E-BOOK

proofpoint.

# Deploying Microsoft 365 Copilot Securely

## INTRODUCTION

# Microsoft 365 Copilot: transformative capabilities, serious data risks

In the span of just a few years, generative AI (GenAI) has established itself as a transformative technology. Awakened to GenAI's potential by the 2022 launch of ChatGPT—and driven by the promise of business transformation and productivity gains—organizations across all industries have rapidly integrated GenAI tools into their business workflows. Microsoft 365 Copilot has become one of the most widely deployed enterprise AI tools.

By embedding GenAI directly into Microsoft applications such as Word, Outlook, Excel, PowerPoint, Teams, and SharePoint, Copilot can accelerate and streamline how employees create content, analyze data, and collaborate. Instead of switching between applications or manually searching for information, users can now use Copilot to quickly synthesize and surface organizational knowledge.

However, these capabilities also introduce new security, privacy, and governance challenges. Copilot is not a standalone application; it's deeply integrated in an organization's Microsoft 365 environment. That means Copilot's outputs are only as secure as the data it can access. If sensitive information is over-shared, misclassified, or poorly governed, Copilot can surface that data in ways that increase exposure and risk.

For many organizations, the challenge is not whether to adopt Copilot, but how to do so safely and responsibly. Leaders must balance productivity gains with the need to protect intellectual property, customer data, and regulated information. Traditional security controls designed for static files and human-initiated access are insufficient when applied to AI systems that dynamically retrieve, summarize, and transform data at scale.

Recognizing these challenges, this eBook explores how to deploy Microsoft 365 Copilot securely in your enterprise.

**You'll learn:**

- The most common data security risks organizations face when enabling Copilot
- Best practices for Copilot data governance, visibility, and control
- How Proofpoint helps organizations strengthen their data security posture for Copilot and continuously monitor Copilot usage to prevent data leakage, misuse, and compliance violations

CHAPTER ONE

# Data security risks in Copilot deployments

For organizations that use Microsoft's enterprise software suite, Copilot can be a significant business enabler. However, Copilot also transforms how enterprise data is accessed and consumed. Instead of users manually opening files or emails, Copilot retrieves and aggregates information to generate responses. While this enables powerful new workflows, it also magnifies weaknesses in data security, permissions, and governance. This section describes some of the primary data security risks that can arise in poorly secured Copilot deployments.

## Over-permissioned access

Copilot respects Microsoft 365 permissions, but it doesn't question whether those permissions are appropriate. If a user has access to sensitive SharePoint sites, OneDrive files, or Teams channels—intentionally or not—Copilot can include that information in generated responses. In many organizations, permissions have formed over years of collaboration, mergers or acquisitions, and reorganizations. Users might retain access to projects long after their involvement ends. Copilot amplifies this problem by making it easier to retrieve and summarize sensitive information that users technically have access to but should not be using.

## Over-shared files and folders

Related to over-permissioned access, broad sharing settings are another common source of Copilot exposure. Files or folders shared with the *Anyone with the link* or *Anyone in the organization* options are especially risky. While these settings are intended to simplify collaboration, they dramatically increase Copilot's retrieval surface.

When Copilot indexes broadly shared content, it can surface sensitive information to users that were never intended to see it. Content that previously required deliberate searching can now appear instantly in a Copilot-generated summary, increasing the chances of accidental disclosure.

## Data misclassification

Copilot relies on existing data labels and classifications to apply granular usage rights beyond basic access controls. Unfortunately, many organizations struggle with inconsistent or incomplete data classification. Sensitive documents can lack appropriate labels, while others might be mislabeled.

Without accurate classification, Copilot can't reliably distinguish between low- and high-risk content. This increases the likelihood that confidential data, such as financial forecasts, legal documents, or personal information, will be included in AI-generated responses.

## Indirect prompt injection attacks

Email remains one of the most common attack vectors, and Copilot's integration with email tools such as Outlook introduces new risks. Threat actors can execute indirect prompt-injection attacks by embedding manipulative prompts in emails. These hidden prompts can instruct Copilot to perform insecure actions, retrieve sensitive data, or ignore safeguards. This type of zero-click exploit can occur without a user ever interacting with the malicious email.

