

Microsoft trainable classifiers

With Microsoft Purview trainable classifiers, you can automate your data discovery, classification, and governance at scale.





01 /

Introduction

02 /

Use trainable classifiers
to protect data

- » The benefits of trainable classifiers
- » AI classifiers for different
business functions

03 /

Explore Microsoft Purview

- » Know your data
- » Information Protection sensitivity auto-labeling
- » Data Loss Prevention
- » Data Lifecycle Management

04 /

Summary

05 /

Learn more

Introduction

Today, organizations across various industries are generating massive amounts of data, and the volume grows exponentially each year. With the adoption of a cloud and remote work model, data is no longer locked behind your corporate network's perimeters and is instead spread across many nodes.

Recent statistics tell us that 80 percent of business data is dark¹, which means it's unclassified and unprotected, while another study shows that 38 percent of sensitive information incidents involve inadvertent internal misuse² (Figure 1). Because of this, detecting and protecting sensitive data is at the core of our focus.

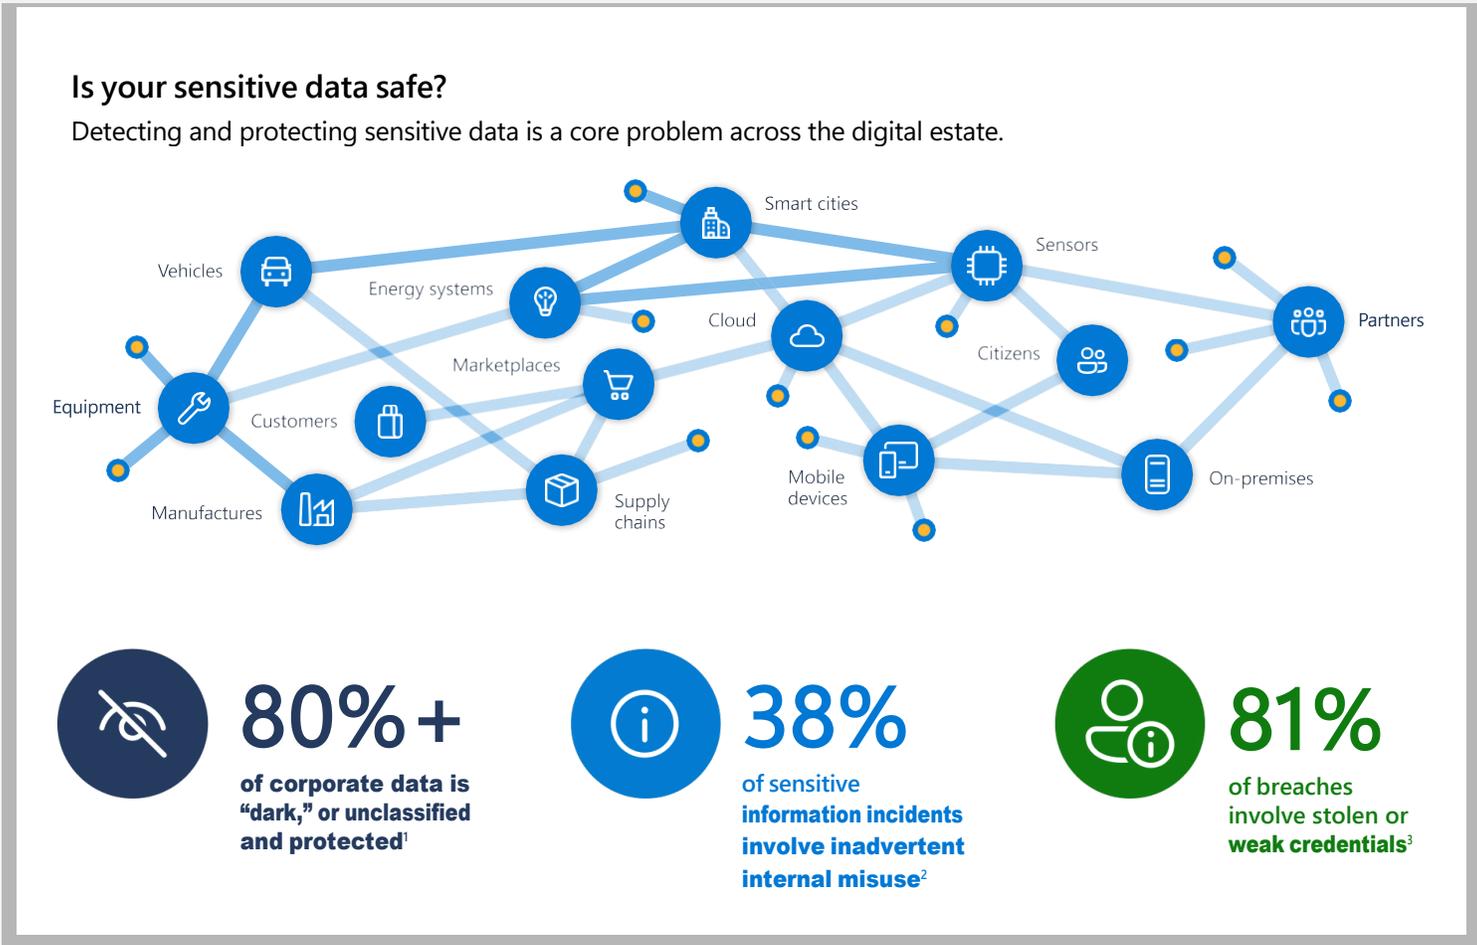


Figure 1: Data across the digital estate

¹ Andrew Trice, "The Future of Cognitive Computing," The IBM Cloud Blog, November 23, 2015.

² Jeff Pollard, "Security Budgets 2019: The Year Of Services Arrives," Forrester Research, December 17, 2018.

³ Clare Ward and Nilesh Pritam, "Cyberespionage and ransomware attacks are on the increase warns the Verizon 2017 Data Breach Investigations Report," Verizon News Center, April 27, 2017.

You need a comprehensive governance plan that can help you decide what business data to protect, retain, or delete, but first you must effectively identify the data. Data classification is the process of organizing data into categories so that it can be protected and handled correctly, acting as the starting point for an information protection discipline.

But no matter how large your workforce, manual and rule-based approaches to data classification cannot work effectively on their own. A better approach is to automate data classification with machine learning technology that can train a model to predict the class of new, unseen data, giving you more efficient data protection and minimizing false positives more efficiently as compared to manual approaches.

What is machine learning?

Machine learning (ML) is the process of applying mathematical concepts to data to help a computer learn. ML uses algorithms to identify patterns within data and attempts to learn these patterns to create a model. Once an ML model has been trained to find these patterns, it can be used to make predictions on unseen data. The more data the model is exposed to during the training process, the better the model will be able to perform, just as humans improve with more practice.

Manual approaches to classification cannot scale to handle the massive data that organizations have across various business functions. Instead, you can automate your data classification with trainable classifiers in Microsoft Purview Information Protection, an artificial intelligence (AI)-based solution that identifies the type of content by analyzing the elements of the content itself. Advanced classification algorithms powered by state-of-the-art intelligence can quickly adapt to changes in regulatory and dynamic business contexts. (Figure 2)

How can we solve this problem with intelligence?



Scale

Manual or rule-based approaches can't effectively work for large volumes of data



Automation

AI and machine learning models automate workflows and greatly improve productivity

