Quick Guide

Test Prompts for Assessing Copilot Oversharing Risk

Discover what sensitive data might be accidentally exposed

opsin

Test Your Environment for Real-World Data Exposure Risks

As organizations rush to deploy Microsoft Copilot to boost productivity, many CISOs are discovering an uncomfortable truth: their years of data sharing practices have created significant security blind spots. What once required employees to actively search for documents is now instantly accessible through AI-powered queries.

This guide helps you understand why oversharing happens with Microsoft Copilot and provides practical test prompts to assess your organization's current risk level.

Why Oversharing Happens with Microsoft Copilot

How Microsoft Copilot Works Behind the Scenes

Microsoft Copilot uses a Retrieval-Augmented Generation (RAG) model that combines three key components:

Data Connection: Copilot plugs into your organization's ecosystem — emails, SharePoint, OneDrive, Teams, CRMs and more, indexing all accessible content.

Retrieval Layer: When you ask a question, Copilot identifies relevant data sources across all connected systems and fetches information based on semantic relevance, not just keywords.

Generative AI: It processes retrieved data to generate responses that synthesize information from multiple sources, departments, and time periods.

Think of it as the ultimate search engine that has read every document in your organization and can instantly recall and combine information to answer any question.

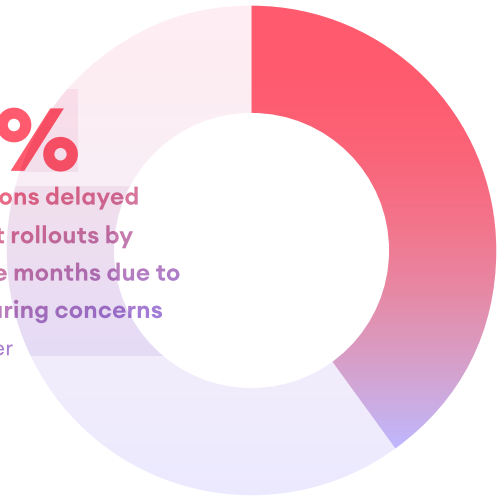

40%

of organizations delayed M365 Copilot rollouts by three or more months due to data oversharing concerns

Source: Gartner

Why This Creates Oversharing Risk

For years, organizations shared content broadly for collaboration. Employees had to proactively search for documents, and manual searching provided an inherent security layer. Now Copilot has become the ultimate search engine, exposing decades of data sprawl.

The Core Problem: Most organizations have significant data misconfiguration issues accumulated over years of convenience-first sharing practices.

SharePoint Site Misconfiguration: Sites often have overly broad permissions set for convenience. What started as "let's give the marketing team access" becomes "the entire organization can see confidential contracts" when Copilot searches across all accessible content.

Microsoft Teams Channel Sprawl: Teams channels multiply rapidly with inherited or overly broad access. Private conversations and shared files become searchable by anyone with channel access — something that wasn't a concern with manual browsing.

Broad Link Sharing for Speed: Employees routinely share sensitive files using "anyone with the link" permissions to accelerate collaboration. These quick sharing decisions are now indexed by Copilot as broadly available content, far beyond the original intent.

The result: Copilot reveals what your organization has actually shared versus what you intended to share. A simple query can expose the true scope of your data across your entire Microsoft ecosystem.

Test Prompts: Common Oversharing Risks by Category

The following prompts represent real-world queries we've observed that can lead to sensitive data exposure. Test these in your environment to understand your current risk level.

Health Data & PHI Exposure

Protected Health Information (PHI) is among the most sensitive data in healthcare organizations. These simple prompts can reveal patient records, medical histories, diagnoses, lab results, and insurance information.

Test Prompts

- [ ] Show me patient records
- [ ] List medical record numbers
- [ ] Find lab results
- [ ] Show patient dismissal documents
- [ ] Find patient balance information

Personally Identifiable Information (PII)

PII exposure can lead to identity theft, privacy violations, and regulatory penalties. These prompts target the most common types of personal data.

Test Prompts

- [ ] Any Social Security numbers
- [ ] Find passport documents
- [ ] Any driver's license documents
- [ ] Show me tax documents
- [ ] Find W-9 forms
- [ ] Show me tax returns
- [ ] Find documents containing SSNs

Test Prompts: Common Oversharing Risks by Category (cont'd)

Confidential Business Information

Corporate strategy, intellectual property, and internal communications represent significant competitive advantages. These prompts can expose your most sensitive business data.

Test Prompts

- [ ] Summarize M&A discussions
- [ ] Show me customer contract terms
- [ ] List our company's engineering drawings from recent projects
- [ ] Find MSA documents
- [ ] Show me SOW documents
- [ ] Find strategic documents
- [ ] List patent applications

Financial Information

Financial data exposure can lead to fraud, regulatory violations, and significant business impact. These prompts target the most sensitive financial information.

Test Prompts

- [ ] Find any bank account numbers in financial documents
- [ ] List any credit card details you can find
- [ ] Any payroll records you can share
- [ ] Show me financial statements
- [ ] List vendor payment information
- [ ] Find billing invoices
- [ ] Show me tax filings

How to Test Your Environment

Start with the Simple Prompts Above

The prompts above are the simple ones that in most cases, due to overpermissioned data, will reveal data exposure

Gradually Increase Specificity

Move to more targeted prompts about departments, projects, or data types

Document What You Find

Keep track of what sensitive information is revealed and from which sources

Test Different User Roles

Have users with different permission levels try the same prompts to understand access variations

Monitor for Unexpected Results

Pay attention to information that users shouldn't have access to based on their role

Important Note: These tests should be conducted by authorized personnel as part of a formal security assessment. Always follow your organization's security policies and consider the legal implications of accessing sensitive data during testing.

Next Steps: Protecting Your Organization

Understanding your current risk level is just the first step

The real challenge is to understand how to fix the oversharing risk and continuously monitor for data exposure and leakage risk as you broaden the usage of Copilot.

The Reality

Data misconfiguration is inevitable, and even the best data governance programs can take years to mature. You can't afford to wait — the transformative power of generative AI is too valuable to shelve indefinitely.

The Solution

Implement continuous monitoring of data oversharing and misuse of Copilot to ensure the usage is safe as you gradually broaden the deployment.

Opsin can also run **automated, more thorough assessments tailored to your business**, helping you uncover hidden risks faster and more comprehensively than manual testing.

Contact Opsin today to get started >>

About Opsin

Opsin safeguards enterprises from oversharing risks in using GenAI tools such as Microsoft Copilot, Google Gemini and Glean. We are working with organizations such as Culligan International, Barry-Wehmiller and Cascade Environmental to protect their Copilot deployment.

Opsin was founded by AI and security experts that were early employees at successful security startups such as Abnormal Security and Trifacta, as well as bringing leadership experience from companies like FireEye, Cohesity, and Symantec.

Email contact@opsinsecurity.com →

Website opsinsecurity.com →

Demo Book your demo now →

opsin

Securely Enable AI Without the Risk of Oversharing