nedscaper

Secure AI in practice: how to make your organisation Copilot-Ready

A proven path to Copilot Readiness: from risk assessment to safe AI adoption in practice.

NEDSCAPER.COM

Table of Contents

03 Secure AI in practice: how to make your organisation Copilot Ready

05 Different roles, different perspectives
- Chief Information Security Officer
- Compliance officer
- Security-analyst
- Business employees

06 5 key challenges for secure AI
- Lack of policy
- Lack of data classification
- Oversharing
- Legacy data
- Shadow-AI

07 Why legacy data can be a problem
A simple example

08 The shadow-AI epidemic

10 Step-by-step plan: 6 steps to Copilot Readiness
- **Step 1:** Start with awareness and stakeholder involvement
- **Step 2:** Establish policies and rules
- **Step 3:** Inventory data, access rights and AI usage
- **Step 4:** Implement foundational technical controls
- **Step 5:** Train and support employees
- **Step 6:** Monitor, evaluate and continuously improve

12 How Nedscaper prepares your organisation for secure AI

Secure AI in practice: how to make your organization Copilot-Ready

AI tools such as Microsoft Copilot are ushering in a new era of productivity and efficiency. By enabling employees to access information faster, automate processes, and make better-informed decisions, these technologies have the potential to significantly enhance organizational performance. However, alongside these opportunities come new risks. In this whitepaper, we explore how organizations can prepare themselves to manage these risks while maximising the value that AI can deliver.

The promise of AI is significant, but realising its value comes with challenges. Organizations are eager to innovate, yet must remain vigilant when it comes to data security, governance, and regulatory compliance. As an IT or security leader, you play a critical role in balancing innovation with effective risk management. This means ensuring a secure and compliant environment in which employees can adopt AI responsibly, without exposing the organization to unnecessary risk.

This whitepaper provides a proven path to Copilot readiness. It offers insight into the key challenges and considerations that most organizations face, along with a pragmatic step-by-step approach and practical guidance drawn from real-world experience. With this framework, you can begin taking concrete steps towards enabling AI in your organization with confidence.

Different roles, different perspectives

For the **Chief Information Security Officer** (CISO), data security remains a top priority. The introduction of Copilot and other AI tools presents an opportunity to help the organisation work smarter and more efficiently. At the same time, it raises concerns about potential data breaches and unauthorised access to sensitive information. The key question for the CISO is therefore: how can AI be adopted safely without compromising the security of organisational data?

Security-Analysts are responsible for monitoring and responding to incidents related to AI usage. They need clear visibility into which AI tools are being used within the organisation, what types of data are being shared, and where potential risks may arise. This insight is critical to detect, investigate, and resolve AI-related security incidents effectively.

The **Compliance Officer** focuses on the legal and ethical implications of AI adoption. For this role, it is essential that the use of Copilot and AI agents fully complies with applicable laws and regulations, such as POPIA and GDPR. Compliance leaders must ensure that sensitive or special categories of personal data are not unintentionally processed or exposed through AI tools, and that the organisation can clearly demonstrate adherence to regulatory obligations.

Business users meanwhile, are primarily focused on the opportunities Copilot provides to work faster and more efficiently. However, they may also encounter challenges around the reliability of AI-generated outputs. For example, outdated or poorly governed documents can influence Copilot's responses, leading employees to question whether they can fully trust the information provided. As a result, many business users seek ways to improve data quality and maximise the value of Copilot—sometimes without fully understanding the security or governance implications.

5 key challenges for secure AI

Most organisations have already implemented measures to address privacy, security, and compliance. Policies have been defined, governance frameworks established, and the necessary technology put in place. So why does the introduction of AI still create new challenges? In practice, many organisations encounter five recurring issues:

01 Lack of policy

Many organisations still lack a clear and comprehensive policy governing the use of AI and the handling of data in AI-driven tools. As a result, employees are often unsure about what is permitted and what is not, increasing the risk of unintended misuse. The absence of clear guidelines also makes it more difficult to assess and respond effectively to incidents.

02 Lack of data classification

Data classification is frequently incomplete or inconsistently applied across both new and existing data. When data is not properly classified, sensitive information may unintentionally be accessed or processed by tools such as Copilot, creating unnecessary risk for the organisation.