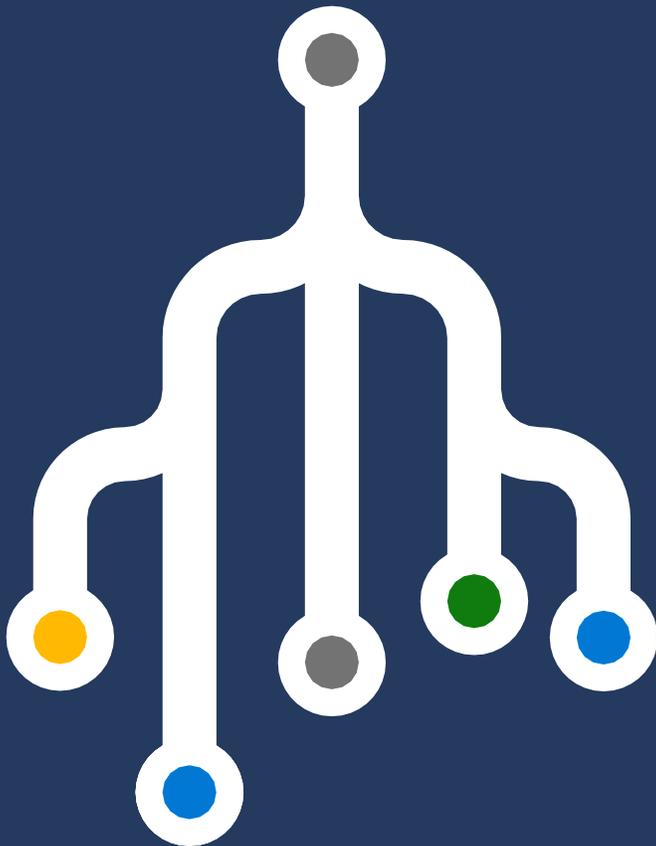


Breadth and coverage

Intelligence can be quickly expanded for growing business contexts and needs

Figure 2: How AI can address business challenges in protecting sensitive data

Use trainable classifiers to protect data



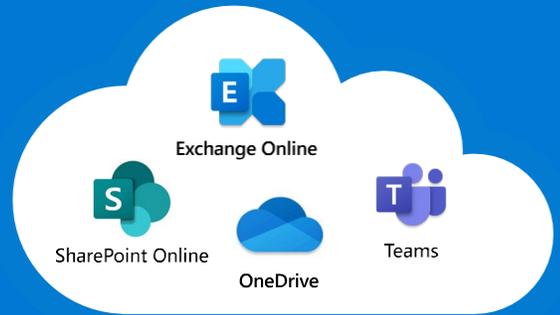
Information Protection simplifies this process with a unified set of capabilities for data classification, labeling, and protection. Our solution addresses information stored in Office apps as well as other popular productivity services where the information resides, such as Microsoft Teams, SharePoint Online, Exchange Online, and endpoint devices. Microsoft is focused on delivering built-in, intelligent, unified, and extensible solutions to protect sensitive data across your digital estate—in Microsoft 365 cloud services, on-premises, in third-party software as a service (SaaS) applications, and more.

Information Protection trainable classifiers are unique compared to other offerings in the market.

- Our out-of-the-box, trainable classifiers for sensitive business document discovery and classification have been trained using a wide sample of real-life data to minimize false positives.
- The ML-enabled trainable classifiers are purpose-built and can quickly and accurately discover and classify sensitive data at scale.
- You can also create customized trainable classifiers to meet your unique content labeling and categorization requirements.
- Document fingerprinting enables systems admins to create a fingerprint sensitive information type (SIT) of a specific document, which can be used later to detect if the same document or part of the same document is found elsewhere in the organization.

How are trainable classifiers uniquely designed for auto-classification?

-  **Breadth and coverage across crucial business functions**
-  **Trained using a large, diverse number of real-world examples**
-  **Crawl and scan data across all workloads**
-  **Out of the box, ready to use, and built in**
-  **Fully integrated with other Microsoft compliance solutions policy authoring**



Microsoft SaaS apps



Endpoint

Figure 3: Microsoft trainable classifiers are optimized for auto-classification.

The benefits of trainable classifiers

Information Protection provides out-of-the-box, ready-to-use classifiers. This means that you can directly use these classifiers in any policy. In addition, you have the ability to build your own custom models to detect proprietary data.

Information Protection trainable classifiers provide the foundation for automating capabilities to detect and classify data effectively. Once data is classified, you can start to govern your data by deciding what to protect, retain, or delete. Our trainable classifiers transform the way your data is classified and categorized.

Scalability

Through our trainable classifiers, you can use the power of machine learning to identify more data categories with increased performance and quickly classify massive volumes of data. All these classifiers have been trained across a large, diverse number of real-world samples and have used some of Microsoft's latest AI technology to build and improve these models.

Breadth and coverage

Significantly improve the speed, performance, and coverage of sensitive data identification at an enterprise scale. We provide coverage for broad common business categories required by global enterprise customers with our ready-to-use optimized classifiers.

Automation

Customers need a solution to work behind the scenes automatically and adapt to ever-changing business needs and regulatory context. Our trainable classifiers can also be used for auto-labeling policies to automatically label and protect sensitive data in key business categories. They're also fully integrated with different Microsoft Purview compliance solutions, such as information protection, data loss prevention, and data lifecycle management, that can help your organization effectively respond to and protect against unauthorized access. (Figure 4)



