


Security risks for customers using other AI tools to egress Microsoft 365 data

A focused whitepaper on Microsoft SharePoint file access via other AI connectors and the resulting security gaps, with recommendations for a safer future state.



Executive summary

This whitepaper explains how other AI applications use connectors to Microsoft 365 APIs to access Microsoft SharePoint content on behalf of a signed-in user, and why those patterns can increase the risk of data egress and policy bypass compared to Microsoft 365 Copilot (Copilot). It summarizes observed behavior differences across common file conditions (permissions, sensitivity labeling, and encryption), highlights security gaps customers should understand before enabling connectors from other AI tools and recommends a future state in which access decisions consistently honor Microsoft Purview controls, sensitivity label encryption semantics, and timely permission changes.

Introduction

Customers are increasingly connecting other AI tools to Microsoft 365 to summarize, search, and generate insights from enterprise content. A top use case is granting these tools access to Microsoft SharePoint files via connectors. While connectors can improve productivity for end users, they can also create new pathways for Microsoft 365 data to leave the tenant boundary and be processed in external systems posing a business risk for Microsoft and its customers. This document describes connector access patterns, compares how Copilot and several competitor offerings behave when retrieving SharePoint files via Microsoft APIs, and outlines the security implications customers should evaluate before enabling other AI connectors.

Scope



In scope

- Microsoft SharePoint Online (SharePoint) files that are accessed by other AI tools through connectors, including both *synced* patterns (copying content out of the tenant) and *federated* patterns (real-time queries/calls).
- Authentication and authorization basics for “on behalf of user” access via Microsoft APIs.
- Observed differences in responses when files are unencrypted vs. encrypted with sensitivity labels.
- Security risks related to stale permissions, data egress, and inconsistent enforcement of Purview controls and label permissions.

Out of scope

- Other Microsoft 365 content types (Exchange email, Teams messages, etc.).
- Non-security concerns (e.g., cost, UX quality, general model performance).

How connectors work

In this context, a “connector” is an integration that allows other AI applications to call Microsoft APIs (commonly Microsoft Graph and SharePoint/OneDrive endpoints) to retrieve content from a customer tenant. Typically, the user signs in with Microsoft Entra ID, consents to an application, and the third-party service receives tokens that enable it to query files the user can access. The connector then either retrieves data in real time (federated) or copies data into the vendor’s index/store (synced) to support faster search and AI workflows.

Common connector patterns

Synced (copied) access

The connector extracts files (or file-derived text/metadata) and stores a copy outside the Microsoft 365 tenant boundary for indexing. This can introduce a window where SharePoint permission changes, sensitivity label changes, or revocations are not reflected immediately in the third-party store. The duration and scope of this window depends on factors such as connector sync frequency and whether it’s a full content or partial sync.

Federated (real-time) access

The connector queries Microsoft APIs on demand to fetch content. Authorization is evaluated by Microsoft at request time, but data can still egress (content and/or metadata) to the third-party service as part of the response and downstream processing.

Authentication and “on behalf of user” access

Most connectors operate “on behalf of” a user by using delegated permissions: The other AI tool receives an access token representing the signed-in user and the app and then calls Microsoft APIs to retrieve resources that the user is authorized to access. In practice, customers should treat this as extending the user’s access path to include an external processor: Even when Microsoft correctly enforces access at the API boundary, any content returned can be stored, transformed, summarized, or redistributed by the other AI tool according to its architecture and governance.

Comparative behavior: Copilot vs. other AI tools when accessing SharePoint files

The table below summarizes observed behaviors of Copilot vs. other AI tools when they access SharePoint files on behalf of a user. We asked the same prompt to each of the AI tools to try to access the same file under different file conditions specified in the first column: "Summarize [file name]" with link to specific file.

This is intended to illustrate policy and control mismatches that can lead to unintended exposure. Testing results as of May 18, 2026; products and behaviors may change over time.

