Pro CISO ®

AI Security Microsoft Copilot

Enterprise Risk Assessment

Understanding threats, data flows, and mitigations for organisations deploying Microsoft Copilot at scale

Confidential — For Client Use

March 2026

About Pro CISO®

Your Augmented Cybersecurity Team

Pro CISO®

Pro CISO® delivers expert cybersecurity leadership through its *Pro CISO-as-a-Service* model, providing organizations with a front-facing senior security expert and an entire team of specialists and solutions, to rapidly build and sustain cybersecurity resilience.

ISO 27001:2022 & ISO 9001:2015 Certified

prociso.com

ADVISORY

Strategic cybersecurity guidance aligned to ISO 27001, NIST 2.0, NIS2, GDPR, DORA.

OPERATIONS

Fully managed security services: Detection & Response, Threat Intel, Vulnerability Mgmt.

TESTING

High-quality security assessments & penetration testing — web, mobile, IoT, cloud.

CA/CR® METHOD

Proprietary Continuous Assessment / Continuous Remediation methodology for resilience.

OVERVIEW

Agenda

What this report covers

01 **What is Microsoft Copilot?**

Copilot Pro · M365 Copilot · Copilot Studio — capabilities & access model

02 **Data Exfiltration Risks**

Where data goes, Microsoft Graph exposure, retention & access purpose

03 **Prompt Injection Threats**

External & internal injection via email, SharePoint, Teams & plugins

04 **Plugin & Code Pull-in Risks**

Power Platform, Copilot Studio connectors & supply chain threats

05 **Built-in Safeguards**

Microsoft Purview, tenant isolation & admin controls in place

06 **Mitigation Framework**

Pro CISO recommendations to manage residual Copilot risk

SECTION 1

What is Microsoft Copilot?

An AI assistant deeply integrated into the Microsoft 365 ecosystem — three distinct deployment tiers

Microsoft Copilot embeds AI across productivity, enterprise, and development workflows. Each tier has a different capability footprint, data access boundary, and risk profile.

COPILOT PRO

Personal Microsoft 365 AI assistant ($30/user/month). Integrates with Word, Excel, PowerPoint, Outlook, and OneNote. Accesses personal M365 data via Microsoft Graph. Limited enterprise admin controls.

Graph API · Office Apps · Bing · Personal M365

LOW

M365 COPILOT

Enterprise deployment across all M365 applications. Uses the Microsoft 365 Semantic Index to query SharePoint, OneDrive, Exchange, and Teams. Full admin controls, Purview DLP, audit logging, and tenant governance.

Semantic Index · SharePoint · Teams · Purview · BizChat

MEDIUM

COPILOT STUDIO

Low-code platform to build custom Copilot agents and chatbots. Connects to external APIs, databases, and Power Platform. Enables custom plugins and connectors. Can trigger Power Automate flows autonomously.

Custom Agents · Power Platform · External APIs · Plugins

HIGH

CAPABILITIES

What Copilot Can Access in Your Organisation

Understanding the data access surface before assessing risks

Microsoft Graph & Identity

- All emails, calendars, contacts in Outlook/Exchange
- OneDrive & SharePoint files the user has access to
- Teams messages, channels, and meeting transcripts
- User directory and Entra ID metadata

Productivity & Documents

- Draft, edit, and summarise Word documents
- Generate charts, formulas, and reports in Excel
- Create and reformat PowerPoint presentations
- Summarise and action OneNote notebooks

Web & External Data

- Search and summarise the web via Bing integration
- Access external data via approved Copilot plugins
- Connect to third-party APIs in Copilot Studio agents
- Process attachments and files shared in Teams/email

Automation & Power Platform

- Trigger Power Automate flows from natural language
- Build Power Apps via Copilot Studio agents
- Query databases via Power Platform connectors
- Orchestrate multi-step autonomous workflows

RISK CATEGORY

Data Exfiltration

02

Where does your data go when Copilot processes it? Microsoft Graph, semantic index, and third-party plugins.