Benefits of trainable classifiers

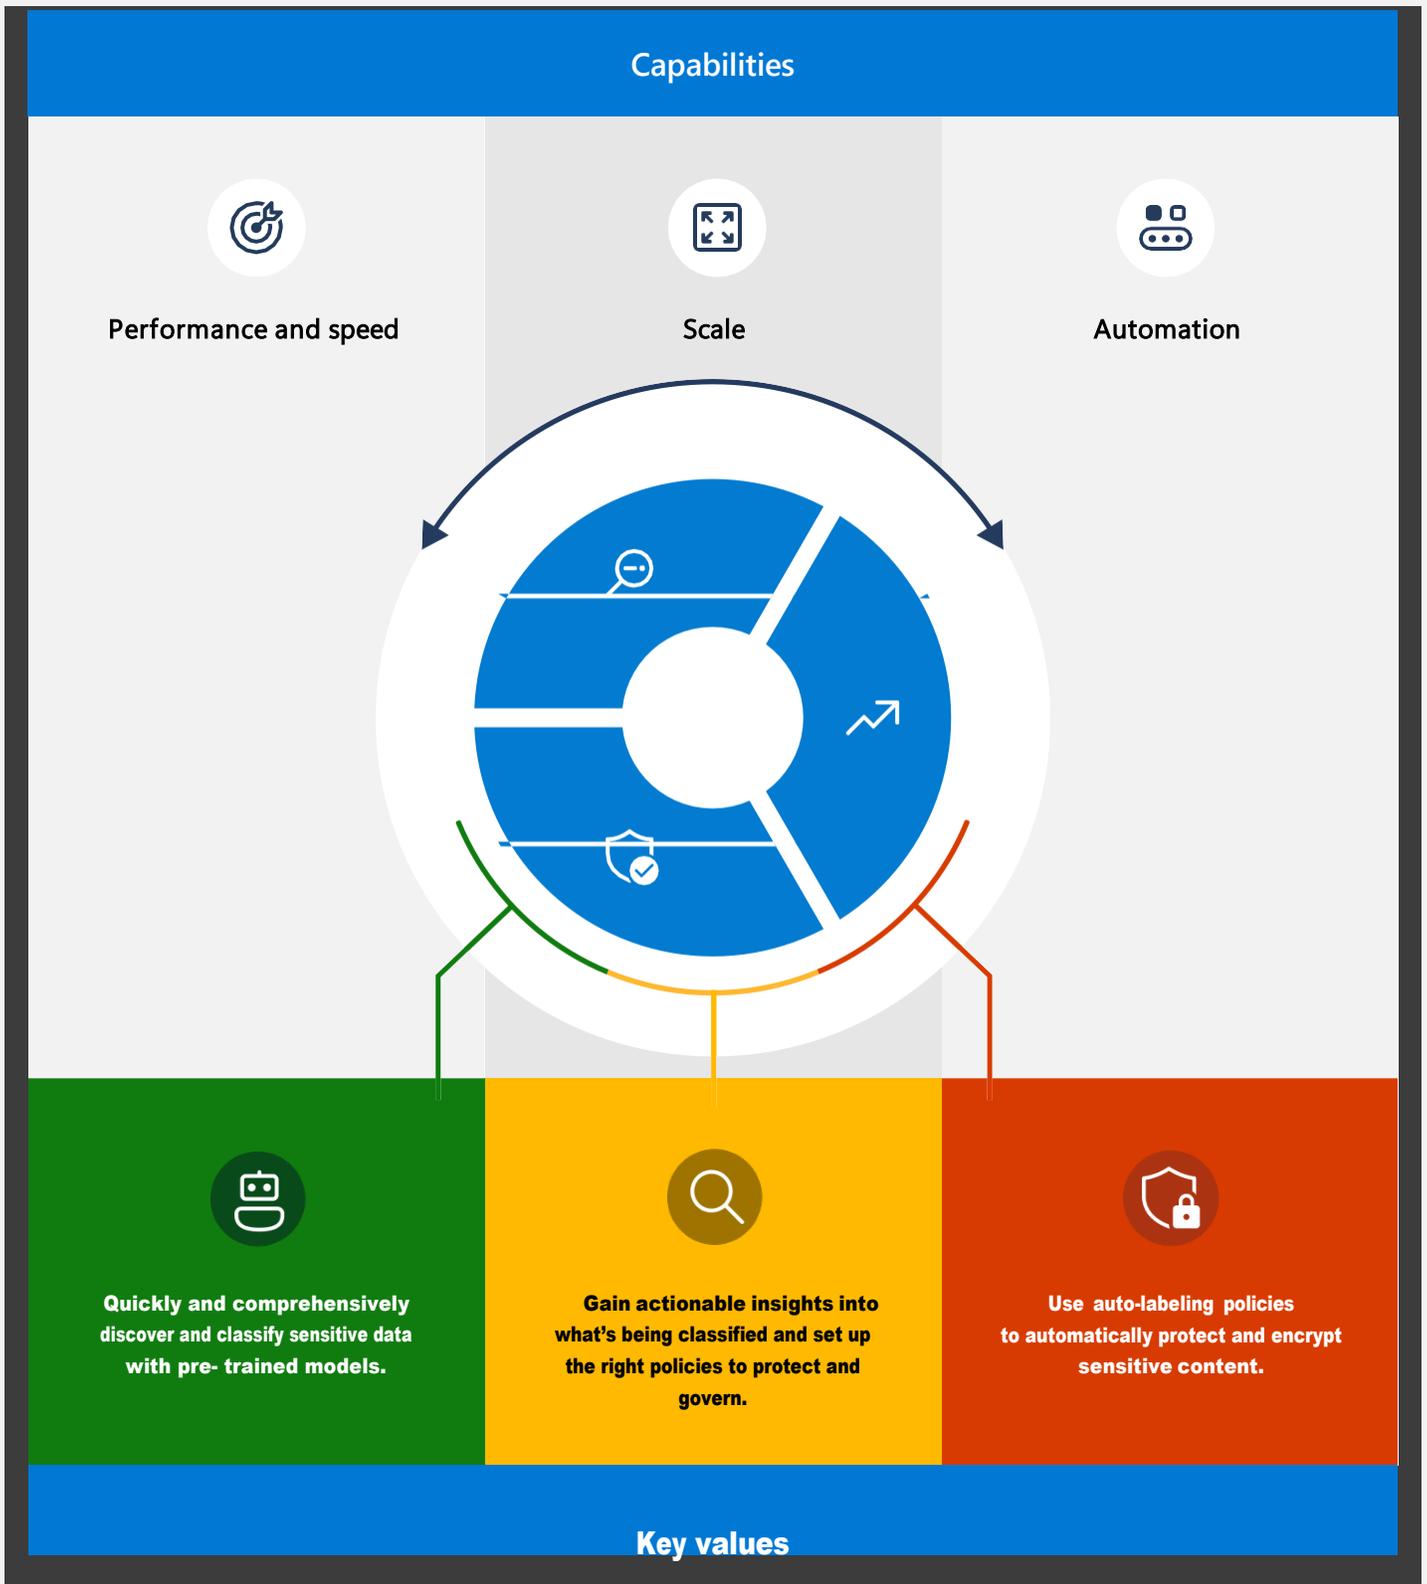


Figure 4: How trainable classifiers can help

AI classifiers for different business functions

A typical data map of an organization comprises data originating from multiple data sources. AI-based applications are best suited for this data geography because they can adapt to meet dynamic requirements. For example, organizations should use pre-trained, ready-to-use classifiers to discover and protect generic documents and

data for common business functions like legal, human resources, sales and marketing, research and development, and finance. For proprietary, organization-specific, or market vertical-specific documents, it's best to use custom classifiers