

Generative AI Security: Ensuring a Secure Microsoft Copilot Rollout

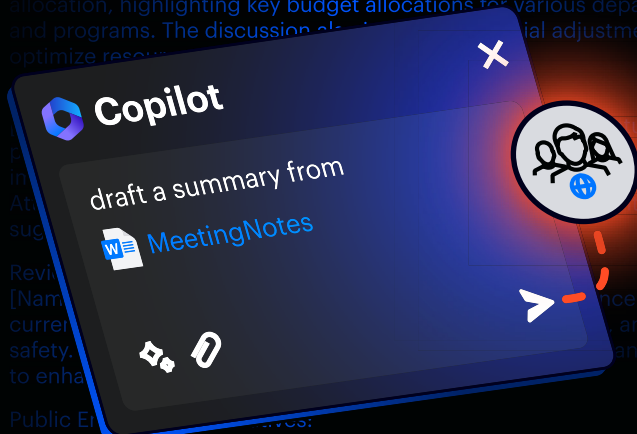
Decoding how Copilot's security model works

The meeting commenced with Cameron Hubbard welcoming all participants and briefly introducing the meeting's purpose.

Updates on City Resource Allocation:
Fahima Morales presented an overview of the current city resource allocation, highlighting key budget allocations for various departments and programs. The discussion also touched upon potential adjustments to optimize resource usage.

4. Review of Current Safety Initiatives:
[Name] presented a report on the current status of city safety initiatives, including the implementation of new security measures and public safety programs. The report highlighted the effectiveness of these initiatives and ways to further enhance public safety.

5. Public Engagement Strategies:
The meeting discussed strategies to increase public engagement and gather feedback from residents regarding city resource allocation. Initiatives such as town hall meetings, surveys, and social media campaigns were proposed to foster better communication and community involvement.





Introduction

Microsoft Copilot has been called one of the most powerful productivity tools on the planet.

Copilot is an AI assistant that lives inside each of your Microsoft 365 apps — Word, Excel, PowerPoint, Teams, Outlook, and so on. Microsoft’s dream is to take the drudgery out of daily work and let humans focus on being creative problem-solvers.

Copilot is a different beast than ChatGPT and other AI tools because it has access to everything you’ve ever worked on in Microsoft 365. Copilot can instantly search and compile data across your documents, presentations, email, calendar, notes, and contacts.

And therein lies risks for information security teams. Copilot can access all the sensitive data that a user can access, which is often **far too much**. Approximately 10% of a company’s M365 data is open to all employees.

Copilot can also rapidly generate net-new sensitive data that must be protected. Prior to the AI revolution, humans’ ability to create and share data far outpaced the capacity to protect it. Just look at data breach trends. Generative AI pours kerosine on this fire.

There is a lot to unpack when it comes to generative AI as a whole: model poisoning, hallucination, deepfakes, etc. In this whitepaper, however, we will focus specifically on data security and how your team can ensure a safe Copilot rollout.

BETTER TOGETHER: VARONIS + MICROSOFT

Varonis and Microsoft forged a new strategic collaboration to help organizations safely harness one of the most powerful productivity tools on the planet — Microsoft Copilot for Microsoft 365.

Together, Varonis and Microsoft are helping organizations confidently roll out AI while continually assessing and improving their Microsoft 365 data security posture behind the scenes before, during, and after deployment. So you can trust your AI rollout is secure and compliant — and stays that way. [Learn more.](#)