Pro CISO ®

DATA FLOW

Data Flow Through Microsoft Copilot

All interactions route through Azure OpenAI — tenant-scoped but cloud-transmitted

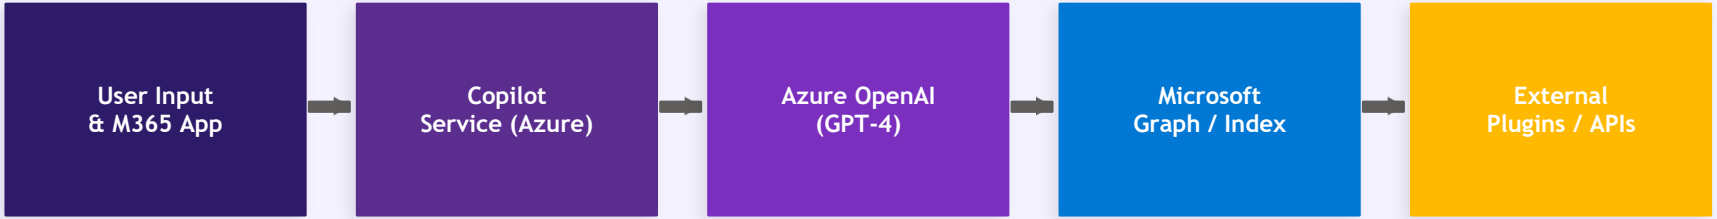

🔒 **Encrypted in transit (TLS 1.3) · Tenant-scoped · No cross-tenant data sharing**

Destination

Microsoft Azure (primarily within your M365 tenant). Copilot Pro routes via Microsoft 365 services. EU Data Boundary available for enterprise customers.

Retention Period

M365 Copilot interaction history is stored in Exchange Online. Retention controlled by your organisation's M365 retention policies. Defaults vary by plan.

Who Can Access

Microsoft engineers under strict access controls (JIT/JEA). Your M365 admins via audit logs. Law enforcement on valid legal request only.

Purpose of Use

Service delivery only. Microsoft does NOT use M365 Copilot/Copilot Pro commercial customer data to train foundation AI models. Copilot Studio may vary.

COPILOT PRO

Copilot Pro — Data Exfiltration Risks

Personal M365 data accessed via Microsoft Graph with limited enterprise governance controls

Personal M365 Data via Microsoft Graph

HIGH

Copilot Pro reads emails, OneDrive files, and calendar data via Graph API. Without enterprise DLP, no guardrails restrict what sensitive data is transmitted to Azure OpenAI.

Third-Party Plugin Data Sharing

HIGH

Copilot Pro supports Microsoft Store plugins (e.g., OpenTable, Kayak) that receive user context and query data. Plugin developer data retention is outside Microsoft's control.

Conversation History Persistence

MEDIUM

Copilot Pro chat history is stored in M365 and synced across devices. Shared or compromised devices expose the full conversation history including sensitive content.

No Enterprise DLP Controls

HIGH

Copilot Pro plans do not include Purview DLP integration. PII and financial data patterns in M365 content cannot be blocked from Copilot processing at the personal plan level.

Key Facts

- Graph API accesses all personal M365 data
- Plugin data goes to plugin developers
- No enterprise DLP at personal plan level
- Chat history stored in Exchange Online
- Commercial data protection: no AI training
- Bing search results sent to Azure OpenAI

Mitigation Priority

HIGH

Upgrade to M365 Copilot (Enterprise)

M365 COPILOT

M365 Copilot (Enterprise) — Data Exfiltration Risks

Overpermissioned access to the Semantic Index amplifies the blast radius across the organisation

Overpermissioned Data Surface

HIGH

M365 Copilot can access all content the authenticated user has permission to see — including files they have never opened. Years of legacy broad permissions now become Copilot's effective data scope.

SharePoint Over-Sharing Amplification

HIGH

Broadly shared SharePoint sites (e.g., 'Everyone' permissions) mean sensitive HR, legal, or financial documents are accessible to any Copilot user in the tenant via natural language queries.

Business Chat (BizChat) Data Aggregation

HIGH

BizChat can simultaneously query emails, Teams messages, SharePoint files, and calendar data. A single prompt can aggregate and return sensitive data from multiple M365 data sources.