that are trained using organizations' own document examples. For regulatory functions, such as the pharmaceutical, banking, and insurance industries that have standard regulatory templates and policies, fingerprinting is best suited to discovering and protecting these standard documents. (Figure 5)

Typical data map of organizations and Microsoft Purview AI solutions for each use case

Organizations have massive volumes of data across various functions and categories.

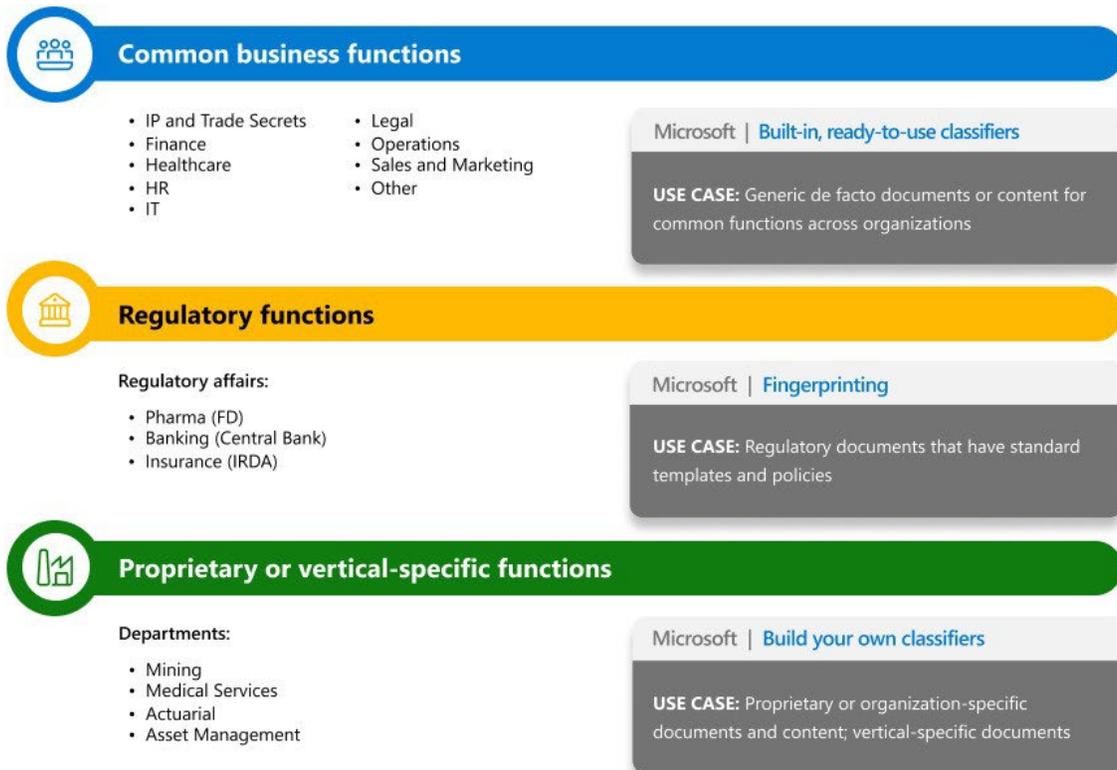


Figure 5: Data map of organizational functions and categories

With Information Protection, organizations can use our advanced classifiers to automatically identify and label files and emails as sensitive in a scalable way.

