GUIDE

Why Data Security Will Make or Break Copilot Success

What every organization needs to know about data security related to their Microsoft 365 Copilot rollout

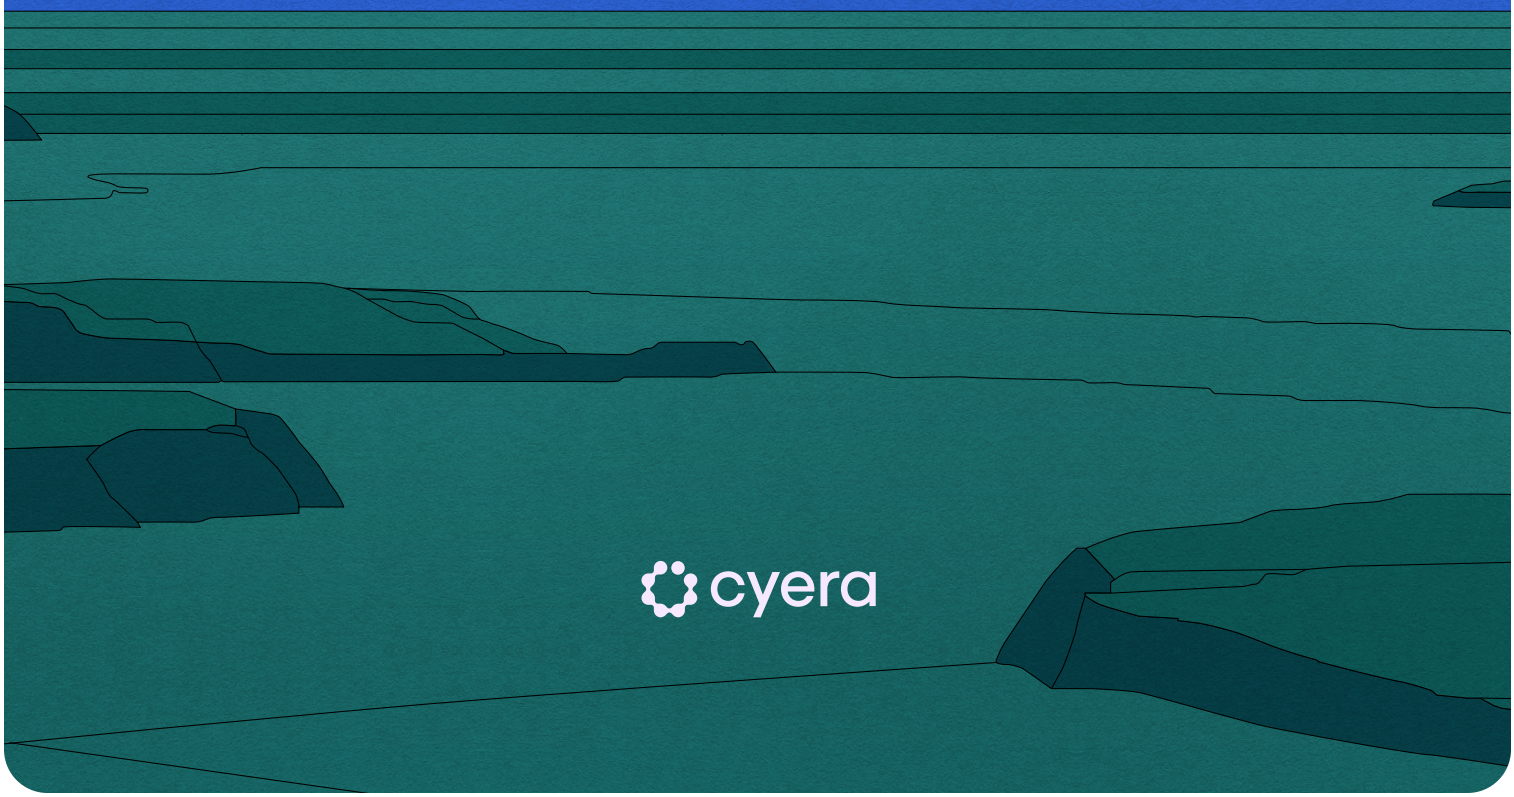

cyera

Executive Summary

Microsoft Copilot doesn't just search files. It reasons across them at machine speed, connects dots between documents, and surfaces answers from data your team didn't know was exposed.

> The problem isn't Copilot. It's what Copilot can reach.

