



Unlock AI Potential with Secure Foundations



Table of Contents

- 01** The AI Opportunity and Security Imperative
- 02** Good Intentions with Potential Risk: Example AI Scenarios
- 03** Building a Secure, AI-Ready Foundation with Microsoft 365 E3
- 04** Continue Your AI Adoption Journey





INTRODUCTION

The AI Opportunity and Security Imperative

There is a growing productivity capacity gap amongst today's workforce. In the recent Work Trend Index report, **more than 80%** of knowledge workers surveyed across 31 countries – including both employees and leaders – reported feeling they are lacking in the time or energy to do their work¹. At the same time, **more than 50%** of business leaders surveyed said that productivity needs to increase to meet goals¹. This presents an opportunity to assess the ways we work and potential solutions to help bridge this gap. AI, including agents, holds a lot of promise as a solution. **More than 80%** of surveyed leaders reported feeling confident that their organisations will start using AI agents to expand workforce capacity in the next 12–18 months¹.

This rush toward AI comes with a caveat: innovation should be balanced with robust security and data protection to mitigate potential risks. In a recent iSMG study, **more than 70%** of business leaders surveyed reported that ensuring data privacy and security is the primary challenge for integrating generative AI with existing IT infrastructure². Leaders face a dual imperative: accelerate AI adoption to drive productivity and protect data in this new era.

This eBook will provide insight into example AI scenarios that could pose potential risks, tips to mitigate those risks with Microsoft 365 E3 and a checklist to help organisations get started or continue to prepare for AI adoption at scale. Get ready to embrace AI with confidence: enabling your organisation to innovate with AI, without compromising your organisation's security or compliance standards.





Good Intentions with Potential Risk: Example AI Scenarios

The biggest risks with AI use and adoption often stem from well-intentioned employees trying to do their job and drive greater impact. In the rush for increased productivity, team members may inadvertently overshare sensitive information with AI tools or use applications that are not yet approved by IT. This may create risks and could expose the organisation to new threats.

Let's look at two example scenarios to illustrate this point – one with an office-based information worker and the other with a frontline worker.

In these fictional examples, you'll learn 1) how well-intentioned AI use could lead to security compromises if not properly governed; and 2) how these potential risks could be mitigated with a robust security foundation.