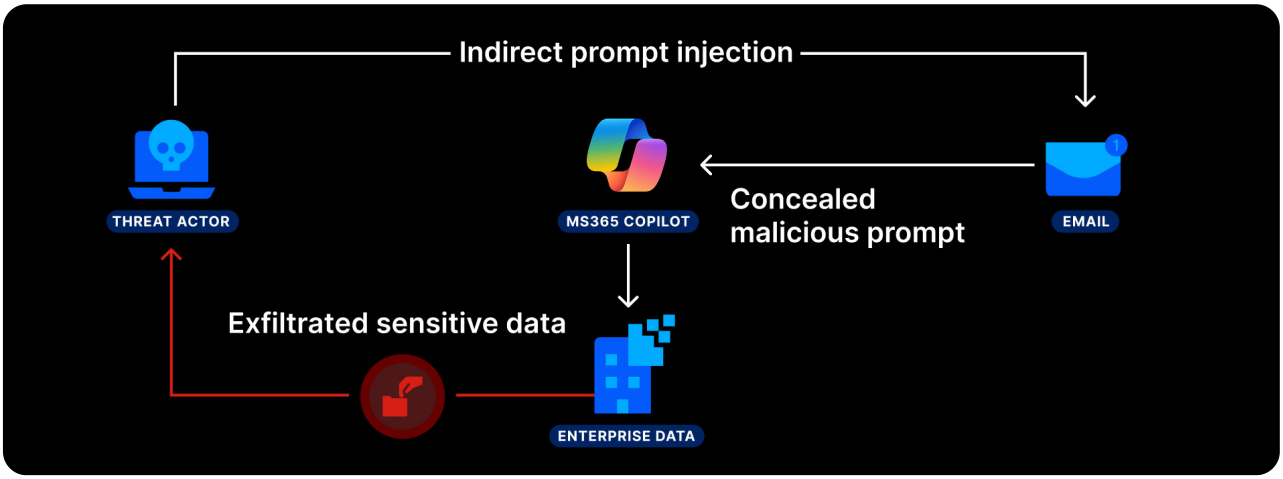


**Figure 1:** Indirect prompt-injection attacks can exploit the integration of Microsoft 365 Copilot and email tools such as Microsoft Outlook.

Because indirect prompt-injection attacks use natural language rather than traditional malware, they're harder to detect using legacy security tools. As AI becomes more deeply integrated in email workflows, organizations must rethink how they protect against these types of AI-assisted threats.

## Incomplete user training

Uneven levels of AI competency across an organization are another major risk factor. Improperly trained users might assume that Copilot will automatically block sensitive or unauthorized data. In reality, Copilot responds to prompts based on available permissions and data context. Users who don't understand these limitations might inadvertently request or submit sensitive information.

Without proper training, employees might treat Copilot as a trusted authority rather than a productivity assistant. This increases the risk of data misuse and compliance violations.

## Data theft by malicious insiders

AI tools such as Copilot can also be misused intentionally. Malicious insiders can attempt to exfiltrate sensitive data by using Copilot to summarize large datasets, reformat information, or transform content in ways that bypass security controls. Because Copilot interactions are conversational, malicious behavior can be harder to detect than conventional data theft.

## Incident response and visibility gaps

When sensitive data is processed by an AI tool, organizations often lack the visibility needed to respond quickly. Tracing what data Copilot accessed, which prompts were used, and who received the output can be difficult without specialized monitoring.

This lack of visibility affects incident response and governance. It also increases the risk of regulatory non-compliance. Organizations might struggle to demonstrate appropriate controls under regulations such as the General Data Protection Regulation (GDPR) and Health Insurance Portability and Accountability Act (HIPAA). Underlining these concerns, a 2025 Gartner research report revealed that 47% of IT leaders have low confidence in their ability to manage Copilot's security and access risks.[1]

**47%**

of IT leaders report they are either not very confident or have no confidence at all in their ability to manage Copilot's security and access risks.