Audit Logging Gaps by Default

MEDIUM

Not all Copilot interactions are captured in the Unified Audit Log without additional Purview configuration. Sensitive data retrieved via Copilot may not be forensically traceable.

Highest Exposure Surface

- ⚠ SharePoint (all user-accessible sites)
- ⚠ Exchange Online (full mailbox history)
- ⚠ Teams messages & call transcripts
- ⚠ OneDrive for Business files
- ⚠ Microsoft 365 Groups & Planner
- ⚠ Copilot interaction metadata

Mitigation Priority

HIGH

Requires SharePoint access review + Purview DLP

COPILOT STUDIO

Copilot Studio — Data Exfiltration Risks

Custom agents with external connectors can move sensitive data outside the M365 trust boundary

External Connector Data Exfiltration

HIGH

Copilot Studio agents connect to external APIs via custom connectors. Sensitive organisational data passed in plugin payloads leaves the M365 trust boundary and enters third-party infrastructure.

Overprivileged Service Principal Access

HIGH

Studio agents run as service principals, often with broad Microsoft Graph permissions. Misconfigured agents can access far more data than the use case requires — a major blast radius risk.

Power Automate Flow Data Transit

HIGH

Agents can trigger Power Automate flows that write data to external systems, send emails with M365 content, or post data to external webhooks — all outside standard DLP monitoring scope.

User Input Context Sharing

MEDIUM

User conversational input and retrieved M365 data are included in API calls to external connectors. Without strict connector governance, this data may be logged or retained by third-party services.

Trust Boundary Exposure

- External API connectors & webhooks
- Power Automate external actions
- Third-party plugin data pipelines
- Custom authentication flows
- Service principal Graph permissions
- User prompt content in API payloads

Mitigation Priority

HIGH — Requires connector allowlist + DLP policy

RISK CATEGORY

Prompt Injection

03

Malicious instructions in emails, SharePoint, Teams, and plugins — silently manipulating Copilot behaviour.

Pro CISO ®

INJECTION 101

Understanding Prompt Injection in Copilot

When M365 content manipulates Copilot's behaviour without user awareness

Prompt injection occurs when malicious instructions are embedded in content that Copilot reads — in emails, SharePoint documents, Teams messages, or plugin responses — causing Copilot to override its intended behaviour, leak data, or perform unintended actions.

Direct Injection

- User crafts prompts to bypass Copilot safety controls
- Jailbreak attempts via roleplay or hypothetical framing
- Requests to ignore system instructions or summarise all email
- Primarily an insider risk or deliberate misuse scenario

Indirect (Content) Injection

- Malicious instructions embedded in SharePoint or email content
- When Copilot summarises or processes the file, it follows them
- User is unaware — the attack is invisible to the victim
- Primary enterprise threat: exploits Copilot's broad M365 access

Indirect injection via M365 content is the primary enterprise concern — it is invisible to users and requires no malicious intent on their part.

EXTERNAL INJECTION

External Prompt Injection Vectors

Threats entering the M365 environment from outside the organisation via Copilot-processed content

Phishing Emails in Outlook

HIGH

External attackers send emails with embedded injection instructions. When Copilot drafts a reply or summarises the email thread, it reads and acts on the hidden payload — without the user's awareness.

Mode: M365 Copilot · Copilot Pro

External SharePoint Links & Files

HIGH

Documents shared from external M365 tenants or public URLs may contain injection instructions. When Copilot processes the shared file, it follows embedded commands — potentially exfiltrating data.

Mode: M365 Copilot

Bing Web Search Results

HIGH

Copilot fetches and summarises web content via Bing integration. Malicious web pages with hidden instructions (white text, HTML comments) can redirect Copilot's behaviour during web-assisted tasks.

Mode: Copilot Pro · M365 Copilot

Third-Party Plugin API Responses

HIGH

Copilot Studio plugins call external APIs and process their responses. A compromised or malicious API can return injection instructions embedded in JSON payloads that Copilot then executes.

Mode: Copilot Studio

External Calendar Invites

MEDIUM

Meeting invites from external parties may contain instruction text in the body. When Copilot processes calendar events to generate summaries or action items, it processes the injected payload.

