


# Data security as a foundation for secure AI adoption

A step-by-step guide to prepare for AI deployment and secure AI adoption



# Table of Contents

03/

Introduction

04/

Current generative AI landscape

05/

Four steps to prepare your data for AI

- Know your data
- Clean up your data
- Protect your data
- Prevent data loss

8/

Why your AI choice should be Copilot for Microsoft 365

9/

How to secure and govern usage of Copilot in three steps

- Discover how Copilot is being used
- Protect data interacting with Copilot
- Govern the use of Copilot

12/

Addressing employees bringing their own AI to work

13/

Conclusion



# Introduction

The inception of generative AI brings tremendous value to organizations, helping them boost productivity and increase creativity. With this growing popularity, organizations are utilizing data in vastly different ways, which has reshaped the data landscape and their associated needs.

Despite all the benefits of AI, many customers considering AI worry about how to keep sensitive data safe. With AI being integrated into business operations, organizations are grappling with new data worries that could expose sensitive data.

End-users are using AI to seek access to sensitive data that isn't labeled, trying to gain access to confidential information. They might accidentally expose sensitive data, such as customer personal information, to AI applications. Or, they're leveraging AI to simplify content creation but without appropriate guardrails, are creating unethical or high-risk material. All these scenarios introduce new security risks.

On top of that, organizations are struggling with comprehensive visibility of their data. They don't know where—or even what—their sensitive data is. It's important they have that awareness because 83 percent of organizations experience more than one data breach in their lifetimes.<sup>1</sup> The good news is you can get ahead of data risks and protect your sensitive data.

Even before implementing AI, you can prepare your data environment for AI deployment by getting your data ready. This means identifying your sensitive data to clean it up, protect it, and prevent data loss. Then, you should prioritize secure AI adoption of a solution like Microsoft Copilot for Microsoft 365, referred to as Copilot later in this white paper. You should also take steps to secure and govern the usage of your data and ensure your compliance with government regulations and company policies.

This white paper discusses the emergence of new and evolving data security needs in the age of AI, provides strategies on preparing for AI adoption, and explores solutions that can strengthen your data security as you take advantage of Copilot and AI.

## Top data challenges



### Data oversharing

Users may gain access to sensitive data via AI apps that they're not authorized to view or edit because of a lack of labeling policies or the right access controls.



### Data leakage

Users may inadvertently leak sensitive data to unsanctioned AI apps. Sanctioned apps pose a risk as well if you haven't ensured the AI-generated responses inherit the data protection controls of the files referenced.



### Noncompliant usage

Users may generate high-risk content or content that doesn't abide by ethics standards with AI apps, such as documents created to hide insider trading, money laundering, or other illegal activities.

## What we've heard from customers



Shadow AI

# 58%

of organizations are concerned about the lack of visibility into the unsanctioned use of gen AI.<sup>1</sup>



Data oversharing  
and data leaks

# 80%+

of leaders cited leakage of sensitive data as their main concern.<sup>1</sup>



Increased  
regulatory liability

# 55%

of leaders lack understanding of how AI is and will be regulated and are seeking guidance on how to adhere to these requirements.<sup>2</sup>

1. PRNews wire, portal 26 report, November 2023.

2. ISMG, First Annual Generative AI study: Business Rewards vs. Security Risks, Q3 2023, N=400

# Current generative AI market landscape

The year 2024 has been monumental for AI. A growing number of organizations are integrating AI into their business workflows, and with a whopping 75 percent of knowledge workers using AI at work today,<sup>2</sup> the benefits are getting attention. In a recent Microsoft and LinkedIn survey<sup>3</sup>, people were quick to share how AI at work helps them to:

- Save time (90 percent).<sup>2</sup>
- Focus on their most important work (85 percent).<sup>2</sup>
- Be more creative (84 percent).<sup>2</sup>
- Enjoy their work more (83 percent).<sup>2</sup>

While employee enthusiasm is inspiring and a boon for AI user adoption, many are so eager to realize the productivity gains of AI that they're not going through the proper channels to implement it. Approximately 78 percent of AI users are bringing their own AI tools to work, and shadow IT is even more common at small and medium-sized companies (80 percent).<sup>2</sup>

Unfortunately, many organizations don't have the visibility needed to prevent sensitive data from going to unsanctioned places. Some organizations are reacting out of fear, with approximately 48 percent blocking AI entirely.<sup>2</sup> Even if your organization hasn't yet adopted or deployed AI at the enterprise

level, your employees are actively bringing their own AI tools to work today. And those tools may lack the built-in security necessary to protect your organization, not to mention they may not prevent major risks like data oversharing, data leakage and non-compliant usage.