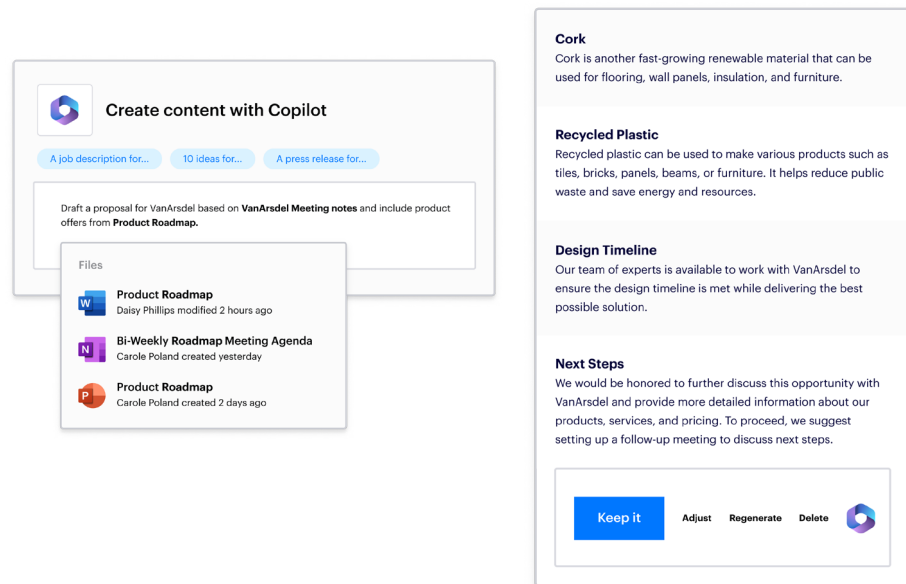


M365 Copilot use cases

The use cases of generative AI with a collaboration suite like M365 are limitless. It's easy to see why so many IT and security teams are clamoring to get early access and preparing their rollout plans. The productivity boosts will be enormous.

For example, you can open a blank Word document and ask Copilot to draft a proposal for a client based on a target data set, including OneNote pages, PowerPoint decks, and other office docs. In a matter of seconds, you have a full-blown proposal.

All it takes is one valid set of stolen credentials, and a hacker can access everything a user can.



HERE ARE A FEW MORE EXAMPLES MICROSOFT GAVE DURING THEIR LAUNCH EVENT:

- + Copilot can join your Teams meetings and summarize in real time what's being discussed, capture action items, and tell you which questions were unresolved in the meeting.
- + Copilot in Outlook can help you triage your inbox, prioritize emails, summarize threads, and generate replies for you.
- + Copilot in Excel can analyze raw data and give you insights, trends, and suggestions.

How M365 Copilot works

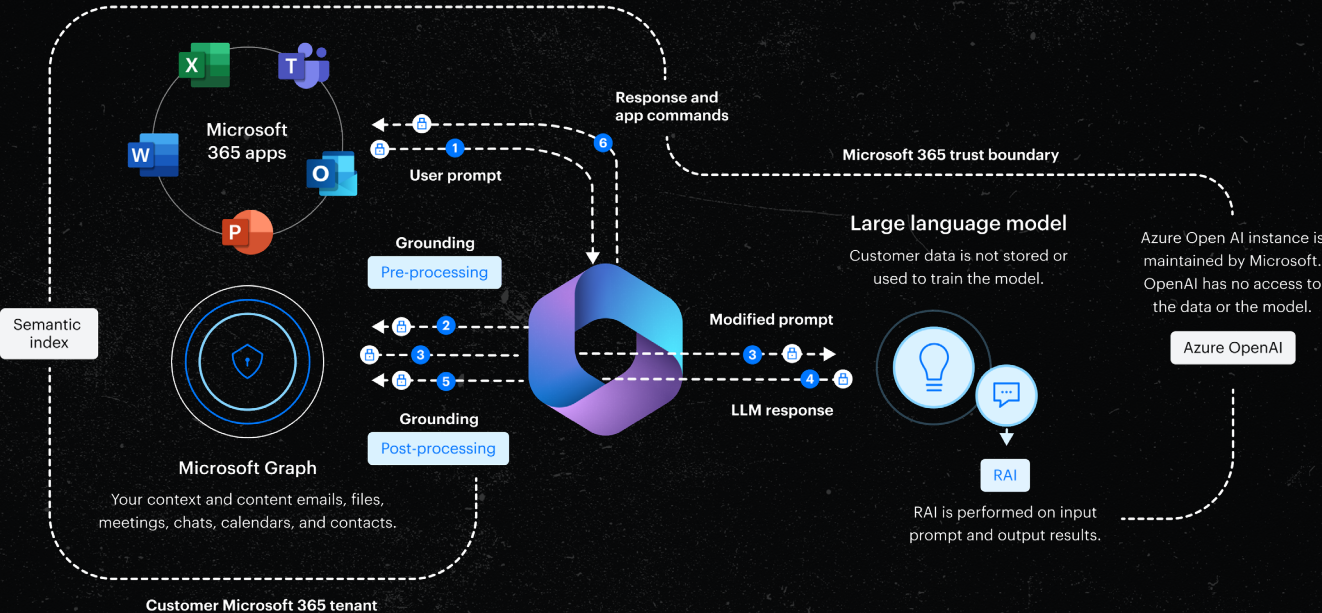
Here's a simple overview of how a copilot prompt is processed:

- + A user inputs a prompt in an app like Word, Outlook, or PowerPoint.
- + Microsoft gathers the user's business context based on their M365 permissions.
- + The prompt is sent to the large language model (like GPT4) to generate a response.
- + Microsoft performs post-processing responsible AI checks.
- + Microsoft generates a response and commands back to the M365 app.

MICROSOFT 365 COPILOT

Data flow (🔒 = all requests are encrypted via HTTPS)

- 1 User prompts from Microsoft 365 Apps are sent to Copilot.
- 2 Copilot accesses Graph and Semantic index for pre-processing.
- 3 Copilot sends modified prompt to large language model (LLM).
- 4 Copilot receives LLM response.
- 5 Copilot accesses Graph and Semantic index for post-processing.
- 6 Copilot sends the response, and app command back to Microsoft 365 apps.





M365 Copilot security model

With collaboration tools, there is always an extreme tension between productivity and security.

This was on display during the coronavirus when IT teams were swiftly deploying Microsoft Teams without first fully understanding how the underlying security model worked or how in-shape their organization's M365 permissions, groups, and link policies were.

WHAT MICROSOFT HANDLES FOR YOU:

- + **Tenant isolation.** Copilot only uses data from the current user's M365 tenant. The AI tool will not surface data from other tenants that the user may be a guest in, nor any tenants that might be set up with cross-tenant sync.
- + **Training boundaries.** Copilot **does not** use any of your business data to train the foundational LLMs that Copilot uses for all tenants. You shouldn't have to worry about your proprietary data showing up in responses to other users in other tenants.

WHAT YOU NEED TO MANAGE:

- + **Permissions.** Copilot surfaces all organizational data to which individual users have at least view permissions.
- + **Labels.** Copilot security relies heavily on accurate and current labels. Manual labeling is error-prone and automated labeling is only as good as the classification engine that powers it.
- + **Humans.** Copilot's responses aren't guaranteed to be 100% factual or safe; humans must take responsibility for reviewing AI-generated content.

Let's take what you need to manage one by one.



Permissions


Granting Copilot access to only what a user can access would be an excellent idea if companies were able to easily enforce least privilege in Microsoft 365.

Microsoft states in its [Copilot data security documentation](#):

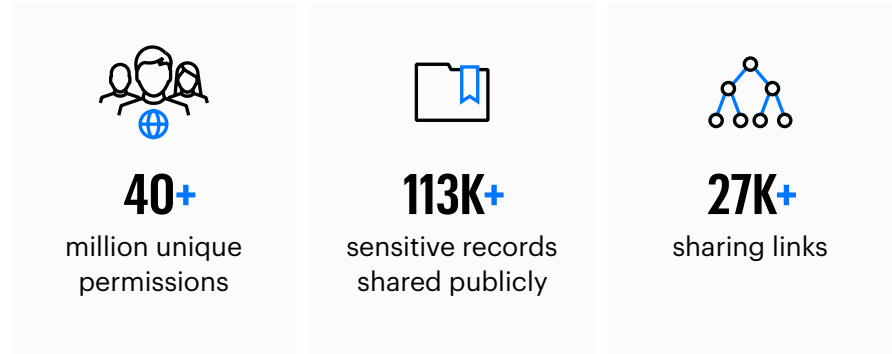
“It’s important that you’re using the permission models available in Microsoft 365 services, such as SharePoint, to help ensure the right users or groups have the right access to the right content within your organization.”