*Source: Gartner*

1. Gartner. *How to Secure and Govern Microsoft 365 Copilot at Scale*. January 2025.

CHAPTER TWO

# Best practices for governing Copilot

Secure Copilot deployments require a proactive, structured governance approach. Organizations must view Copilot as a new data access layer that demands increased oversight. This section highlights some key best practices for governing Copilot deployments.

## Audit and refine permissions

At the foundation of Copilot governance is strong identity and access management. Organizations must audit permissions across SharePoint, OneDrive, Teams, and Exchange to identify excessive or outdated access. Enforcing role-based access control (RBAC) and least-privilege principles reduces the amount of sensitive data Copilot can retrieve.

## Manage permitted SharePoint sites

It's possible to restrict which SharePoint sites Copilot can index. This is especially valuable during phased rollouts, when Copilot will support highly regulated teams such as Legal, Finance, or Mergers & Acquisitions, or when there are particular concerns about exposing intellectual property (IP). By limiting Copilot's scope, organizations can reduce risk while still enabling productivity gains.

## Discover and classify sensitive data

Effective governance requires knowing where sensitive data lives. Data security posture management (DSPM) tools can help organizations discover structured and unstructured data across the enterprise and classify it based on sensitivity, privacy, and compliance requirements. Full-featured DSPM tools can also automatically apply Microsoft Information Protection (MIP) labels. Organizations should then enforce these data classifications through data loss prevention (DLP) policies.

## Classify and label Copilot output

Governance must extend beyond source data to also cover Copilot-generated output. Applying sensitivity labels to AI-generated content ensures that it inherits the same security posture as the underlying source content. This prevents sensitive information from being inadvertently shared or stored insecurely.

## Establish automated governance

Manual processes don't scale for AI. Organizations must define clear policies for Copilot usage, permissible data types, and remediation workflows—and automate these wherever possible. Alignment with legal, HR, and compliance teams can ensure policies are enforceable and defensible.

## Monitor access and usage

A dedicated AI data governance tool can provide visibility into Copilot prompts, file uploads, data access, and outputs. Monitoring should focus on sensitive data usage, unusual behavior, and signs of malicious intent. Alerts and investigations should be prioritized based on risk.

## Educate users on responsible AI use

Finally, governance must include people. Training programs should teach users how Copilot works, how to design safe prompts, and how to recognize sensitive data. Reinforcing that Copilot is a productivity aid, not an authority, helps set appropriate expectations and reduce risk.