M365 has strong governance built in. But covering your full data estate requires extending those foundations. Sensitivity labels, the basis of Microsoft's Copilot governance model, don't apply to 40% of enterprise file types. CSVs, ZIP archives, images, source code. Copilot reads all of them. Labels never touch them.

Cyera closes that gap. AI-powered classification covers your full data estate. Omni DLP protects every file type, labeled or not. The result is 100% coverage. Not most files. All of them.

Data Risk in M365 Copilot

Across hundreds of enterprise M365 environments, Cyera consistently uncovers the same red flags when it comes to data. These aren't edge cases.

23K+

sensitive files accessible to every employee by default

83%

of sensitivity labels are wrong or under-labeled

40%

of M365 file types structurally cannot receive sensitivity labels

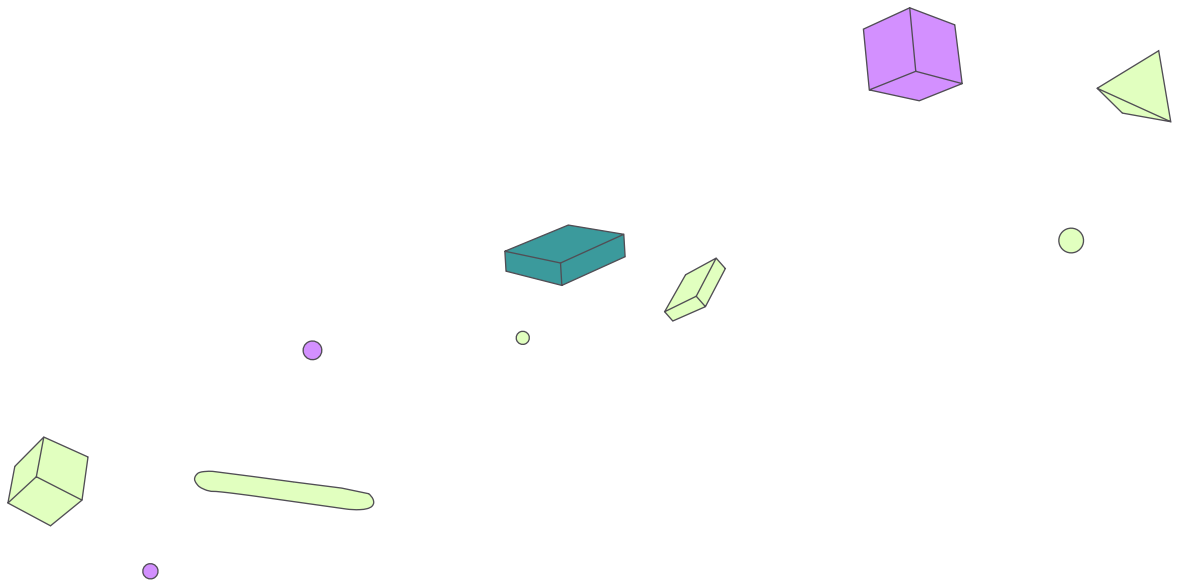

Why Labels Aren't Enough

Microsoft provides strong guidance on Copilot security with recommendations on data governance, security controls, and compliance posture. Much of this relies on Microsoft Information Protection (MIP) sensitivity labels to govern what Copilot can access and what DLP policies apply in Purview. However, MIP labels have a structural limitation nobody talks about openly: they don't work on approximately 40% of file types within the enterprise.

> When MIP labels aren't applied to the data, Copilot can read all of it

Time is another issue. Manual classification and labeling at enterprise scale takes 6 to 12 months. But the business and Copilot can't wait. This results in organizations deploying Copilot and exposing it to data estates that simply aren't ready for AI.

Lastly, even when labels do exist, 83% of labels are wrong. That's because manual labeling is inconsistent. So when an HR spreadsheet isn't protected or an IP document isn't restricted, Copilot can access that data and make decisions. Incorrect labels introduce data risk even though the AI was working as intended.

Three Gaps That Break Copilot Security

Understanding these gaps is the difference between a compliance checkbox and a defensible security posture.

Gap 1: Coverage

40% of your M365 files can't receive a sensitivity label. No label means no DLP policy. No DLP policy means no Copilot restriction. A CSV export of your customer database. A ZIP of your source code. A PNG of a signed contract. All of them reachable. None of them governed.

This isn't a configuration issue. It's a structural limitation of how sensitivity labeling works. No policy fixes it.