As a result of these risks, organizations are more eager than ever to secure their data and experience peace of mind when considering implementing the full capabilities of AI. We're hearing from organizations that express apprehension about cybersecurity and data loss risks. The good news is that 75 percent of organizations are more likely to adopt AI apps—and realize all the transformative benefits—when they come with assurance mechanisms for secure and compliant use.

To succeed, you need to get ahead of risks like data oversharing and data leakage by exercising good data hygiene—before even deploying AI. Let's start with the four steps to prepare your data for AI.


## How can we securely adopt AI to enable innovation?



of organizations believe the use of AI will result in a range of benefits—improved efficiency, innovation, and more.<sup>1</sup>



of organizations perceive significant potential risks from AI—cybersecurity, harmful use, and more.<sup>1</sup>



of users are bringing their own AI tools to work.

1. KPMG, Trust in Artificial Intelligence, 2023.

2. ISMG, First Annual Generative AI Study: Business Rewards vs. Security Risks, 2023.

3. <https://www.microsoft.com/en-us/worklab/work-trend-index/ai-at-work-is-here-now-comes-the-hard-part>

# Four steps to prepare your data for AI

You want to implement AI in your organization, but you want to minimize the risks we mentioned around data oversharing, data leakage, and noncompliant usage. But it can be challenging to discover and manage vast amounts of data, potentially leaving your sensitive data vulnerable to oversharing and data leakage if you don't address this before deploying AI. Taking steps can help you better prepare your organization and your data.

The [AI shared responsibility model](#) establishes the responsibilities of your organization and those of your service provider for using more secure generative AI and complying with regulations. This model recognizes the three layers of AI functionality:

- The **AI platform** that delivers capabilities to the applications
- The **AI application** that accesses these capabilities to deliver the service or interface the user wants
- The **AI usage** that covers how these capabilities are consumed

Understand these responsibilities before adopting AI. You also must prepare your data properly for its use in your AI tools to ensure you understand all your data that needs protecting, clean up data and permissions, protect your sensitive data, and prevent data loss. Take these four steps to increase your data security and ensure your data is ready when you implement AI tools.

## 1 Know your data

With data as the lifeblood of every organization and with cybersecurity threats multiplying, securing your data is essential. It can be challenging to discover and manage vast amount of data in your org, which could leave your sensitive data vulnerable to oversharing and leakage if you don't address this before deploying AI.

Protecting and governing data is a complex and multifaceted effort, and it's especially challenging if you don't know where all your data lives.

Unfortunately, 30 percent of decision-makers say they lack visibility for all their business-critical data.<sup>3</sup>

Microsoft Purview Information Protection provides tools to help you know your data. Take advantage of content explorer and activity explorer to locate sensitive data and identify risky activities performed on sensitive data across your organization. Admins can use Microsoft Purview's built-in sensitive information types (SITs) or custom SITs to classify and label your sensitive data. Alternatively, users can use built-in labels in their enterprise apps (e.g., Microsoft 365) to manually self-label their own sensitive files as they work.

### Why do this before AI adoption?

You can better protect your organization if you know, before you adopt AI, the potential sensitive data that could be referenced or at risk of overexposure by AI apps.

### Checklist:

- ☐ Take advantage of Microsoft Purview content explorer and activity explorer to locate sensitive data and identify activities done to this data across your organization.
- ☐ Classify and label sensitive data using SITs in Purview.
- ☐ Use built-in labels in enterprise apps such as Microsoft 365 to apply labels to sensitive content.