03 Oversharing

Oversharing remains a widespread challenge. In many organisations, access controls are too broad, giving employees access to more information than they require for their roles. This increases the likelihood that sensitive data may be surfaced or shared through AI tools.

04 Legacy data

Organisations often hold vast amounts of historical data that have accumulated over time. Much of this data is poorly governed or lacks proper classification. Reviewing, cleaning up, and appropriately classifying this legacy data is a significant—but essential—task when preparing for AI adoption.

05 Shadow-AI

Shadow AI has emerged as a growing concern. Employees may use AI tools outside the visibility of IT or security teams, making it difficult to manage data flows, enforce governance policies, and maintain compliance. This significantly increases the risk of data leakage and non-compliant behaviour..

In the following sections, we explore several of these challenges in more detail. The section “Step-by-step plan: Six steps to Copilot readiness” provides practical guidance on how organisations can address these five challenges in a structured and effective way.

Why legacy data can be a problem: A simple example

Legacy data can create an additional challenge for AI adoption: it can distort the output generated by AI tools. The following example illustrates how this works in practice.

Imagine your organisation hosts a company party every year. The event takes place annually, but the timing varies—sometimes in November, sometimes in December. If you ask Copilot the question, "When is the company party?", there is a good chance the response will reference past events rather than the upcoming one.

Why does this happen? Because Copilot finds far more references to company parties from previous years—2020, 2021, 2022, 2023, and 2024—than it does for the upcoming event in 2025. As a result, the abundance of historical data can overshadow the most relevant and current information.

However, if you ask a more specific question such as "When is the company party in 2025?", the likelihood of receiving the correct answer increases significantly. The difference lies in the prompting technique.

This example highlight two important insights:

Legacy data can influence AI outputs. This does not mean Copilot is performing poorly; rather, it reflects the simple reality that organisations typically store far more historical data than recent information.

Employees must learn how to interact with AI effectively. Understanding how AI works and how to ask precise questions is essential to obtaining reliable results and using AI responsibly.

It is important to note that this does not mean organisations must first solve all legacy data challenges before adopting Copilot. In fact, the volume of data organisations generate continues to grow rapidly each year. Waiting until all historical data is fully organised may delay innovation unnecessarily.

From a security and governance perspective, it is often more effective to start by focusing on new data. Classifying newly created information and applying appropriate exclusions helps establish a solid foundation of data hygiene. This approach enables organisations to begin using Copilot safely while progressively improving the governance of legacy data over time.

The shadow AI epidemic

Many organisations underestimate the scale of shadow AI. Although many organisations formally restrict or discourage the use of AI tools, research and practical experience show that nearly 60% of employees still use AI applications, often without the knowledge of IT or security teams.

This phenomenon—essentially a new form of shadow IT that can be described as shadow AI—highlights the limitations of simply trying to prohibit the use of AI. Employees naturally seek ways to work faster and more efficiently. When approved tools are unavailable or difficult to use, they will turn to whatever alternatives they can find.

In practice, we see this pattern repeatedly. At many organisations, employees are not only using tools such as Copilot or ChatGPT, but also a wide range of lesser-known AI applications. In many cases, these tools are completely unknown to IT departments, and there is little or no visibility into how they process or store organisational data.

For this reason, organisations must focus not only on establishing rules, but also on enabling safe and guided adoption of AI. This requires investment in:

- Awareness and training to promote responsible and secure AI usage
- Improved alternatives that employees can use safely and productively
- Monitoring and visibility into the AI tools being used within the organisation
- Guidance and governance, rather than relying solely on restrictive controls

Step-by-step plan: 6 steps to Copilot Readiness

The successful and secure deployment of Copilot requires a balanced approach that combines policy, technology, and organisational adoption. The following steps provide a practical roadmap that enables organisations to introduce AI in a controlled and demonstrably compliant manner.

Step 1

Start with awareness and stakeholder involvement

Begin by building alignment among key stakeholders, including IT, security, compliance, HR, and executive leadership. Organise a kick-off session to discuss both the opportunities and risks associated with Copilot, and jointly define the objectives and guiding principles for adoption.It is important that all stakeholders understand why policy frameworks, data classification, and technical controls are critical to the responsible use of AI.