Gap 2: Accuracy

Of the labels that exist, 83% are wrong. Most are under-labeled — assigned a classification that understates the sensitivity of the data. Your DLP policies are drawn against an inaccurate map. That's arguably more dangerous than having no labels at all.

Manual labeling is why. At enterprise scale, human classification introduces error and inconsistency that compounds over time. AI fixes the root cause.

Gap 3: Speed

Microsoft's recommended path to readiness takes 6 to 12 months. That's a defensible timeline for a deliberate data governance program. But not when it comes to enterprise AI adoption.

Security teams can't halt Copilot rollouts. They can only hold them off for so long. Chances are Copilot is already live in production reading files that have never been classified or labeled, serving answers to questions your security team hasn't anticipated, and doing all of it at a speed no human audit process can match.

Copilot is live now. Your readiness program is not.

What 'Copilot Risk' Actually Means

Copilot risk isn't a checklist but rather a security posture that can be measured, benchmarked, and continuously monitored.

Microsoft defines minimum thresholds for safe Copilot deployment. Your Microsoft Secure Score should be 75 or higher. At least 95% of your data should be classified. And 100% of sensitive data should be covered by active DLP policies.

Most organizations aren't close to those benchmarks. And it comes down to a measurement problem. You can't govern what you can't see. You can't label what you don't know exists. And you can't protect data that your tools structurally cannot reach.

Addressing your Copilot risk means four things: your data estate is fully visible, your sensitive data is accurately classified, your DLP coverage extends to every file type regardless of label status, and your identity access model reflects who actually needs to reach what.

The Cyera Copilot Risk Report

Cyera offers a Copilot Risk Report that gives you a verified risk score before you flip the switch. No agents. No production impact. Results in days. Full readiness in approximately 8 weeks versus 6 to 12 months using Microsoft-native tools alone.

The methodology follows five phases:

SEE	Cyera scans your entire M365 environment and creates a complete data inventory including where your data lives, which platforms and databases it's concentrated in, and who has access to it. This isn't a sample. It's your actual data estate, at petabyte scale. Most organizations find 23,000 or more sensitive files accessible to every employee on day one.
CLASSIFY	Cyera's AI classification engine ensures coverage of every file type, including the 40% that Microsoft Purview labels can never reach. Classification is dependent on business context, not file name. And this happens with 95%+ accuracy with no manual rules. Manual classification is why so much enterprise data stays unlabeled. AI fixes the root cause at scale.
LABEL	Accurate MIP labels are applied at scale through the Microsoft Graph API. Cyera does this in weeks, not months. Cyera also corrects the 83% of existing labels that are wrong today, so your DLP policies, encryption, and Copilot access restrictions apply to the right files.
PROTECT	Cyera's Omni DLP covers every file type, labeled or not, at rest and in motion. Labeling and Omni DLP ensure complete coverage so you can govern and protect your data accordingly.
MONITOR	Runtime controls continuously track what Copilot accesses and enforce policy in real time. Stale accounts and over-permissioned identities are flagged automatically. Copilot activity is auditable from day one.

What the Copilot Risk Report Delivers

The report is structured around nine analytical components, with each one tied directly to a remediation outcome, not just a finding.

1. Data Landscape Overview

A complete inventory of your M365 data estate, mapped to platforms and databases. This ensures you have a full picture of Copilot's access to data that you previously didn't have visibility into.

2. Current Security Posture

Classification coverage, DLP policy coverage, and your Microsoft Secure Score — benchmarked against the thresholds Microsoft defines for safe Copilot deployment (95% classified, 100% sensitive data covered by DLP, Secure Score ≥ 75).

3. Gap Analysis Framework

Progress against required security controls, with control domain, current state, remaining gap, and risk level. Includes a department-level view of which teams are accessing which categories of sensitive data. It surfaces questions like 'why does Customer Support have access to Intellectual Property?'

4. Identity Risk Analysis

Correlates identities with the data they can access. Shows identity type (human, service, external, admin), records accessible, data category distribution, and last active date. This allows you to pinpoint stale or over-permissioned accounts before Copilot amplifies the risk.

5. Critical Risk Scenarios

Highlights real-world risks specific to your environment — data leakage, regulatory non-compliance — with specific records at risk and remediation paths. Not generic threat categories. Your actual exposure.

6. Cyera-Accelerated Readiness

Side-by-side comparison of your time to readiness using Microsoft native tools alone versus Cyera. Configurable by enterprise size and license type (E3 or E5).