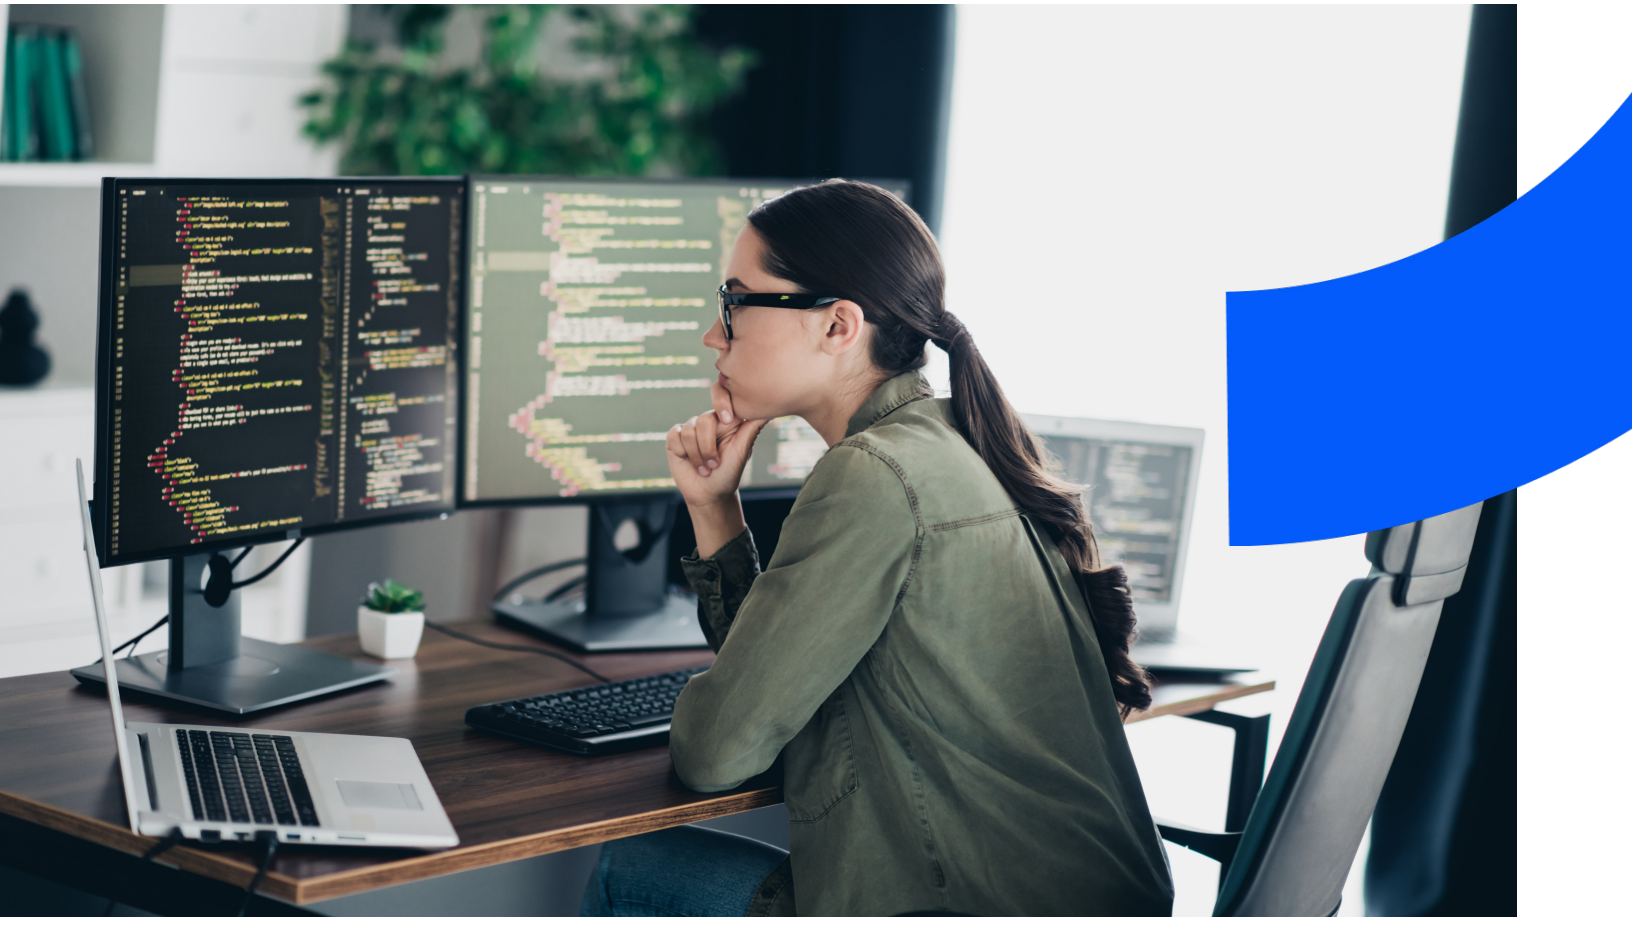

CHAPTER THREE

# How Proofpoint secures Copilot deployments

Proofpoint's AI-ready data security solution secures Copilot deployments by strengthening your organization's data security posture and providing continuous, intelligent monitoring to prevent sensitive data leakage. This section describes the Copilot-ready capabilities of our unified data security solution.

## Strengthen data security posture for Copilot

### Classify sensitive data

Proofpoint enables organizations to classify sensitive data across SharePoint, OneDrive, and Teams using Autonomous Custom Classifiers that have over 95% accuracy. These AI-powered classifiers identify business-critical documents, uncover unknown risks, and integrate seamlessly with MIP labeling.