## 2 Govern your data

Non-compliant AI usage can result in severe regulatory violations and hefty fines. Prepare for compliant usage by cleaning up your data and permissions.


This means taking stock of your SharePoint sites and the permissions to them, running reports to identify which sites and files have open permissions, and then remediating them, applying SharePoint-wide policies for effective content management and deleting old or obsolete data.

### Why do this before AI adoption?

If you don't ensure you have strong compliance measures in place, you could be at risk later on and face fines for non-compliance.

### Checklist:

- ☐ Use Microsoft Purview machine learning classifiers to detect and mitigate risks.
- ☐ Clean up data and permissions in SharePoint.
- ☐ Apply SharePoint-wide policies for content management and delete old or obsolete data.





## 3 Protect your data

Protecting sensitive data across your digital estate is important prior to deploying Copilot. Copilot recognizes labeled and protected files and any generated output responses will inherit the highest sensitivity level of referenced files that the user has access to view.

Microsoft Purview Information Protection enables you to classify, label, and protect your data based on its sensitivity. Users can self-label their own files and content within enterprise apps and system admins can automatically classify data using SITs, autolabeling, and protect content. The label will determine whether protections such as rights management, watermarks, and encryption will be applied and who has access.

This helps enable secure collaboration based on the sensitivity level of the data itself, where data labeled "general" can be freely shared, but files labeled "highly confidential" are encrypted and can be limited to a small group of employees, for example. This unified labeling solution works across Microsoft apps, services, security solutions, and devices.

### Why do this before AI adoption?

The labels and protections of sensitive files and content referenced in Copilot prompts are inherited and automatically applied to Copilot-generated outputs for end-to-end protection.

### Checklist:

- ☐ Use Microsoft Purview's unified labeling solution to label sensitive data.
- ☐ Set your permissions to "general" if content can be shared broadly and "highly confidential" if access to files is encrypted and limited.

## 4 Prevent data loss

One crucial component of data security is ensuring you prevent data loss of your business critical data. Data loss prevention policies can help you prevent data loss, including the loss of any data generated by AI.

Microsoft Purview Data Loss Prevention helps prevent data exfiltration activities across various channels – including uploading to cloud, uploading to USB, sharing externally and more. Microsoft Purview capabilities can also be extended to data used on Windows 10 computers, the Chrome browser, on-premises files shares and SharePoint folders and document libraries, and Teams chat and channel messages.

### Why do this before AI adoption?

Ensuring you have data loss prevention policies in place can prevent the loss of AI-generated data and prevent users from sending sensitive data in AI prompts.

### Checklist:

- ☐ Establish data loss policies in Microsoft Purview to prevent unintentional sharing of sensitive data.
- ☐ Extend Microsoft Purview capabilities to where you share data.

# Why your AI choice should be Copilot for Microsoft 365

When you adopt generative AI, look for apps like Copilot that provide built-in security controls to prevent data oversharing. You can benefit from AI-driven advantages while protecting your company from the data security risks associated with AI. And you can achieve both in ways appropriate for your business. Copilot should be your top choice for several reasons.

First, it features powerful built-in security contained in a simple design. For all the value that Copilot brings to individual users and their broader organizations, the architecture itself is fundamentally simple. Copilot combines an orchestrator and an LLM, built on top of the existing Microsoft 365 solution. Copilot lets you create personalized documents, emails, chats, and other content using LLMs.

In addition, Copilot builds on and applies your existing privacy, security, and compliance commitments. It only has access to the content already available to the user. And it's managed with the same tools and standards that you use today. And as a core service, it's subject to the same terms and conditions that our customers expect.

## How Copilot for Microsoft 365 works

Copilot is an add-on available for Office 365 E3 and E5 and for Microsoft 365 E3 and E5 suites. Choose a path—Core or Best-in-class—based on your preferred level of security and compliance controls.

This data protection means:

- You control your data. Your data is encrypted and not used to train foundational LLM models in Copilot.
- You control your data location, data residency for data at rest, and European Union data boundary storing and processing.
- You receive commercial data protection for web-grounded prompts using the latest web data.