7. ROI Breakdown

A customized ROI analysis for your environment. Forrester research shows a 25,000-employee organization can expect 116% ROI from Copilot deployments, equivalent to a net present value of nearly $20 million dollars. Every month of delayed readiness is a month of deferred value.

8. Success Metrics & KPIs

Baseline values and targets for each key readiness metric, with automated weekly progress tracking through a Cyera dashboard.

9. Executive Recommendations

A prioritized remediation list with timelines and milestones that are structured for both board-level reporting and CISO accountability.

Risk Remediation Simulator

Every critical security gap — enforcing MFA, removing stale external users, removing excessive permissions, deploying sensitivity labels — is quantified by its impact on your overall Readiness Score. As you resolve each one, your score updates in real time. Remediation becomes a managed program, not a wishlist.

Why Cyera

With Cyera you can lower your Copilot risk in 8 weeks. AI classification covers your entire data estate in the first few weeks. Automated label application follows. Omni DLP closes the coverage gap. And runtime controls are in place before most traditional readiness programs have finished their discovery phase.

Every week your Copilot deployment is delayed is a cost to the business. Using only native tools takes 6 to 12 months of manual classification, phased rollout, and policy configuration to reduce data risk. If Copilot is already live, you're living with unnecessary and unknown risk throughout that period.

6-12 mo

How long it takes to address Copilot Risk without Cyera and via manual classification

~8 wks

How long it takes to address Copilot Risk with Cyera and AI-powered classification

116%

ROI from Copilot successful deployment (Forrester, 25K-employee organization)

Get Your Score for Free

Deploying Copilot without knowing what it can reach isn't a calculated risk. It's an uninformed one.

The data is already in your environment. The gaps are already there. The question is whether you find them before Copilot does.

Cyera's Copilot Risk Report gives you a verified risk score in days. No agents. No production impact. No six-month readiness program standing between your organization and the productivity gains you've already paid Microsoft to deliver.

Request Your Free Copilot Risk Report

See your actual risk score. Find your gaps. Get to readiness in ~8 weeks.

Learn more at www.cyera.com

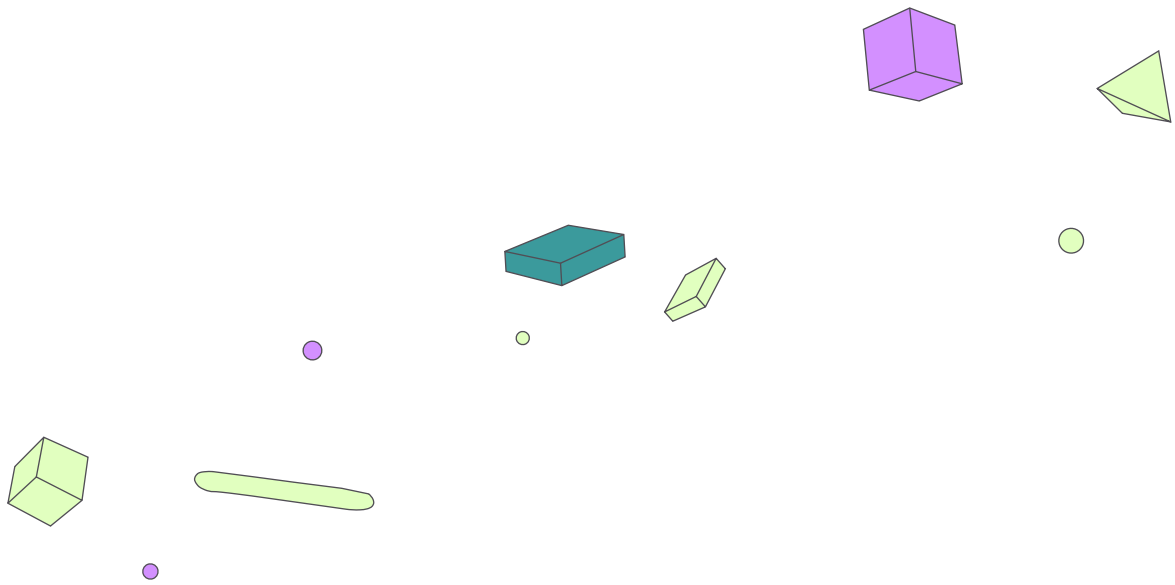

Data in this guide is based on hundreds of enterprise M365 environments assessed by Cyera.

Forrester ROI data sourced from Forrester Research, 2024.