

# Cleared for Takeoff? Copilot Legal and Technical Preflight Checklist

## WRITTEN BY

Jason Lichter

---

*This article was originally published on June 25, 2025 on [Legaltech News \(LTN\)](#) and is republished here with permission.*

Millions of companies use the Microsoft 365 suite of tools every day to create, communicate, and collaborate, but far fewer have adequately grappled with the legal risks introduced by Copilot, the powerful generative AI assistant embedded in those same applications. While Copilot can enhance employee productivity, creativity, and connectivity, it may do so at the expense of privacy, security, and compliance without adequate planning and oversight. In this article, I will first identify the potential pitfalls of a laissez-faire approach to deploying Copilot and then provide actionable recommendations on how to navigate those hazards.

One key distinction to make at the outset is between the consumer-facing version of Microsoft Copilot and the commercial-facing version known as Microsoft 365 Copilot. To add to the complexity, the consumer and commercial flavors of Copilot each have both free and paid tiers. Because Microsoft's [enterprise data protection](#) controls and commitments apply only to the commercial iteration of Copilot, companies should take steps to ensure that authorized users are directed exclusively to that version of the standalone Copilot chatbot in the ordinary course, and the balance of this article will focus on the commercial Copilot offering.

When composing its responses, Copilot leverages the incredible reach of the Microsoft Graph to reference the litany of emails, documents, Teams chats, meeting transcripts, and other data repositories that each user has permission to access within (and sometimes beyond) the Microsoft 365 ecosystem. But since many organizations rely more on "security by obscurity" than on a true "zero trust" or "least-privilege" approach to access controls, Copilot could surface sensitive content that a user technically always had permission to view if they knew where to look but would never otherwise have had reason to stumble upon. Microsoft has provided repeated assurances that the prompts, responses, and data accessed through Microsoft Graph are not used to train the foundation LLMs used by Copilot, but the risks of oversharing within the Copilot user base remain very real.

The tendency of many organizations to hoard information also means that Copilot is likely to encounter vast quantities of redundant, obsolete, and trivial (ROT) data as it traverses the Microsoft Graph, with no reliable way to set the ROT aside in favor of sources that are more reliable, accurate and useful. Just as a junior attorney might select an outmoded sample from a disorganized document management system when tasked with preparing their first memorandum, the quality of Copilot's output can suffer as it struggles to separate the wheat from the chaff. In fact, Copilot could theoretically contribute to the problem by generating the very content that it then encounters when responding to future user prompts. This side effect of poor data hygiene is a variant on the phenomenon more broadly known as model collapse, wherein generative AI models trained primarily on their predecessors' output produce increasingly inferior results.

Given its prodigious attack surface, Copilot has become an attractive target for threat actors. On June 11, 2025, cybersecurity researchers disclosed a zero-day vulnerability—since dubbed EchoLeak and already patched by Microsoft—that could bypass security controls and trick Copilot into exfiltrating sensitive corporate data simply by sending a benign-looking email with embedded prompt instructions that are executed during subsequent, seemingly-unrelated Copilot conversations. Frighteningly, this particular exploit did not depend on the recipient of the nefarious email clicking a link or taking any other action we all are assiduously trained to avoid.

The good news is that, with some effort and thoughtful planning, organizations can mitigate many of these risks without remaining on the Copilot sidelines. Microsoft Purview provides companies with the tools necessary to implement sound data classification, organization, and retention practices, including a labeling taxonomy that restricts Copilot's access to sensitive data, but retroactively applying governance to terabytes of preexisting SharePoint site content may be daunting. Enabling [Restricted SharePoint Search](#) can be an effective stopgap measure by restricting Copilot's scope of collaborative location coverage to an allowed list of curated SharePoint sites. Relatedly, while corporations have understandably sought to restrict the sprawl of repositories that SharePoint and Teams can unleash, imposing too many barriers on new site creation can lead users to commingle disparate information in a handful of preapproved locations, which can be equally problematic. Striking the right balance between structure and flexibility is critical.

As soon as Copilot is turned on, early adopters are bound to use it extensively, while others may proceed more warily. But companies need not take a “one-size-fits-all” approach to their Copilot deployment strategy. A phased rollout may be more prudent, perhaps starting with a small number of tech-savvy pilot participants, then expanding to those business functions likely to see the greatest productivity benefit at the lowest risk (e.g., sales and marketing), and eventually enabling teams who regularly handle the most sensitive data (e.g., human resources and legal) only after the governance model has been validated.

Once the floodgates are open, Copilot is virtually guaranteed to generate a steady stream of artifacts, much of which is likely to be discoverable (although precedent is still scant). The best time to enable reasonable data lifecycle management policies is before a user base grows accustomed to infinite retention. Absent a legal duty to preserve, the optimal duration to keep prompts, responses, document versions, meeting recaps, and other Copilot interactions will vary by organization and sometimes even by department, but adopting Microsoft's generous out-of-the-box retention periods should be an explicit decision, not a dereliction.

Recognizing the challenges confronting organizations, Microsoft recently announced [Copilot Control System](#), billed as a robust, coherent system of integrated enterprise-grade controls for Copilot and its agents covering

security, governance, management, measurement, and reporting. But even with Microsoft's assistance, companies should consider engaging experts to help assess their Copilot technical readiness; comply with international data protection regimes; avoid a damaging implementation misconfiguration; update applicable policies and procedures (including as to record retention, responsible AI use, and incident response), prepare and deliver real-world training; and assemble and advise cross-functional AI Centers of Excellence.

## RELATED INDUSTRIES + PRACTICES

- [eDiscovery + Data Management](#)