Before rolling out Copilot, complete an optimization assessment. This online questionnaire queries you about your situation, licensing status, data security posture, and other relevant matters. Based on your questionnaire answers and licensing and deployment status, you'll get a recommendation from Microsoft for your best deployment path. Your options are either the Core path or the Best-in-class path to elevate your data security.

Copilot is built with a comprehensive approach to security, compliance, privacy, and responsible AI, but there's much more value to be realized. On top of Copilot, you can layer on additional data security and compliance capabilities—like configuring label policies and applying sensitivity labels—to discover AI risks, protect sensitive data, and govern Copilot usage.



## Copilot for Microsoft 365

Built on Microsoft's comprehensive approach



Security



Compliance



Privacy



Responsible AI



# How to secure and govern usage of Copilot for Microsoft 365 in three steps

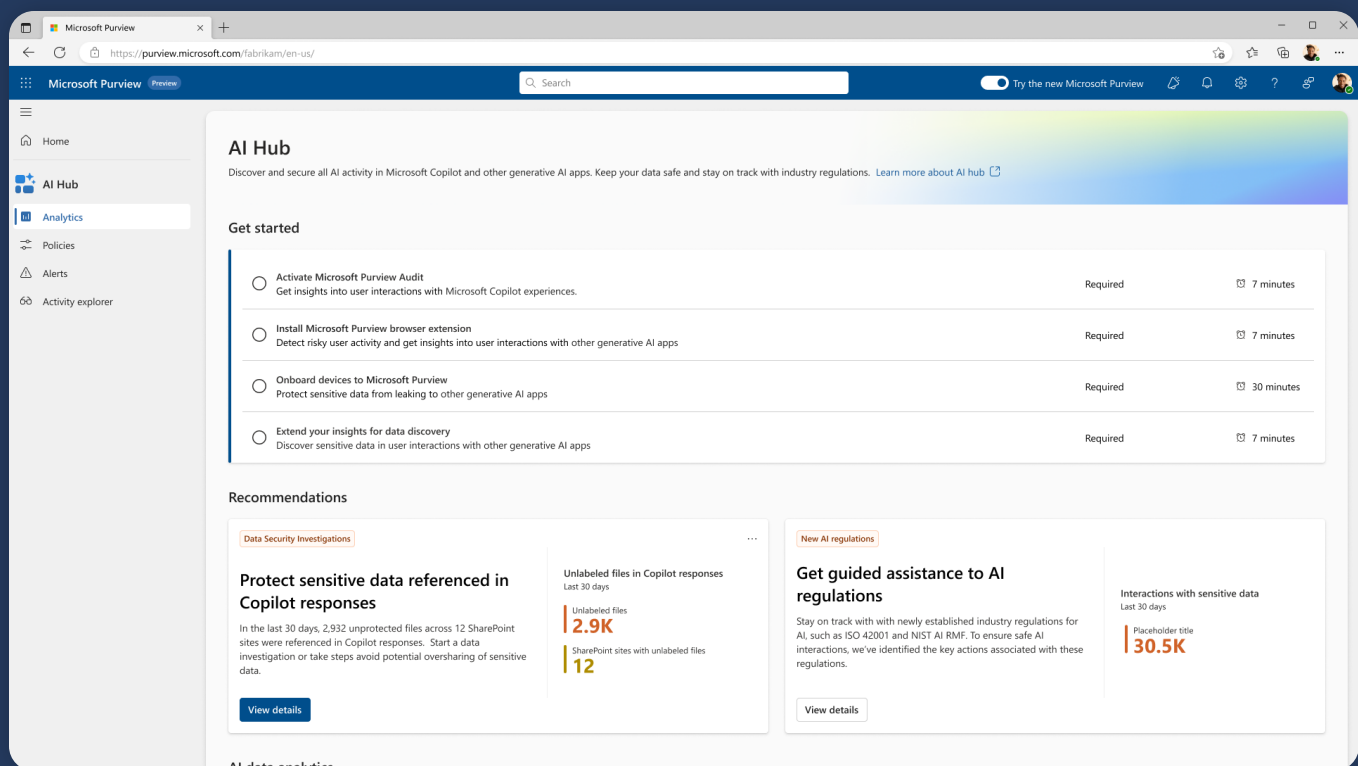
Congratulations! You've deployed Copilot, and your organization is on its way to using AI responsibly! Now, it's time to make sure you can manage its usage. With Copilot's built-in security or the other protection layers you added through Microsoft Security, you can now easily discover, protect, and govern your data to optimize your AI use.