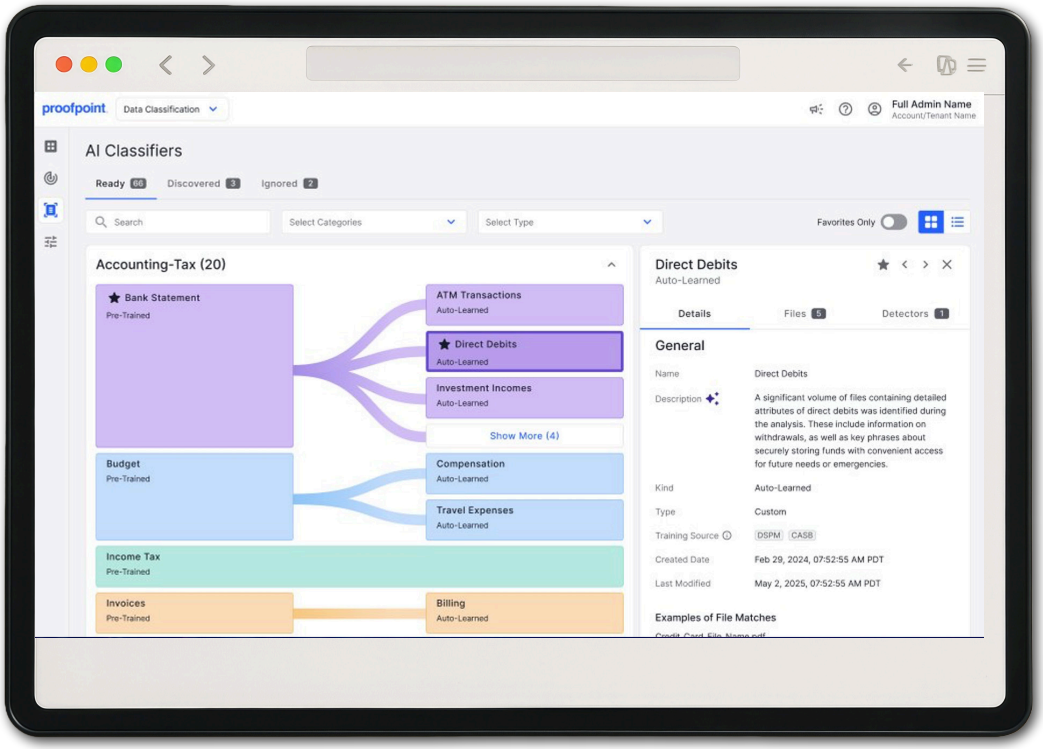


**Figure 2:** Proofpoint Autonomous Custom Classifiers discover and classify sensitive data across your Microsoft 365 environment with over 95% accuracy.

### Identify overexposed data

Proofpoint identifies files and folders that are over-shared. That includes those accessible via the *Anyone with the link* or *Anyone in the organization* options or those shared with overly permissive groups. Real-time risk prioritization highlights exposed, high-value data so teams can act quickly.

### Apply one-click DLP policies and remediation actions

With Proofpoint, organizations can apply DLP policies for Microsoft 365 with minimal effort. Risky access can be revoked, file-sharing permissions adjusted, and exposed files quarantined. Proofpoint also flags unconfigured SharePoint site policies, missing sensitivity labels, and other governance gaps.

## Governing data access

Proofpoint's data risk map visualizes where sensitive data lives, how it moves, and who accesses it. This helps security teams govern and mitigate access, exfiltration, and configuration risks right across the enterprise environment.

Proofpoint also provides workflows to remediate excessive file-sharing permissions, including reviews by administrators or file owners. Access reviews can be delegated, reminders automated, and progress tracked and reported over time, making governance sustainable.