Condition on a SharePoint file	Microsoft 365 Copilot	ChatGPT Enterprise*	Gemini Enterprise*	Claude Enterprise
User does not have access to the file	Has no access to the file	Has no access to the file*	Has no access to the file*	Has no access to the file
User has access to the file; no sensitivity label; does not contain sensitive information types (SITs)	Has access to the file	Has access to the file*	Has access to the file*	Has access to the file
Contains SITs; sensitivity label applied; Microsoft Purview DLP policy configured to block Copilot processing	File is blocked from Copilot processing	Has access to the file Does not honor Purview DLP policy, no DLP blocking observed	Has access to the file Does not honor Purview DLP policy, no DLP blocking observed	Has access to the file Does not honor Purview DLP policy, no DLP blocking observed

Products tested:

- Microsoft 365 Copilot, model used: "Auto"
 - ChatGPT Enterprise + M365 connectors; model used: GPT 5.5 Thinking
- Gemini Enterprise + M365 connectors; model used: Gemini Auto
 - Claude Enterprise + M365 connectors; model used: Opus 4.7

* When using ChatGPT Enterprise or Gemini Enterprise sync connectors, access reflects SharePoint permissions as of the last sync

Condition on a SharePoint file	Microsoft 365 Copilot	ChatGPT Enterprise*	Gemini Enterprise*	Claude Enterprise
Contains SITs; no sensitivity label applied	Has access to the file SITs in files are not blocked solely due to SIT presence; Copilot response and/or web search may be blocked by Purview DLP if text prompt contained specified SITs in the policy	Has access to the file	Has access to the file	Has access to the file
Encrypted sensitivity label applied; user has access	Has access to file	Does not decrypt the file for permissioned user	Does not decrypt the file for permissioned user	Does not decrypt the file for permissioned user
Encrypted sensitivity label applied; DLP policy configured to block Copilot processing	File is blocked from Copilot processing	File is not processed because it does not decrypt the file for permissioned users	File is not processed because it does not decrypt the file for permissioned users	Decrypts "content preview" Does not honor Purview DLP policy, no DLP blocking observed
Encrypted sensitivity label applied; extract restricted when file is open**	File is blocked from Copilot processing	Extracts and summarizes file content despite extract restrictions	Not applicable: no access to file via add-in	Extracts and summarizes file content despite extract restrictions

Products tested:

- Microsoft 365 Copilot, model used: "Auto"
 - ChatGPT Enterprise + M365 connectors; model used: GPT 5.5 Thinking
- Gemini Enterprise + M365 connectors; model used: Gemini Auto
 - Claude Enterprise + M365 connectors; model used: Opus 4.7

* When using ChatGPT Enterprise or Gemini Enterprise sync connectors, access reflects SharePoint permissions as of the last sync

** Using the AI tools with the file open, Add-in file access tested in Excel using Edit with Copilot in Excel, Claude and ChatGPT for Excel add-ins

Condition on a SharePoint file	Microsoft 365 Copilot	ChatGPT Enterprise*	Gemini Enterprise*	Claude Enterprise
Restricted Content Discovery is enabled on the site the user the user has access and recent interaction with the content	File is discoverable	File is discoverable; (required multiple prompts and explicit file and site references to surface)	File is discoverable	File is discoverable
Restricted Content Discovery is enabled on the site, and the user the user has access but no recent interaction with the content	Copilot works when the file is open in an Office app, but it is not discoverable in Copilot/search	File is undiscoverable	File is undiscoverable	File is undiscoverable
User has access to the file; Restricted Access Control (RAC) is enabled on the site; user IS in allowed groups	Has access to the file	Has access to the file	Has access to the file	Has access to the file
User has access to the file; Restricted Access Control (RAC) is enabled on the site; user is NOT in allowed groups	Tells the user the file is not accessible (i.e. cannot summarize)	No access to the file	No access to the file	Tells the user the file is found and provides a high-level summary based on search metadata

Products tested:

- Microsoft 365 Copilot, model used: "Auto"
 - ChatGPT Enterprise + M365 connectors; model used: GPT 5.5 Thinking
- Gemini Enterprise + M365 connectors; model used: Gemini Auto
 - Claude Enterprise + M365 connectors; model used: Opus 4.7

* When using ChatGPT Enterprise or Gemini Enterprise sync connectors, access reflects SharePoint permissions as of the last sync

Security gaps when using other AI tools to egress Microsoft 365 data

01

Data egress beyond tenant boundary

Other AI tools inherently introduce an external processing boundary. Synced connectors may copy content to vendor-controlled stores, expanding the data footprint and increasing breach/exfiltration impact. Across both synced and federated tools, returned or generated content may be processed or stored outside the Microsoft tenant boundary, where DPA controls may not apply.

02

Stale authorization in synced connectors

When access is cached outside Microsoft 365, SharePoint permission changes may not be honored immediately (or at all) until the vendor re-syncs and re-evaluates access, creating a “stale permissions window.”

03

Inconsistent enforcement of Purview controls

Observations suggest other AI tools may not consistently honor policies that customers expect to apply to Microsoft 365 content (e.g., DLP-style controls in the interaction flow, or other governance layers), resulting in potential exposure of sensitive information.