Mode: M365 Copilot · Copilot Pro

Shared Files via Teams External Access

MEDIUM

Files shared in Teams meetings by external guest participants may contain injection content. Copilot processing post-meeting transcripts and files can encounter and act on the embedded payload.

Mode: M365 Copilot

INTERNAL INJECTION

Internal Prompt Injection Vectors

Threats originating from within the organisation's own M365 tenant content

SharePoint Documents & Wikis

HIGH

Internal documents, policy files, or SharePoint pages with embedded instructions. When Copilot summarises or answers questions from these sources, it may follow hidden payloads — modifying output or leaking adjacent content.

Mode: M365 Copilot

Teams Channel Messages

HIGH

A malicious internal actor or compromised account posts injection content in a Teams channel. When Copilot summarises channel threads or generates meeting catch-ups, it reads and executes the embedded payload.

Mode: M365 Copilot

Outlook Email Threads

HIGH

An internal email chain includes injected instructions. When Copilot drafts a reply or summarises the thread for another user, it processes the embedded instructions and may alter the content it generates.

Mode: M365 Copilot · Copilot Pro

OneNote & Shared Notebooks

MEDIUM

Shared OneNote notebooks may contain injected content (e.g. instructions in a hidden cell or page). Copilot processes these during document Q&A and may follow the embedded commands.

Mode: M365 Copilot

Planner & SharePoint List Items

MEDIUM

Task descriptions, metadata, or list item fields with injection content are processed when Copilot queries Planner or SharePoint lists to generate summaries or status reports.

Mode: M365 Copilot

Teams Meeting Transcripts

MEDIUM

An attacker verbally states injection instructions during a recorded Teams meeting. When Copilot generates meeting summaries or follow-up emails, it processes the transcribed payload.

Mode: M365 Copilot

SECTION 4

Plugin & Code Pull-in Risks

Copilot Studio connectors, Power Platform, and third-party plugins introduce supply chain risk

When Copilot invokes plugins, triggers Power Automate flows, or generates code and apps via Copilot Studio, it may introduce untrusted logic, overpermissioned connectors, or supply chain vulnerabilities.

Unvetted Microsoft Store Plugins

HIGH

Third-party Copilot plugins from the Microsoft Store may contain malicious functionality, data harvesting logic, or vulnerabilities. Plugin code is not subject to the same scrutiny as Microsoft's own services.

Power Automate Malicious Flows

HIGH

Copilot agents can trigger Power Automate flows. Malicious or poorly designed flows can write data to external systems, send emails with M365 content, or post data to external HTTP endpoints.

Copilot-Generated Power Apps

MEDIUM

Power Apps built via Copilot Studio assistance may include insecure configurations, overpermissioned connectors, or inadvertent data exposure — especially when deployed by non-technical users.

Overpermissioned Custom Connectors

HIGH

Copilot Studio custom connectors are frequently built with OAuth scopes that are broader than required. A connector with full Graph permissions becomes an unmonitored data exfiltration channel.

GitHub Copilot Code Vulnerabilities

MEDIUM

Code suggestions from GitHub Copilot (often used alongside M365 Copilot) may include vulnerable dependencies, hardcoded credentials, or insecure patterns deployed directly to production systems.

Third-Party Plugin Supply Chain

HIGH

Plugins from external vendors are subject to supply chain attacks. A compromised plugin update can silently exfiltrate data from every interaction across the organisation.

EXISTING CONTROLS

Built-in Safeguards

05

What Microsoft and M365 Copilot already do to protect users and organisations.

Pro CISO ®

MICROSOFT CONTROLS

Microsoft's Built-in Safeguards

Platform-level protections embedded in Copilot and Azure OpenAI by design

Commercial Data Protection

✓ CONTROL

Microsoft commits contractually that Copilot Pro, M365 Copilot, and Copilot Studio do NOT use customer data to train foundation AI models. This is governed by the Product Terms and the Customer Data Protection commitment.