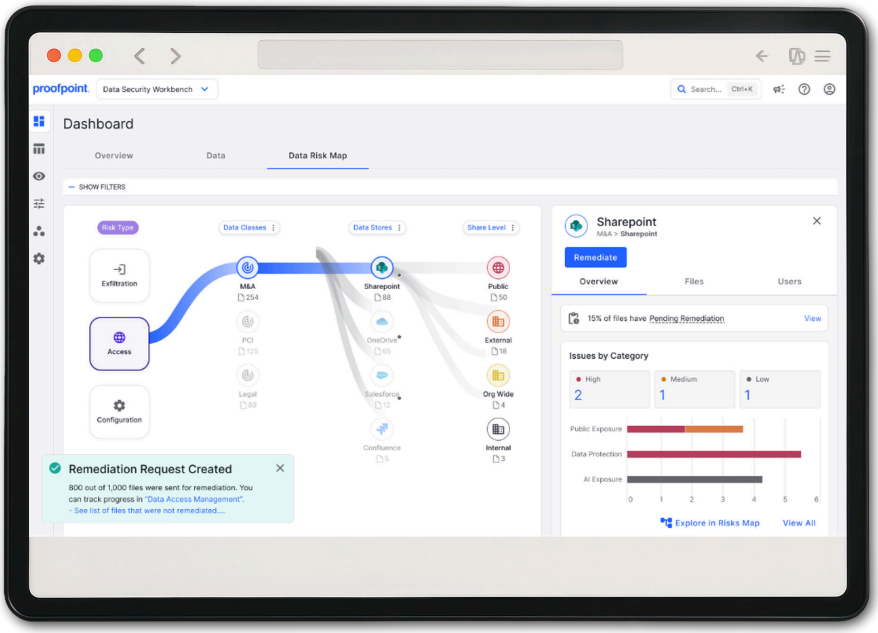


**Figure 3:** Proofpoint's data risk map visualizes access, exfiltration, and configuration risks across the enterprise environment.

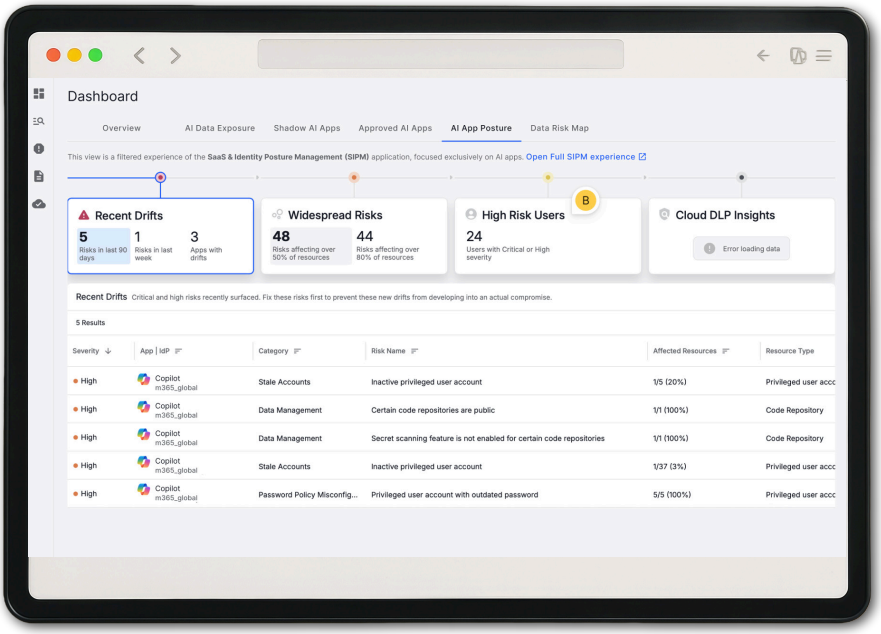


**Figure 4:** Proofpoint continuously monitors Microsoft 365 and Copilot configurations for risks that could lead to data exposure.

## Monitoring for Copilot misuse

Proofpoint continuously monitors Copilot activity for signs of misuse, including sensitive data in prompts, file uploads, or responses; malicious prompt patterns; and unusual volumes of activity. Context such as network location further enhances detection accuracy.

## Continuous configuration monitoring

Beyond Copilot itself, Proofpoint monitors Microsoft 365 and Copilot configurations for risks that could lead to data exposure. This includes excessive SharePoint site-level permissions, missing DLP or retention policies, and unsanctioned third-party AI applications.