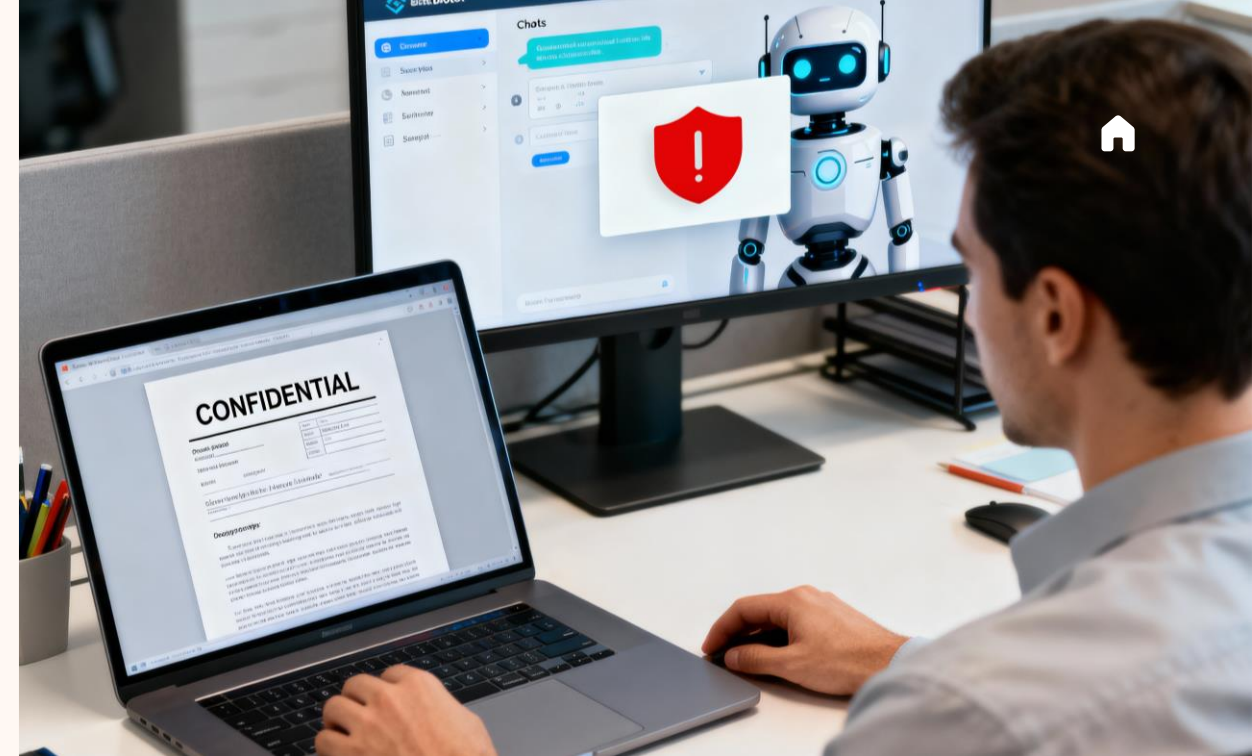
Example Scenario 1⁴

A Product Marketing Manager leverages an AI tool to draft a press release

A Product Marketing Manager at a manufacturing company is under a tight deadline to draft a press release about a new product launching in the coming months. To brainstorm ideas and speed up the writing process, this employee copies and pastes excerpts from a confidential product roadmap and an internal financial projection into a free online AI writing assistant. The AI tool quickly generates some polished paragraphs, which the employee then uses in the press release draft. However, this tool has not yet been vetted or approved by their organisation.

Potential Risk:

By using this external AI service, the employee has uploaded sensitive company data (product plans and financial information) to a third-party system that the IT team does not have visibility into and cannot manage or control. Once the data is shared with the AI tool, the company loses control over it as it is beyond authorised boundaries. The organisation is now unable to know how long the data is stored, who has access to it or how it is used by that third party. What began as a quick productivity win has turned into a data leakage incident. This scenario reflects potential internal risk when a well-intentioned employee shares sensitive data with an unsanctioned AI tool, creating vulnerability for the organisation.



Risk mitigation with Microsoft 365 E3:

Microsoft 365 E3 includes several features and capabilities that could help avert this outcome.

Microsoft Purview Information Protection Plan 1 allows you to apply sensitivity labels and encryption to documents like the product roadmap and financial project documents. Had these documents been classified and encrypted with sensitivity labels like *Confidential* with encryption, the documents' access could have been restricted to authorised users only. Additionally, **Cloud app discovery** can help identify Generative AI cloud apps that are being used (from the many apps in the cloud app catalogue), by which user and how frequently. IT could have uncovered that employees were using this particular AI tool and taken steps to limit its use with **Microsoft Defender for Endpoint P1** or provide a sanctioned alternative. Microsoft 365 E3 provides the tools and capabilities to help organisations improve visibility and mitigate risk.

Example Scenario 2⁴

A retail store manager uses AI to draft employee schedules and performance notes

A store manager at a large retail chain is overwhelmed with administrative tasks and discovers a free AI productivity app that promises to help with scheduling and documentation. To save time, the manager uploads a spreadsheet containing employee names, phone numbers, hourly wages, availability preferences and recent performance notes into the AI tool. They ask the AI app to create an optimised work schedule and help draft performance review summaries for upcoming one-on-one meetings.