Out-of-the-box classifiers for nine different common business categories

Microsoft has created and pre-trained multiple classifiers that can help increase the coverage and performance of data classification while reducing false positives. Over 46 pre-trained, ready-to-use trainable classifiers can identify 100 types of sensitive content in English. Common and critical categories will soon be released in other global languages, allowing you to scan generic de facto documents with content from everyday organizational functions—including finance, IT, intellectual property and trade secrets, legal, healthcare, human resources, and operations.

These classifiers are pre-trained using a diverse number of real-world samples to ensure they provide broad coverage of various types of business functions. You can use them to discover and automatically label files and documents, apply retention labels for data lifecycle management, establish conditions in data loss prevention policies, and monitor inappropriate content for communication compliance. These classifiers can also be used to discover and classify standard types of sensitive data found across the business functions shown in Figure 6.

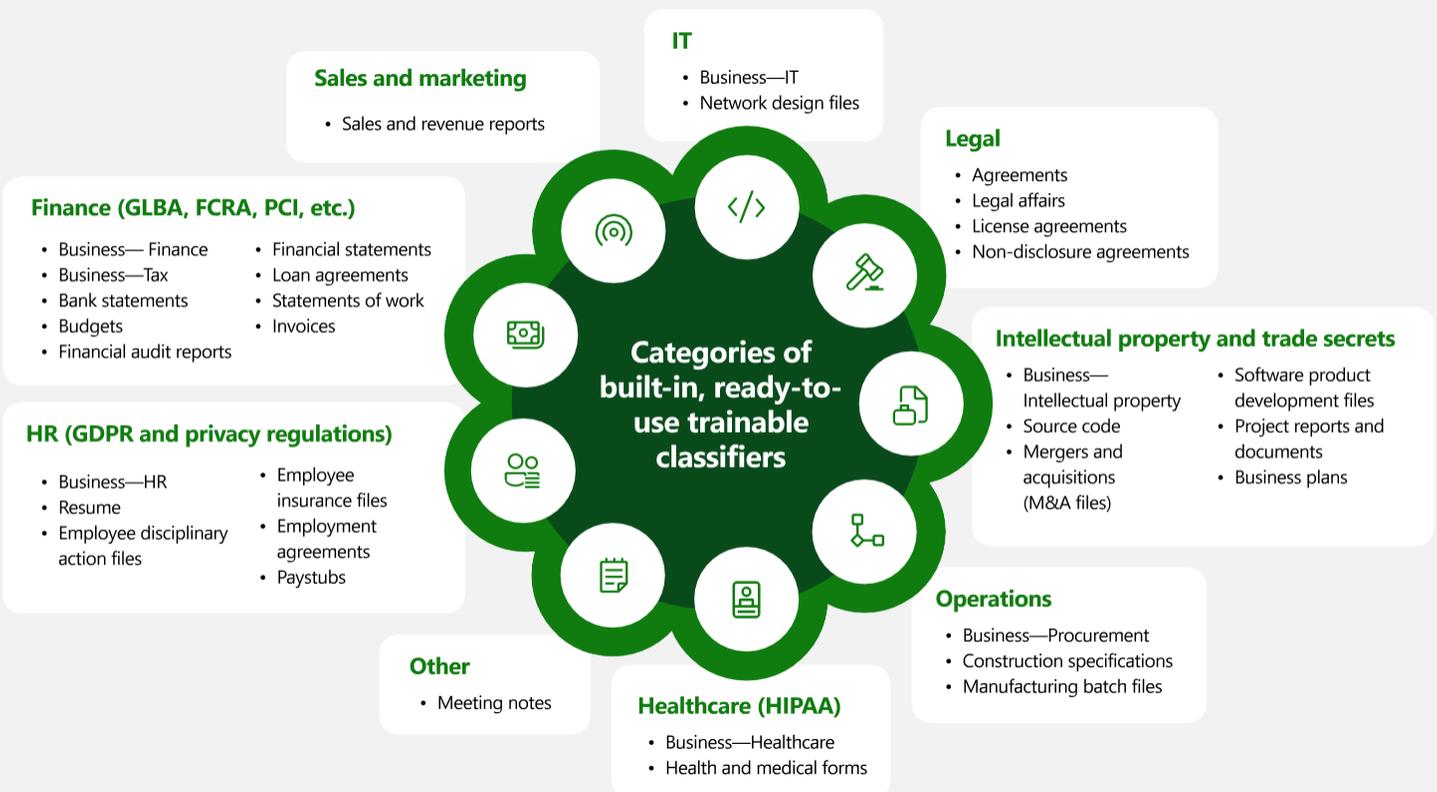


Figure 6: Categories of built-in, ready-to-use trainable classifiers

Explore Microsoft Purview

By efficiently categorizing and labeling content, Information Protection trainable classifiers enable organizations to protect sensitive data across multiple fronts. (Figure 7)

Use of trainable classifiers



Know your data in content explorer

Discover various categories of content that match the trainable classifiers in your content explorer, with no need to create policies to discover content.



