot of controls that are ticked off by Microsoft. These can be seen on the tab “Microsoft Actions”.

That's about it there is to tell about the basics of Compliance Manager. Some last tips to help you on your way to compliance managing:

- Your Microsoft 365 license determines how many regulation licenses you get “for free”. If you want to use additional regulation licenses you can buy them separately. You can check the current usage status at the assessments tab, regulation licenses used, view details.
- I would advise you to grant stakeholders permission to this part of the compliance portal by navigating to compliance manager, compliance manager settings in the upper right hand corner, user access.