Azure OpenAI Content Filtering

✓ CONTROL

All Copilot interactions route through Azure OpenAI with Microsoft's Responsible AI content filters. Harmful categories (CSAM, violence, weapons) are blocked at the model layer and cannot be bypassed by standard prompts.

Microsoft Purview Sensitivity Labels

~ PARTIAL

M365 Copilot respects Microsoft Purview sensitivity labels. Files classified as Confidential or Highly Confidential can be excluded from Copilot processing, preventing sensitive content from being surfaced via AI queries.

Tenant Isolation by Design

✓ CONTROL

M365 Copilot is strictly tenant-scoped. Data from one organisation cannot cross into another tenant. All Copilot interactions remain within the user's authenticated Entra ID (Azure AD) boundary.

Permission Inheritance (Zero Privilege Elevation)

✓ CONTROL

Copilot can only surface data the authenticated user already has permission to see. It cannot access files outside the user's current permission scope — no privilege escalation is possible through Copilot alone.

Unified Audit Log (Enterprise)

~ PARTIAL

M365 Copilot interactions are logged in the Microsoft 365 Unified Audit Log (with appropriate licensing). Security teams can investigate which data was retrieved, what prompts were used, and by whom.

ADMIN CONTROLS

M365 Admin & Compliance Safeguards

Administrative controls available to enterprise customers to govern Copilot access and data handling

Microsoft Purview DLP for Copilot

DLP policies can be configured to prevent Copilot from processing content matching sensitive information types (SSNs, PII, credit card data, health records) defined in Purview.

Entra ID Conditional Access

Access to Copilot can be governed by Conditional Access policies — requiring MFA, device compliance (Intune), network location restrictions, and risk-based access controls.

Copilot Usage Policies (Admin Centre)

M365 admins can enable or disable Copilot for specific users, groups, or apps. Plugin and connector access can be restricted to Microsoft-published apps only.

Power Platform DLP (Copilot Studio)

Data Loss Prevention policies in the Power Platform admin centre govern which connectors Copilot Studio agents can use — blocking connections to unapproved external services.

Communication Compliance (Purview)

Microsoft Purview Communication Compliance can monitor Copilot interactions for policy violations, insider threats, and regulatory compliance requirements.

EU

EU Data Boundary

European M365 customers can enable the EU Data Boundary to ensure Copilot processing, inference, and data storage remain within the European Union to meet GDPR data residency requirements.

Risk Summary Matrix

Consolidated risk rating by threat category and Copilot deployment tier

Risk Category	Copilot Pro	M365 Copilot	Studio
Conversation / Prompt Data Exposure	MEDIUM	HIGH	HIGH
Over-permissioned M365 Data Access	MEDIUM	HIGH	HIGH
File & IP Exfiltration via Graph	HIGH	HIGH	HIGH
External Prompt Injection	MEDIUM	HIGH	HIGH
Internal Prompt Injection (SharePoint/Teams)	LOW	HIGH	MEDIUM
Malicious / Unvetted Plugin	MEDIUM	MEDIUM	HIGH
Power Automate Flow Exploitation	LOW	MEDIUM	HIGH
Third-Party Connector Supply Chain	MEDIUM	HIGH	HIGH
Data Residency / Training Use Risk	MEDIUM	LOW	LOW

Risk Level: HIGH | MEDIUM | LOW

PRO CISO FRAMEWORK

Mitigation Framework

06

Practical controls for organisations to manage Microsoft Copilot risk effectively.

Pro CISO ®

TECHNICAL CONTROLS

Technical Mitigation Controls

Recommended technical measures to reduce Microsoft Copilot risk

Microsoft Purview DLP & Sensitivity

- ✓ Enable DLP policies to block Copilot processing of sensitive content types
- ✓ Deploy sensitivity labels across M365 — configure labels to restrict Copilot access
- ✓ Implement auto-labelling rules to classify SharePoint content before Copilot deployment

Applies to: M365 Copilot · Copilot Pro

SharePoint Access Governance

- ✓ Run SharePoint Access Reviews — remove 'Everyone' and 'All Staff' permissions
- ✓ Enable SharePoint Advanced Management for oversharing reports and site access controls
- ✓ Restrict external sharing to approved domains; block anonymous link access