Potential Risk:

By uploading this employee data to an unvetted AI app, the manager has exposed sensitive personal information of their employees to a third-party system outside the company's control and oversight. This could create potential privacy violations and compliance issues with labour laws. The company's HR and IT departments have no visibility into where the employee data now resides, how long it's stored or who may access it. What began as an attempt to be more efficient has created serious legal, privacy and employee trust issues.



Risk mitigation with Microsoft 365 E3:

Microsoft 365 E3's device and application management capabilities could help mitigate frontline scenario risks like this. With **Microsoft Intune P1**, the retailer's IT team could enforce application protection policies to help apply rules to applications that access work data, such as restricting copy/paste and data transfer to unmanaged apps. If the retailer creates Conditional Access policies via **Microsoft Entra ID P1**, they may be able to help block sign-in attempts to corporate accounts from unmanaged apps or devices. In this scenario, if the manager's phone is not compliant or the app is unmanaged, these policies could help prevent them from accessing and uploading restricted company files to a shadow AI application.

KEY INSIGHT: In both scenarios, the employees' aim was positive – improve work efficiency with AI. The outcome, however, could be negative if proper guardrails and policies aren't in place. These examples emphasise why adopting AI without a safety net can bring about new and additional risks for organisations. The goal is not to discourage AI use, but to enable it responsibly. The next section examines the fundamental capabilities to consider when evaluating solutions that may improve productivity capacity and the Microsoft 365 E3 capabilities that help organisations set guardrails that help keep sensitive data in approved, secure channels.

Building a Secure, AI-Ready Foundation with Microsoft 365 E3

Is your organisation ready to scale AI? Review the checklist below and learn how Microsoft 365 E3 can help.

As AI adoption accelerates, it is increasingly important to evaluate your organisation's security posture and identify potential gaps before scaling AI. Microsoft 365 E3 is designed to be an AI-ready productivity and security foundation – combining a broad range of identity, security, compliance and management tools that work together to empower you to scale your AI adoption with confidence.

Use the following checklist as a starting point to understand your organisation's current position. If any items are not yet fully implemented, your organisation may be exposed to unnecessary risk.



Enforce strong identity and access management policies like multifactor authentication (MFA) and conditional access.

Microsoft 365 E3 includes **Microsoft Entra ID P1** which enables enforcement of robust identity and access management across both cloud and on-premises environments. Solutions like MFA and conditional access policies verify each sign-in and restrict or allow access based on criteria such as user role, device health and location, helping to mitigate unauthorised access and data oversharing. For example, you could restrict access to AI applications to specific employee groups and require MFA for login, reducing the risk of account compromise. Conditional access policies can further protect sensitive data by preventing employees from accessing unsanctioned applications with corporate credentials and shared devices can enforce sign-out rules to protect against unauthorised session use.



Enable secure collaboration across managed endpoints and devices.

Microsoft 365 E3 delivers intelligent endpoint security through solutions like **Intune P1, Defender for Endpoint P1 and Windows E3**. These tools help manage and increase security for PCs, mobile devices and apps. Organisations can enforce web filtering to block risky sites (including unsanctioned AI tools) and keep devices updated with security patches via **Windows Autopatch and Autopilot**. These controls help keep devices healthy, managed and up to date with security patches, reducing the attack surface. Unified endpoint and device management allows organisations to secure work devices and maintain application compliance, whether employees use their own devices or company-provided ones.



Protect data by using capabilities like sensitivity labelling and data loss prevention policies.

Microsoft 365 E3 provides core **Purview** capabilities – sensitivity labelling, encryption and Data Loss Prevention (DLP) – to protect sensitive information across documents and emails. Once label policies are published, employees can classify files containing financial data, personal information or intellectual property (for example, with a 'Confidential' sensitivity label) applying encryption and access controls. DLP policies may detect and restrict unauthorised sharing across **Exchange Online and SharePoint Online/OneDrive for Business**. Even when using **Microsoft 365 Copilot**, these protections remain enforced: Copilot honours the highest sensitivity label and prevents users without permission from accessing or reasoning over protected content.