## #1 Discover data risks with Microsoft Purview AI Hub

Security practitioners say that a lack of visibility into sensitive data is a deterrent to developing smart plans and actionable strategies to ensure data security. As you know, 30 percent of decision-makers say they don't know where or what their sensitive business critical data is.<sup>3</sup> With users sharing data with generative AI apps, it's essential to get that visibility into how sensitive data is flowing through AI, unethical interactions with AI, unlabeled files and SharePoint sites referenced by Copilot.

The Microsoft Purview AI Hub helps you discover how AI applications, including Copilot and third-party apps, are being used in your organization. You can benefit from ready-to-use policies for data protection. You'll gain a better understanding of user and data activities—including potentially risky ones—with audit logging and checking SharePoint site permissions for overshared content. For instance, you can see total AI interactions and the risk level associated with their usage.

In addition, you can discover sensitive data being shared with Copilot, unlabeled files referenced in prompts/responses, and insights that can help you determine if even more protections are needed.



The screenshot displays the Microsoft Purview AI Hub interface. The left sidebar contains navigation links: Home, AI Hub, Analytics, Policies, Alerts, and Activity explorer. The main content area is titled 'AI Hub' and includes a subtitle: 'Discover and secure all AI activity in Microsoft Copilot and other generative AI apps. Keep your data safe and stay on track with industry regulations. Learn more about AI hub'. Below this is a 'Get started' section with four tasks, each with a radio button, a description, a 'Required' status, and a time estimate:

Task	Required	Time
Activate Microsoft Purview Audit Get insights into user interactions with Microsoft Copilot experiences.	Required	7 minutes
Install Microsoft Purview browser extension Detect risky user activity and get insights into user interactions with other generative AI apps	Required	7 minutes
Onboard devices to Microsoft Purview Protect sensitive data from leaking to other generative AI apps	Required	30 minutes
Extend your insights for data discovery Discover sensitive data in user interactions with other generative AI apps	Required	7 minutes

Below the 'Get started' section is a 'Recommendations' section with four data insights:

- Data Security Investigations**  
**Protect sensitive data referenced in Copilot responses**  
In the last 30 days, 2,932 unprotected files across 12 SharePoint sites were referenced in Copilot responses. Start a data investigation or take steps avoid potential oversharing of sensitive data.  
[View details](#)
- Unlabeled files in Copilot responses**  
Last 30 days  
Unlabeled files: **2.9K**  
SharePoint sites with unlabeled files: **12**
- New AI regulations**  
**Get guided assistance to AI regulations**  
Stay on track with with newly established industry regulations for AI, such as ISO 42001 and NIST AI RMF. To ensure safe AI interactions, we've identified the key actions associated with these regulations.  
[View details](#)
- Interactions with sensitive data**  
Last 30 days  
Placeholder title: **30.5K**

## #2 Protect sensitive data throughout its AI journey

Implementing an AI strategy can be a challenge for organizations that lack controls to detect and mitigate risks, particularly around the leakage of intellectual property through AI tools. One example is a situation where confidential details of a highly sensitive project are inadvertently disclosed to unauthorized users. Unsurprisingly, organizations want to ensure that sensitive data does not get into the hands of people who should not have access to it.

Microsoft Purview provides data security controls to ensure Copilot responses are based on a user's permissions and only those with appropriate access rights can see the content. In addition, Microsoft Purview Information Protection provides controls on prompts and responses such as encryption, watermarking, autolabeling, and label inheritance to prevent sensitive data from being overshared.