Step 2

Establish policy and rules

Define a clear AI policy that outlines which AI tools may be used, what types of data can be processed, and which responsibilities and procedures apply in the event of incidents or questions.

This AI policy should extend the organisation's existing data governance policy, which defines the taxonomy used for data classification. Within the AI policy, specify which data classifications or labels may or may not be used in AI-powered tools.

Regulatory frameworks provide an important foundation for this governance. Legislation such as POPIA and Cybercrimes Act (South Africa) and GDPR, NIS2 (EU) require organisations and their suppliers to assess risks and classify information systems appropriately. In practice, this means that policy, classification, and technical safeguards must work together to ensure regulatory compliance.

Step 3

Inventory data, access rights and AI usage

Develop a clear understanding of what data exists within the organisation, where it is stored, and how it is currently being used.

Tools such as Cloud App Discovery can provide insight into the use of AI tools and SaaS applications within the organisation. Microsoft Entra ID can be used to identify active AI agents and review their permissions. In addition, assess SharePoint environments to identify high-risk sites and determine where sensitive information may reside.

This visibility forms the foundation for informed governance decisions.

Step 4

Implement foundational technical controls

Once policies and governance structures are defined, technical safeguards must be implemented accordingly. Examples include:

- Configure SharePoint Advanced Management to exclude high-risk sites from Copilot indexing
- Implementing Data Loss Prevention (DLP's) to prevent sensitive data (such as National Identification Numbers) from being exposed
- Using Microsoft Defender for Cloud Apps to block unauthorised AI tools and monitor the use of approved applications
- Making document classification mandatory across the organisation

A practical starting point is to begin classifying newly created data according to the defined taxonomy. Older, higher-risk data can temporarily be excluded from Copilot indexing until it has been reviewed and classified. It is equally important to communicate clearly to employees what classification labels exist, what they mean, and how they should be applied.

Step 5

Train and support employees

Technology alone is not sufficient for successful AI adoption. Employees must also understand how to use AI responsibly.Provide training and awareness programmes covering topics such as secure AI usage, the importance of data classification, and effective prompting techniques. Ensure that employees know where to go with questions or concerns, and encourage a culture of accountability and responsible AI use.

Step 6

Monitor, evaluate and continuously improve

Finally, establish mechanisms to monitor AI usage and data flows across the organisation. Dashboards and reporting tools can help track the adoption of Copilot and other AI applications.

Regularly collect feedback from users and stakeholders, and periodically evaluate policies, technical controls, and adoption practices. As the organisation matures in its AI governance capabilities, the approach can gradually be expanded to additional departments or broader use cases.