Data lifecycle management policies

Auto-apply retention labels to content-matching trainable classifiers.



Information protection: sensitivity auto-labeling

Auto-apply sensitivity labels on content matching trainable classifiers.

46 trainable classifiers can identify more than 100 types of sensitive content.



Data loss prevention solution

Use trainable classifiers as a condition in DLP policies.

Figure 7: Using trainable classifiers inside content explorer and with Microsoft compliance solutions

Know your data

Information Protection provides out-of-the-box trainable classifiers for sensitive business document discovery and classification to help you better understand your data. Using [content explorer](#), you gain visibility into the sensitive data that has been discovered and tagged by all our advanced classifiers.

Content explorer gives you a current snapshot of the items that have a sensitivity or retention label and provides:

- Visibility into the amount and types of sensitive data and the ability to filter by label or sensitivity type for a detailed view of the locations in which sensitive data are stored.
- The ability to identify documents that are classified with sensitivity and retention labels.
- The ability to discover and view categories of sensitive data content that match these trainable classifiers and specific files containing sensitive data in SharePoint, OneDrive, Teams, and Exchange. (Figure 8)

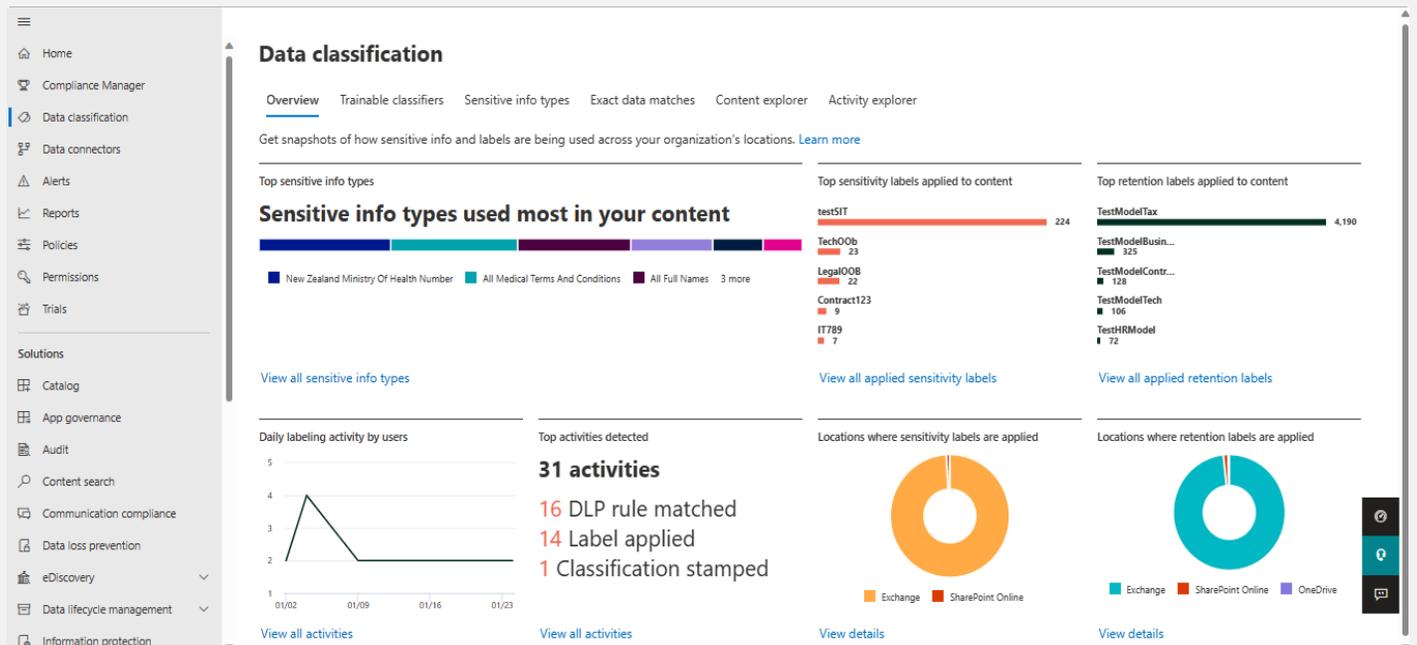


Figure 8: Data classification overview page

Sensitivity auto-labeling

Information Protection can use trainable classifiers in server-side [auto-labeling](#) policies for SharePoint, OneDrive, and Exchange. You can now take advantage of this capability to more quickly and comprehensively discover, label, and protect massive volumes of sensitive data across your digital estate with pre-trained models optimized for performance and scalability. Figure 9 shows how to add trainable classifiers in auto-labeling for files and emails.