We know empirically, however, that most organizations are about as far from least privilege as they can be. Just take a look at some of the stats from Microsoft’s own [State of Cloud Permissions Risk report](#).

- + Multi-cloud environments are complex. There are 40,000+ permissions to manage and over 50% are high-risk.
- + After analyzing over 500 risk assessments, Microsoft found that most identities are greatly over-permissioned, putting organizations’ critical environments at risk for accidental or malicious permission misuse.
- + Workload identities accessing cloud environments are increasing, now outnumbering human identities 10:1.
- + 1% of permissions granted are actually used.
- + Less than 50% of identities are super admins, meaning they have access to all permissions and resources.




This picture matches what Varonis sees when we perform thousands of Data Risk Assessments for companies using Microsoft 365 each year. In our report, The Great SaaS Data Exposure, we found that the average M365 tenant has:



Why does this happen? Microsoft 365 permissions are highly complex. Just think about all the ways in which a user can gain access to data:

- + Direct user permissions
- + Microsoft 365 group permissions
- + SharePoint local permissions (with custom levels)
- + Link access (anyone, org-wide, direct, guest)
- + External access
- + Public access
- + Guest access

To make matters worse, permissions are mostly in the hands of end users, not IT or security teams.



Labels

Microsoft relies heavily on sensitivity labels to enforce DLP policies, apply encryption, and broadly prevent data leaks. In practice, however, getting labels to work is difficult, especially if you rely on humans to apply sensitivity labels.

Microsoft paints a rosy picture of labeling and blocking as the ultimate safety net for your data. Reality reveals a bleaker scenario. As humans create data, labeling frequently lags behind or becomes outdated.

Blocking or encrypting data can add friction to workflows, and labeling technologies are limited to specific file types. The more labels an organization has, the more confusing it can become for users. This is especially intense for larger organizations.

The efficacy of label-based data protection will surely degrade when we have AI generating orders of magnitude more data requiring accurate and auto-updating labels.

ARE MY LABELS OKAY?

Varonis can validate and improve an organization's Microsoft sensitivity labeling by scanning, discovering, and fixing:

- + Sensitive files without a label
- + Sensitive files with an incorrect label
- + Non-sensitive files with a sensitive label

Humans

AI can make humans lazy. Content generated by LLMs like GPT4 is not just good, it's great. In many cases, the speed and the quality far surpass what a human can do. As a result, people start to blindly trust AI to create safe and accurate responses.

We have already seen real-world scenarios in which Copilot drafts a proposal for a client and includes sensitive data belonging to a completely different client. The user hits "send" after a quick glance (or no glance), and now you have a privacy or data breach scenario on your hands.



Getting your tenant security-ready for Copilot

It's critical to have a sense of your data security posture before your Copilot rollout. With Copilot now available to enterprises, now is a great time to get your security controls in place.

Varonis protects thousands of Microsoft 365 customers with our Data Security Platform, which provides a real-time view of risk and the ability to automatically enforce least privilege.

We can help you address the biggest security risks with Copilot with virtually no manual effort. With [Varonis for Microsoft 365](#), you can:

- + Automatically discover and classify all sensitive AI-generated content.
- + Automatically ensure that MPIP labels are correctly applied.
- + Automatically enforce least privilege permissions.
- + Continuously monitor sensitive data in M365 and alert and respond to abnormal behavior.

Ready to experience the Varonis difference?

Reduce your risk without taking any. Contact our team to learn what will be covered in your **free** Data Risk Assessment.

[Contact us](#)

ABOUT VARONIS

Varonis (Nasdaq: VRNS) is a leader in data security, fighting a different battle than conventional cybersecurity companies. Our cloud-native Data Security Platform continuously discovers and classifies critical data, removes exposures, and detects advanced threats with AI-powered automation.

Thousands of organizations worldwide trust Varonis to defend their data wherever it lives — across SaaS, IaaS, and hybrid cloud environments. Customers use Varonis to automate a wide range of security outcomes, including data security posture management (DSPM), data classification, data access governance (DAG), data detection and response (DDR), data loss prevention (DLP), and insider risk management.

Varonis protects data first, not last. Learn more at www.varonis.com.