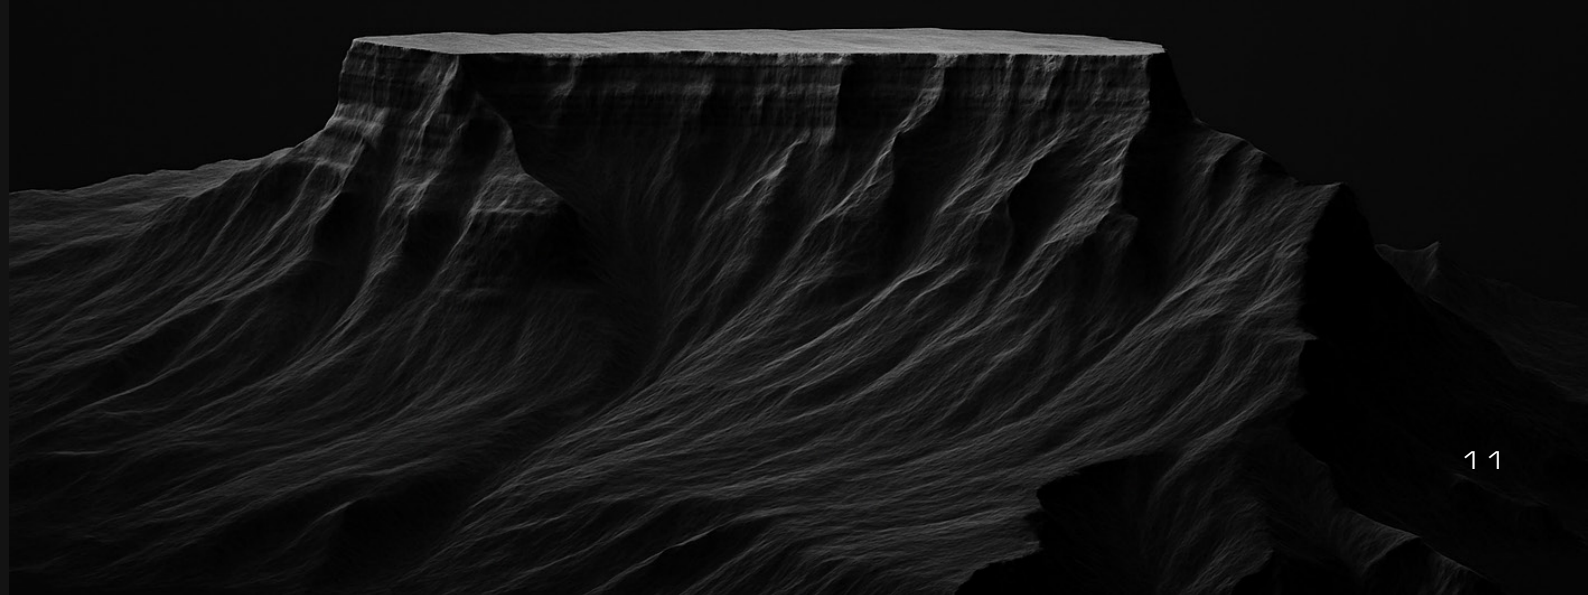

How Nedscaper prepares your organisation for secure AI

Nedscaper helps organisations prepare for secure AI adoption through a structured approach to Copilot readiness. We begin with a comprehensive risk assessment and stakeholder workshop to establish a clear understanding of your organisation's current situation. This is followed by an AI enablement phase, during which policies, data classification, and technical safeguards are implemented in practice.

Within a matter of weeks, your organisation will gain clear visibility into the key risks and opportunities, along with a secure foundation for responsible AI usage.

From there, we support the controlled roll-out of Copilot across the organisation. If required, Nedscaper can also take responsibility for the full integration and lifecycle management of your AI solutions.

Our approach offers several key advantages:

Rapid insight into risks and opportunities
Within the first week, you gain a clear understanding of your current AI readiness and potential quick wins.

A practical, evidence-based roadmap
No abstract theory, but concrete and actionable steps based on proven experience.

Organisation-wide alignment
Driving engagement with key stakeholders across IT and security to HR and executive leadership.

Demonstrable compliance and security
Clear metrics and reporting that provide confidence for auditors, regulators, and leadership teams.

Scalable implementation
Start with a focused pilot, demonstrate value, and expand in a controlled manner.

Microsoft-certified expertise
Our specialists have deep expertise in the Microsoft ecosystem and the technologies that underpin Copilot.

WANT TO KNOW MORE?

Ready to get started with Copilot Readiness?

Book a free consultation and discover how your organisation can safely and sustainably benefit from AI.

Martijn Zantinge

Cybersecurity Architect

+27 87 550 1536
connect@nedscaper.com
nedscaper.com/contact/

About Nedscaper

Nedscaper provides pragmatic, accessible cybersecurity solutions for organizations in Europe and Africa. Our services are built on Microsoft technology and integrate seamlessly with existing systems and processes. With Managed XDR, we offer 24/7 protection: continuous detection of new threats, rapid and effective response, and personalized support to demonstrably improve your security posture.

At Nedscaper, the human factor is central. Our consultants guide clients step by step, simplify complex decisions, and strengthen preventive measures, ensuring that security is effective not only in theory but also in practice. From our offices in Amsterdam, Cape Town, and Johannesburg, and with a growing international network of partners, we help organizations—and the people within them—work with confidence in an increasingly digital world.

In 2025, Microsoft named Nedscaper its Channel Security Partner of the Year: a recognition of our impact, expertise, and scalable approach.