For third-party generative AI apps, the capabilities of Microsoft Purview Data Loss Prevention restrict users from pasting sensitive data into generative AI prompts. And adaptive protection in Microsoft Purview blocks high-risk users from pasting sensitive data into AI apps while allowing low-risk users to do so.

### How does labeling work?

Add sensitivity labels to your data in Microsoft Purview Information Protection for Microsoft 365 apps and services, SQL Server, Azure Data Lake Storage, and Microsoft Fabric. Copilot inherits any labels from these files. Autolabeling makes this even easier by applying sensitivity labels to files and emails based on sensitive data input.

The screenshot shows a PowerPoint presentation titled "Project Obsidian\_Financial Projections" with a "CONFIDENTIAL" watermark. The slide content includes:

**Financial Projections**

When we announce the upcoming merger on May 6th, our analysts have forecasted a 200% increase in our stock price.

Today's Stock Price	May 6th Stock Price	% increase
\$100	\$300	200%

Overlaid on the presentation is a Copilot chat window. The chat history shows a previous response: "is forecasted to result in a 200% increase in stock price". The current prompt is "What are the financial projections for Project Obsidian?". The chat window also displays a sensitivity label: "Confidential/Project Obsidian".

## #3 Govern the use of Copilot and comply with emerging AI regulations

In today's regulatory landscape, compliance and risk managers are increasingly concerned about non-compliant AI usage, which can lead to severe regulatory violations and hefty fines. Regulatory compliance is a priority and Microsoft Purview offers integrated compliance tools specifically designed to address these challenges. Copilot capabilities include:

- Audit to capture when Copilot interactions occur.
- Data lifecycle management to capture the content of interactions and manage their retention and deletion.
- Communication compliance to detect noncompliant usage of Copilot prompts.
- eDiscovery to search Copilot interactions and assist with investigations.

### Comply with regulations and standards

With evolving AI regulations and rapidly advancing technology, organizations must prioritize compliance obligations when developing and using AI applications. When onboarding and deploying AI solutions, you need to comply with these regulations

to avoid penalties and reduce your data security, data compliance, and data governance risks.

Recently, several notable pieces of legislation and frameworks have been introduced. These regulations align with compliance requirements such as detecting AI interactions and preventing data loss in AI applications.

Four new Microsoft Purview Compliance Manager assessment templates will help your organization assess, implement and strengthen its compliance against AI regulations, including EU AI Act, NIST AI RMF, ISO/IEC 23894:2023 and ISO/IEC 42001. Find details in the Microsoft Purview AI Hub.

Being committed to responsible AI is also important to prevent issues. By 2027, Gartner predicts at least one global company will have an AI deployment banned by a regulator for not complying with data protection or AI governance legislation.<sup>4</sup>

The screenshot displays the Microsoft Purview Compliance Manager web interface. The top navigation bar includes the 'WOODGROVE' logo, 'Microsoft Purview', a search bar, and a toggle for 'New Microsoft Purview portal'. The main content area is divided into sections: 'Analytics', 'Policies', and 'Activity explorer'. A 'Get started' section lists three required actions: 'Turn on Microsoft Purview Audit', 'Install Microsoft Purview browser extension', and 'Onboard devices to Microsoft Purview'. A 'Recommendations' section highlights 'Get guided assistance to AI regulations' with a 'View Details' button. A right-hand sidebar provides detailed guidance on AI regulations, including a bar chart showing 'Sensitive information types in AI interactions found across your organization' (Copilot) with categories like Project Obsidian, Stock Manipulation, All Full Names, and Finance. It also lists 'Actions to take' such as 'Monitor logs for AI interactions' and 'Manage retention and deletion policies for AI interactions', along with 'Useful Resources' and buttons for 'Mark as complete' and 'View regulations'.

# Addressing employees bringing their own AI to work

Earlier, we mentioned the challenge of overeager employees bringing third-party AI tools like Google Gemini into the workplace without permission. Shadow IT is a major issue with the growing popularity of AI, and it's one that Microsoft cybersecurity solutions can address.