## CONCLUSION

# Take the right steps to secure Copilot

- To see our unified, AI-ready data security solution in action, **request a free demo.**
- To learn more about how Proofpoint is leading the way in data security designed for the AI age, join us at one of our **Protect Series events.**

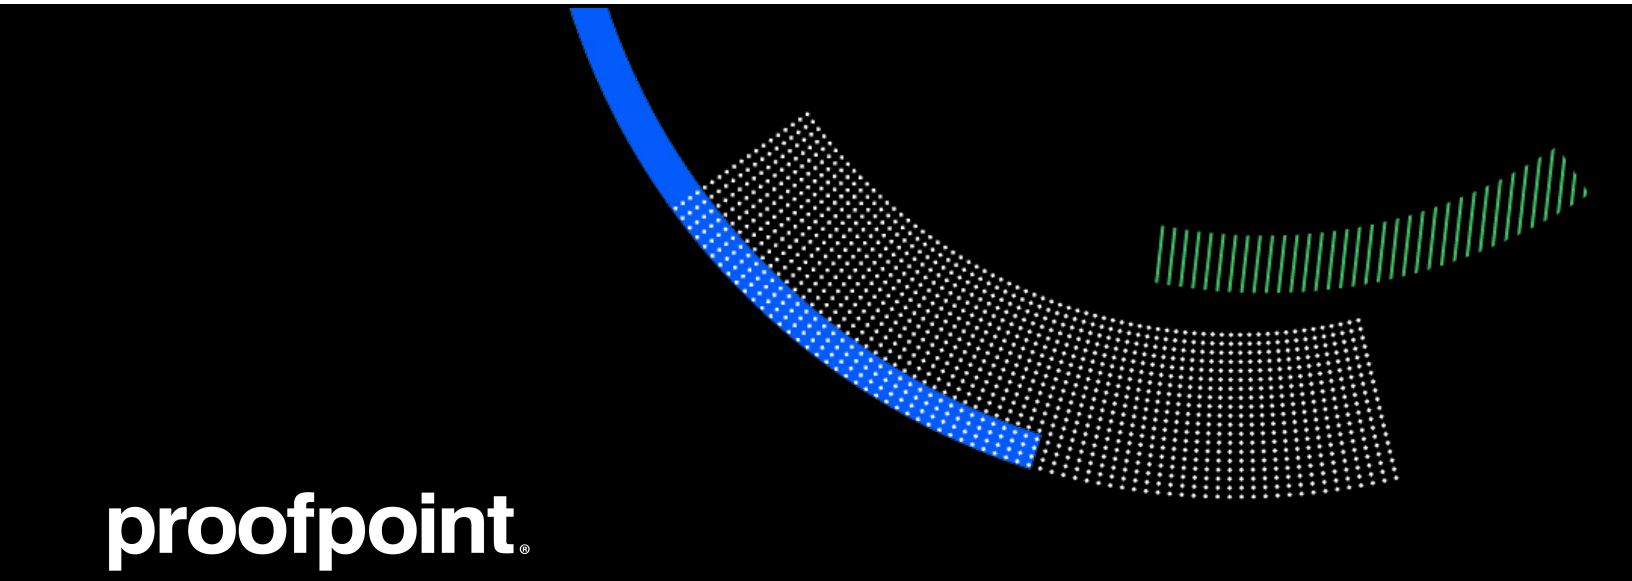

proofpoint.

**About Proofpoint, Inc.** Proofpoint, Inc. is a global leader in human- and agent-centric cybersecurity, securing how people, data and AI agents connect across email, cloud and collaboration tools. Proofpoint is a trusted partner to over 80 of the Fortune 100, over 10,000 large enterprises, and millions of smaller organizations in stopping threats, preventing data loss, and building resilience across people and AI workflows. Proofpoint's collaboration and data security platform helps organizations of all sizes protect and empower their people while embracing AI securely and confidently. Learn more at www.proofpoint.com.

**Connect with Proofpoint:** LinkedIn