Applies to: M365 Copilot

Copilot Studio Connector Controls

- ✓ Maintain an approved connector registry in Power Platform DLP policies
- ✓ Restrict custom connectors to internal services only; block external API endpoints
- ✓ Apply least-privilege OAuth scopes to all service principals used by Studio agents

Applies to: Copilot Studio

Monitoring & Audit

- ✓ Enable Copilot audit logging in the Unified Audit Log (E5/Purview add-on)
- ✓ Configure SIEM integration to alert on anomalous Copilot activity patterns
- ✓ Deploy Communication Compliance policies to monitor high-risk Copilot interactions

Applies to: All Modes

GOVERNANCE CONTROLS

Policy & Governance Controls

Organisational policies and processes to govern Microsoft Copilot use safely

01

AI Acceptable Use Policy

Define which Copilot tiers and capabilities are permitted for which roles. Prohibit entering classified, regulated (GDPR/HIPAA), or commercially sensitive content into Copilot without approval.

02

M365 Copilot Readiness Assessment

Before deploying M365 Copilot, conduct a readiness assessment: audit SharePoint permissions, identify over-shared content, classify sensitive documents, and confirm Purview DLP policies are active.

03

Copilot Studio Governance Process

All custom Copilot Studio agents must be reviewed and approved before deployment. Review service principal permissions, external connector scope, Power Automate flow logic, and data flows.

04

Security Awareness Training

Train staff on Copilot-specific risks — prompt injection via M365 content, how to recognise unusual Copilot behaviour, and the correct process for reporting suspected AI incidents.

05

Incident Response for AI Misuse

Extend the incident response plan to cover Copilot-specific scenarios: data surfaced incorrectly, prompt injection detected in audit logs, compromised plugin, or unauthorised data aggregation via BizChat.

06

Continuous Assessment (CA/CR® Method)

Apply Pro CISO's CA/CR® framework to M365 Copilot: regular threat model reviews as Microsoft adds capabilities, oversharing remediation cycles, and iterative policy alignment to evolving risks.

PURVIEW CONTROLS

Purview — Classify, Block & Report

Using Microsoft Purview to identify sensitive data types, prevent Copilot processing, and surface anomalous activity

01 · IDENTIFY DATA TYPES

Sensitivity Labels & Auto-Labelling

Deploy MIP sensitivity labels across M365. Configure auto-labelling policies to classify SharePoint, OneDrive, and Exchange content before Copilot is enabled — ensuring the Semantic Index reflects your classification.

Content Explorer & SITs / EDM

Map where PII, financial, and health data resides using Content Explorer. Define custom Sensitive Information Types (SITs) and Exact Data Match (EDM) to detect proprietary identifiers unique to your organisation.

Activity Explorer & Data Map

Track label events and Copilot file-access patterns in Activity Explorer. Run Purview Data Map scans to identify high-risk content stores that require DLP or label controls before enabling Copilot.

02 · BLOCK & RESTRICT

Copilot-Scoped DLP Policies

Create DLP policies in Purview Compliance Portal targeting the M365 Copilot workload. Block Copilot from processing or summarising content labelled Confidential or Highly Confidential across all M365 applications.

Label Encryption & Access Control

Apply Rights Management encryption via sensitivity labels to prevent Copilot indexing or summarising protected documents — effective even when SharePoint permissions are over-permissive or misconfigured.

Connector DLP & Endpoint Controls

Block Copilot Studio external connectors via Power Platform DLP policies. Use Endpoint DLP to prevent users from pasting Copilot output into unapproved apps or uploading it to personal cloud storage.

03 · DETECT & REPORT ANOMALIES

Unified Audit Log — Copilot Events

Enable Copilot audit logging in Purview Compliance Portal. Query CopilotInteraction events to see exactly which files each user triggered per session — essential for breach forensics and regulatory reporting.

Insider Risk & Comm. Compliance

Configure Insider Risk Management to flag anomalous Copilot usage: bulk file access, HR/finance queries outside role scope, or activity post-resignation. Set Communication Compliance to review high-risk Copilot prompts.