Get visibility into unsanctioned AI tools being used and take informed action.

Microsoft 365 E3 includes **Cloud app discovery**, which helps you understand what AI applications are being used (from those included in the cloud app catalogue), by whom and the associated risk scores, so that you can determine what action to take based on that information. If deemed too risky, IT can use **Defender for Endpoint P1** to limit individuals' abilities to visit that URL on managed Windows 11 devices. For example, if you discover that 100 users are using a given AI image generator tool that has not been reviewed or approved by IT – IT can investigate it based on the organisation's security and compliance standards. If deemed risky, IT can suggest a sanctioned alternative. These capabilities empower you to mitigate shadow AI risks and make informed decisions related to AI usage that are right for your organisation.



Provide sanctioned AI applications like Microsoft 365 Copilot.

Providing sanctioned AI applications delivers significant value by providing employees access to powerful, approved tools that meet your organisation's security and compliance standards. For example, instead of employees turning to unvetted AI image generators or chatbots, IT can offer sanctioned alternatives – such as **Microsoft 365 Copilot** – that are monitored, supported and integrated into the company's existing policies. **Microsoft 365 Copilot** delivers a secure, enterprise-grade AI solution deeply integrated with the Microsoft 365 productivity suite. **Copilot** empowers employees to work smarter, automate routine tasks and generate insights – all while respecting identity and data protection policies. This approach reduces shadow IT, helps limit data leakage and gives organisations visibility and control over how AI is used.



These capabilities don't operate in isolation – they reinforce each other as an integrated, multilayered defence. Identity, device and data protection policies work together to help you mitigate threats. When security teams know that controls like these are in place, security can become an enabler of innovation rather than a roadblock.

Microsoft 365 E3 empowers you to establish the right guardrails so your organisation can innovate with AI without compromising security.

Continue your AI adoption journey

Every organisation is on a journey to figure out how AI can best serve its goals. The promise is alluring – but AI adoption must go hand-in-hand with robust protection. With Microsoft 365 E3, you don't have to choose one or the other. You get a powerful suite that elevates your security and compliance posture, enabling you to embrace AI with confidence.

Assess Your AI Security Readiness.



Microsoft offers a **Security for AI Assessment**³. This [online assessment](https://security-for-ai-assessment.microsoft.com) (at security-for-ai-assessment.microsoft.com) provides recommendations to enhance your AI security based on the information you provide. It helps you evaluate your current AI security across four key pillars: **Prepare, Discover, Protect** and **Govern**, as well as actionable recommendations.

Learn More & Engage with Microsoft.



Dive deeper into Microsoft 365 E3 capabilities and how they can help you adopt AI with confidence. www.microsoft.com/en-us/microsoft-365/enterprise/e3?activetab=pivot:overviewtab. You can also reach out to your Microsoft account team or partner.

¹. [Work Trend Index Annual Report, 2025: The Year the Frontier Firm Is Born](#) Microsoft, Inc., April 2025. Microsoft analysed survey data from 31,000 full-time employed or self-employed knowledge workers across 31 countries between February 6, 2025, and March 24, 2025, LinkedIn labour market trends, and trillions of Microsoft 365 productivity signals.

². [Second Annual Generative AI Study: Business Rewards vs. Security Risks](#), iSMG, Inc., February 2025 – based on a survey of 360+ business and cybersecurity professionals, conducted in Q3 2024.

³. Security for AI Assessment: security-for-ai-assessment.microsoft.com/ Accuracy of results depends on the accuracy of your input. The report is for informational purposes only and does not guarantee performance or outcomes. It should not be interpreted as a commitment from Microsoft. Actual results may vary based on factors such as location, purchase method and usage. Microsoft makes no express or implied warranties regarding the content of this report or website.

⁴. The scenarios presented are fictional and for illustrative purposes only.