The screenshot displays the 'Data classification' section in the Microsoft 365 compliance center. The 'Trainable classifiers' tab is active, showing a list of 109 published classifiers. A green notification banner at the top indicates that analytics are ready for creating and testing classifiers. The table below lists three classifiers: 'Actuary reports', 'Agreements', and 'Asset Management', all of which are built-in, in English, and ready to use.

Name	Accuracy	Status	Type	Language	Created by	Last modified
Published (109)						
<input type="checkbox"/> Actuary reports	-	Ready to use	Built-In	English	Microsoft	
<input type="checkbox"/> Agreements	-	Ready to use	Built-In	English	Microsoft	
<input type="checkbox"/> Asset Management	-	Ready to use	Built-In	English	Microsoft	1/11/2023

Figure 9: List view of Trainable Classifiers

Data loss prevention

To help protect sensitive data and reduce the risk of data loss, organizations can use Microsoft Purview Data Loss Prevention (DLP) to prevent unauthorized data sharing and data exfiltration. Our DLP solution now supports all advanced classifiers, including trainable classifiers, on various DLP workloads such as SharePoint, OneDrive, Teams, Exchange, and endpoint devices. Visit our [Learn about data loss prevention](#) webpage to check our detailed list of applications and services, workloads, platforms and endpoints, on-premises file shares, and non-Microsoft applications that are supported by our DLP solution.

With our DLP tools, you can:

- Efficiently monitor the activities users take on sensitive items at rest, sensitive items in transit, or sensitive items in use and take protective actions.
- Enable system administrators to create DLP rules specific to a category, such as Health, and designate specific actions when a DLP rule matches specific content or files, like sending incident reports and alerts to system administrators. (Figure 10)

Data loss prevention > Create policy

Choose the information to protect
Name your policy
Locations to apply the policy
Policy settings
Advanced DLP rules
Test or turn on the policy
Review your settings

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

1 item

Name	Status		
Source code demo rule	On		
Conditions Content contains any of these sensitive info types: Source code			
Actions Notify users with email and policy tips Send incident reports to Administrator Send alerts to Administrator			

Back Next Cancel

Figure 10: Creating a DLP policy and rule with Trainable Classifiers

Data lifecycle management

Microsoft Purview Data Lifecycle Management (DLM) helps you govern your data and manage the retention and deletion of business content to meet your legal, business, privacy, and regulatory content obligations. Retaining high-risk or high-value content can help in preventing data from malicious deletion or ransomware and reduces the risk of data breaches.

Data Lifecycle Management provides tools and capabilities to help you retain the content you need and delete the content you don't. A key differentiator

of our capabilities is that they all happen in one place, reducing information silos and the need for multiple copies.

Our DLM solution allows you to create retention policies, as shown in Figure 11, and supports SharePoint, OneDrive, Teams, and Exchange workloads. Visit our [Microsoft Purview Data Lifecycle Management](#) webpage to get more information on our DLM solution.

Create retention label

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

✓ 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.

Done

Figure 11: Creating a retention policy

Summary

Traditional classification techniques such as regular expressions, manual, or rule-based approaches can't easily handle massive volumes of data. These types of approaches are only appropriate for specific use cases in which a small number of highly-sensitive documents are labeled for access by specific groups or individuals.

Using machine learning-powered trainable classifiers enables organizations to quickly and comprehensively discover, label, and protect massive volumes of sensitive data across their digital estate with pre-trained models optimized for performance and scalability. Information Protection delivers a unified set of capabilities for data classification, labeling, and protection, not only in Office apps, but also in other popular productivity services where information resides, like SharePoint Online, Exchange Online, and Microsoft Teams, and endpoint devices.

We invite you to learn more about this game-changing new technology and how your organization can benefit from it.

Learn more

Microsoft Purview websites

[Information protection](#)

[Data loss prevention](#)

[Data lifecycle management](#)

Microsoft Purview blogs

[Information protection](#)

[Data loss prevention](#)

[Data lifecycle management](#)

Technical documentation

[Information protection](#)

[Data loss prevention](#)

[Data lifecycle management](#)

[Trainable classifiers](#)