Adding other protection layers through Microsoft Security can enhance data security even further. Microsoft Defender for Cloud Apps makes discovery of generative AI app usage easier by detecting more than 400 generative AI applications. In the Defender for Cloud portal, choose the Generative AI category for a list of all the apps and the risk profile score assigned to each. Click on a specific app for the risk profile, which includes details like security controls the app supports and regulations the app complies with.

You can also filter apps by risk score and see usage trends for individual apps in Cloud Discovery. In the Defender for Cloud portal, you can develop app use policies based on what you learn about AI app usage; you can even block a problematic app by tagging it as unsanctioned.


Run Cloud Discovery in Microsoft Defender for Cloud Apps to discover unauthorized AI employees bring to work and other apps in use. You can also generate reports that show the risk levels of the visited apps and sanction or unsanction the use of each.

In the AI hub of Microsoft Purview, you can see the number of prompts containing sensitive data shared with third-party apps and other critical metrics to help you shape your policies and address shadow IT. You can even block paste-to-browser to prevent employees from sharing sensitive data with third-party apps.

These strategies will enable you to encourage employees' use of approved AI tools and discourage app usage outside secure and approved channels, protecting both your organization and your people.

The screenshot displays the Microsoft Defender Cloud Discovery interface. The top navigation bar includes 'Microsoft Defender' and a search bar. The main header is 'Cloud Discovery', with a sub-header 'Discovered apps'. Below this, there's a 'Queries' section with a dropdown menu and a 'Save as' button. A filter bar shows 'Apps: Apps...', 'App tag: None', 'Risk score: 0 to 10', and 'Compliance risk factor: Select factors'. The main table lists discovered apps with columns for App, Risk, Tags, Traffic, Upload, Transact..., Users, IP address, Last seen, and Actions. The table shows six apps: Microsoft Bing Chat (Risk 10), Google Bard (Risk 9), OpenAI ChatGPT (Risk 8), Soundful (Risk 5), Nichesss (Risk 4), and Bramework (Risk 3, marked as UNSANCTIONED). The interface also includes a 'Bulk selection' dropdown, a 'New policy from search' button, and an 'Export' button.

App	Risk	Tags	Traffic	Upload	Transact...	Users	IP address	Last seen	Actions
Microsoft Bing Chat Generative AI	10		20 MB	6 MB	167	142	120	Sep 14, 2023	✓ ✕ ⋮
Google Bard Generative AI	9		2 MB	447 KB	13	13	11	Sep 14, 2023	✓ ✕ ⋮
OpenAI ChatGPT Generative AI	8		16 MB	4 MB	134	121	89	Sep 14, 2023	✓ ✕ ⋮
Soundful Generative AI	5		3 MB	962 KB	28	28	21	Sep 14, 2023	✓ ✕ ⋮
Nichesss Generative AI	4		2 MB	550 KB	16	16	9	Sep 14, 2023	✓ ✕ ⋮
Bramework Generative AI	3	UNSANCTIONED	2 MB	481 KB	14	14	9	Sep 14, 2023	✓ ✕ ⋮



## Conclusion

Data security is a necessity as AI adoption increases and organizations explore new ways to benefit from AI. To protect your sensitive data, you must prepare your data. Take measures to secure and govern your data to minimize security risks and gain greater peace of mind when you introduce AI in new areas of your business. These measures include cleaning up your data and permissions and implementing data protection.

This is easier if you securely deploy [Copilot for Microsoft 365](#). This enterprise-ready generative AI features built-in controls that strengthen your data security. Take steps to discover, protect and govern your data. Copilot is even more powerful when combined with [Microsoft Purview](#), which provides additional data and compliance controls to secure and govern Copilot while giving you visibility and protection controls for third-party applications.

Implementing the strategies described in this guide will give you the confidence to innovate and seize exciting new avenues for AI, reaping all the productivity benefits of this industry-transforming technology.

1. IBM, [Cost of a Data Breach Report 2022](#).
2. Microsoft, [2024 Work Trend Index Annual Report from Microsoft and LinkedIn](#), 2024.
3. Microsoft, [Survey of 510 US compliance decision-makers commissioned from agency Vital Findings](#), March 2023.
4. Gartner, ["What is Artificial Intelligence?"](#)