Alert Policies & SIEM Integration

Build custom alert policies for high-volume Copilot activity and DLP matches on Copilot queries. Stream audit events and alerts to Microsoft Sentinel or your SIEM for real-time correlation and incident response.

TENANT ISOLATION

Tenant Isolation & Access Scoping Strategy

Reducing Copilot's blast radius through M365 architecture and permission controls

Properly scoping what Copilot can access — through SharePoint permissions, sensitivity labels, and tenant configuration — is the single highest-value control for reducing M365 Copilot risk before and after deployment.

✓ RISK REDUCTIONS VIA TENANT SCOPING

SharePoint Permission Remediation
Remove Everyone/All Users permissions. Copilot only surfaces content the user explicitly can access.

Sensitivity Label Enforcement
Apply sensitivity labels broadly — exclude Confidential/Highly Confidential from Copilot scope.

Copilot Studio Connector Allowlist
Block external API connections by default. Restrict connectors to an approved allowlist via DLP.

Scoped Copilot Rollout by Role
Enable Copilot incrementally for low-sensitivity roles first, using Entra ID group-based licensing.

⚠ RESIDUAL RISKS THAT REMAIN

Azure OpenAI Still Processes Data
All Copilot interactions route through Azure OpenAI regardless of tenant scoping — it limits what's sent.

Semantic Index Crawls Broadly
Semantic Index crawls continuously — new content added after remediation reintroduces exposure risk.

Plugin & External Connector Gaps
Studio connector scope must be managed separately in Power Platform — not covered by tenant isolation.

Copilot Capabilities Evolve Rapidly
New Copilot features (agents, multi-agent) may bypass current controls — requires ongoing threat review.

Verdict: Tenant scoping + SharePoint permission remediation are essential pre-deployment steps. Combine with Purview DLP + sensitivity labels + scoped rollout.

ROADMAP

Mitigation Roadmap

A phased approach to securing Microsoft Copilot — delivered by Pro CISO

Phase 1
Immediate (Week 1-2)

- Run SharePoint Access Review — remove broad permissions
- Deploy AI Acceptable Use Policy for Copilot
- Enable Copilot audit logging in Unified Audit Log
- Identify and label sensitive SharePoint content

Phase 2
Short-term (Month 1-2)

- Deploy Purview DLP rules for Copilot-processed content
- Configure sensitivity labels to restrict Copilot access
- Implement Power Platform DLP connector allowlist
- Conduct security awareness training on Copilot risks

Phase 3
Medium-term (Month 3-6)

- Deploy Communication Compliance for Copilot monitoring
- Establish AI incident response playbook for Copilot
- Conduct threat model review of all Copilot Studio agents
- Pilot Pro CISO CA/CR® continuous assessment for AI risk

Pro CISO delivers all three phases via our CA/CR® Continuous Assessment / Continuous Remediation programme.

CONCLUSION

Key Takeaways

What every organisation deploying Microsoft Copilot needs to understand

1

Copilot inherits your permission debt

Copilot sees everything the user can see. Years of overpermissioned SharePoint, broad distribution groups, and legacy file shares all become Copilot's data scope. Fix permissions before deploying.

2

M365 Copilot & Studio carry the highest risk

Enterprise Copilot's broad Semantic Index access and Studio's external connector capability create the largest attack surfaces. Personal Copilot Pro is narrower but lacks enterprise controls.

3

Prompt injection via M365 content is real

Emails, SharePoint documents, Teams messages, and meeting transcripts are all injection surfaces. A single malicious document in a shared library is a potential attack vector against any Copilot user.

4

Microsoft's safeguards are necessary but not sufficient

Tenant isolation, sensitivity labels, and Purview DLP significantly reduce risk but require active configuration. Default settings are not hardened — enterprise governance must be explicitly enabled.

Pro CISO ®

Ready to Secure Your Corporate AI Deployment?

Pro CISO® provides comprehensive AI security risk assessments, policy development, and ongoing CA/CR® monitoring to help your organisation harness the power of AI safely and confidently.

Website: prociso.com

Phone: +31 20 211 7467

Email: AaaSK@prociso.com

ISO 27001:2022 & ISO 9001:2015 Certified