04

Metadata leakage in federated access

Even when full content is not retrieved, queries and responses can leak file names, paths, titles, or snippets to third parties. This can disclose sensitive project names or business context.

05

Reduced ability to apply tenant-side safeguards consistently

Once content is returned to a third-party service, tenant administrators have fewer options to enforce retention, auditing, secondary use restrictions, or downstream sharing constraints in a uniform way.

Security & governance advantages of using Copilot with data in Microsoft 365

01 Enhanced security and compliance

Copilot can process data within the Microsoft 365 tenant boundary, helping to ensure sensitive information remains protected by tenant controls, encryption, and compliance policies. This minimizes the risk of data egress and unauthorized exposure, supporting regulatory and organizational requirements.

02 Consistent enforcement of governance policies

All Purview controls, sensitivity labels, and data loss prevention (DLP) policies are built-in integrations when data stays within Microsoft 365. Administrators retain visibility and control over how information is accessed, processed, and shared.

03 Immediate authorization updates

Access permissions and authorization changes are enforced, helping to eliminate the risk of stale permissions windows that can occur with external processing or synced connectors. This helps prevent unintended data leakage and helps to ensure only authorized users can access sensitive content.

04 Helps to reduce risk of metadata leakage

By keeping data and query responses within the Microsoft 365 environment, you can set up the appropriate Purview DLP policies to prevent data loss of file names, paths, or business context to third parties. This helps to maintain confidentiality for projects and sensitive business initiatives.

05 Streamlined management and auditing

Tenant administrators can uniformly apply retention, auditing, and downstream sharing policies. This helps to enable oversight and reduce complexity in managing information across the organization.

Security benefits of implementing Microsoft Copilot connectors

Other AI tools use connectors to route Microsoft 365 data through external systems with varying levels of governance, creating the security gaps described earlier in this document.

Microsoft's Copilot connectors, by contrast, provide a more governable path to extend Copilot to external data sources while preserving Microsoft 365 security fundamentals.

Access is brokered through Microsoft Entra ID with scoped, least-privilege permissions and centralized admin consent, enabling consistent application of Conditional Access and rapid revocation when risk changes. Because connector activity is surfaced through Microsoft's admin and audit experiences, organizations can monitor who accessed which data, apply policy, and demonstrate compliance without relying on vendor-specific logging.

When combined with Microsoft Purview controls (e.g., sensitivity labels and DLP where applicable), Copilot connectors help reduce uncontrolled data egress and create a clearer, tenant-managed governance boundary for AI-driven retrieval.

Recommendations (1/2)

Desired future state principles

01

Keep data in the tenant whenever possible

Prefer architectures that minimize copying and persistently storing Microsoft 365 content outside the tenant boundary.

02

Consistent policy enforcement

Access to SharePoint files via APIs should honor Microsoft 365 authorization, sensitivity label permissions (including encryption rights such as EXTRACT), and customer governance expectations end to end.

03

Minimize egress by default

Where full file content is not required, return the least-privilege/least-data result (e.g., metadata-only) and require explicit user action for broader extraction.

04

Reduce stale-permissions windows

For any cached/synced index, require near-real-time permission validation and rapid revocation behavior.

05

Transparent auditing

Customers should be able to understand what content left the tenant, when, under which user context, and for what purpose.

Recommendations (2/2)

What should happen with other AI tools (*with connectors enabled*)

01

Prefer federated retrieval over bulk sync

If content must be queried, fetch it on demand and avoid building long-lived external copies except where explicitly required and approved.

02

Honor encryption rights without exception

Do not decrypt or show “content previews” for encrypted sensitivity labels when EXTRACT is disabled. Treat this as a hard boundary, not a UI choice.

03

Implement continuous access evaluation

Validate permissions at query time and re-check access for cached items; rapidly purge or re-encrypt cached content when access is revoked or labels change.

04

Offer customer-controlled guardrails

Provide controls to restrict which sites/libraries can be connected, limit which file types are retrieved, and configure whether content, snippets, or metadata can be stored externally.

05

Provide comprehensive auditability

Log which files were accessed, what was stored externally, which users initiated access, and support customer export of these logs for compliance.

06

Consider governance policies for add-ins and plugins from other AI tools

Consider implementing policies to govern add-ins and plugins from other AI tools that access Microsoft 365 data via user-delegated permissions, as these may introduce additional programmatic access paths not consistently governed by existing controls.