JUNE 22, 2026

Microsoft 365 Copilot is the floor, not the ceiling

Owning the licence is the easy part. Copilot's answers are only as good as what it can retrieve from your data — and in most firms, that data is a mess.

By **Dave Taylor**

Just owning and using Microsoft 365 Copilot is not enough on its own. The software is the floor — the cheap, easy part. The magic of Copilot is that it can answer your questions using your company's own data, not just the public internet. But its responses are only ever as good as what it can pull back from what it has access to, and in most firms, that data is messy. Microsoft's own 2023 Work Trend Index found early Copilot users completed tasks 29% faster, and 77% didn't want to give it up. That's how good Copilot can be. Whether your firm gets anywhere near those numbers depends almost entirely on the readiness work underneath the licence — not the model you've paid for.

What is Copilot actually doing when you ask it something?

Out of the box, Copilot is a generic assistant. The intelligence you're paying for sits in the model, but the *relevance* — the part that makes an answer useful to your firm rather than to anyone's firm — comes from retrieval. When you ask a question, Copilot grounds its answer by searching across your tenant: SharePoint, OneDrive, Teams, email, whatever it has permission to see. Microsoft has a name for that connected pile of content — the **Microsoft Graph**: the live index of every file, email, chat, meeting, and calendar in your tenant, plus the relationships between them. The Graph is what Copilot reads. It pulls back the items it judges relevant, then writes an answer from them. This is the same retrieval pattern we've covered before in our primer on how RAG grounds an AI in your own documents.

So the quality of every Copilot answer is set by the quality of what retrieval finds. Garbage Graph, garbage Copilot. If your firm's knowledge is tidy, named, and current, Copilot looks brilliant and the early-user numbers hold. If it's scattered, duplicated, or wrongly shared, Copilot inherits every one of those problems and hands them back with the confidence of a tool that sounds like it knows exactly what it's talking about. Most rollouts stall here — not because the AI is weak, but because the data underneath it was never built for a machine to read on its own. That's the trap with a licence-led purchase: you've paid for the part that was always going to be good, and skipped the part that decides whether it works for you specifically. The model is the bit you can buy in an afternoon. The readiness underneath it is the bit that takes the actual work, and it's where the value of the rollout is either won or quietly lost.

The three ways grounding goes wrong

Grounding fails in three distinct ways, and each one kills a rollout differently.

1. Copilot finds nothing

A firm's real knowledge — how it actually prices a job, solves an unusual problem, manages a difficult client — rarely lives tidily in SharePoint. It lives in bespoke systems, in email threads, and in people's heads. So Copilot grounds on whatever scraps it can reach and returns a generic answer. Staff try it once in week one, decide it's useless, and you never get that first impression back. Adoption dies before it has a chance to start.

2. Copilot finds the wrong thing

Picture ten versions of the same procedures document scattered across a shared drive, three of them from 2019, none of them marked as the master. Copilot has no instinct for which one is current — it cites the stale one and presents it as though it were settled fact. A confidently wrong answer — an out-of-date price, a superseded policy, a rate that changed two years ago — is worse than no answer at all, because someone acts on it. A blank is a non-event; a wrong number that sounds authoritative travels downstream into client work, a quote, or an invoice before anyone thinks to question it. It's also the hardest of the three to catch early, because nothing visibly breaks — the tool looks like it's working, gives fluent answers, and earns trust, right up until a number it served confidently turns out to have come from a document that should have been deleted years ago. Quiet damage is still damage.

3. Copilot finds what it shouldn't

This is the one that turns a rollout into a brand-level incident. Copilot respects permissions correctly — that part works. The problem is that in most firms, permissions are a mess: files shared to "everyone" years ago, broken inheritance, anyone-with-the-link access nobody remembers granting. Those files were always technically accessible; the only thing protecting them was that nobody could find them. Copilot is a search engine that reads everything you can reach and answers in plain English. It makes the buried findable.

So someone types "what does everyone earn?" and gets an answer. For any firm holding sensitive records — financials, contracts, case files, personal data — that's not an IT ticket; it's a question about whether the firm can be trusted with data, which is the whole business. Microsoft knows this is the single biggest blocker, which is why it built SharePoint Advanced Management with oversharing and permission reports specifically to surface where rollouts stall. The exposure is real and measurable — Cyberhaven's 2026 telemetry found 39.7% of workplace AI interactions expose sensitive data. None of those over-shared files are new; they were sitting there all along, accessible to anyone with the patience to dig. What's changed is that Copilot reads everything you can reach and answers in plain English, so the digging is now a single typed question. Copilot doesn't create the oversharing problem — it just removes the obscurity that was quietly doing all the work of keeping it hidden.

You already have Copilot. The next step is tightening the floor

Most firms reading this aren't setting Copilot up from scratch — nobody is, anymore. You've had the licences for a while. Some people use it every day; some tried it once and drifted back to old habits. The question now isn't whether to adopt Copilot; it's how to get more out of what you already own, and how to build your own AI tools on top of it without standing them on a mess.

That second part matters more every month. The real value sits in an AI tool stack shaped to your firm: agents that draft your routine documents, assistants that answer in your house style, automations wired into the systems you actually run on. Every one of those tools grounds on the same floor Copilot does — your Graph and your permissions. Build them on a loose Graph and you scale the garbage-in problem across every tool at once.

So before you build up, you tighten the floor. You don't have to clean every site in the tenant first — and you shouldn't try, because that's a year of work. The move is to restrict what your AI is allowed to ground on. Microsoft built Restricted Content Discovery for exactly this: you set an allow-list of vetted "golden source" sites, and Copilot only ever sees the rooms you've already tidied. The partner-pay spreadsheet and the graveyard of 2019 drafts stay outside reach until you've dealt with them properly.

Why a tight Graph is the foundation, not housekeeping

The instinct in most firms is to treat the Copilot licence as the project and the data cleanup as housekeeping you'll get to eventually. That's backwards — and it gets more backwards the moment you start building custom tools. The model is a commodity; Microsoft, your competitors, and the firm down the road all rent the same one. The Graph underneath it — where your knowledge lives, how it's named, who can see it — is the part that's specific to your firm, the part nobody can buy on your behalf, and the part that decides whether everything you build on top earns its keep or quietly gets switched off.

There's a second payoff that's easy to miss. Mapping where knowledge lives, retiring duplicates, closing the oversharing, and naming a single source of truth per knowledge type — the same discipline we walk through in mapping your AI exposure — makes the firm better at finding things with or without AI. You were going to need it the first time a regulator, a client, or a new hire asked where the current version of something actually lived. Tightening the floor for your AI stack just hands you the deadline and the budget to finally do it, and every tool you build afterwards inherits the benefit. This is the readiness work we run with firms before they point custom tools at live data, and it's the cheapest insurance you'll buy against a confidently-wrong answer reaching a client.

Where this leaves you

A Microsoft 365 Copilot licence buys you a capable generic assistant and nothing more — and if you've had it a while, you already know that. What turns it into something your firm relies on, and what lets you build your own AI tools on top of it, is the state of the floor underneath: a tight Graph and correct permissions. The three ways grounding goes wrong — finding nothing, finding the wrong thing, finding what it shouldn't — all trace back to data that was never built for a machine to read. Fix the floor and you fix it once for every tool you'll ever run on it.

That floor is what Part 2 of this series gets into — the Microsoft Graph that Copilot grounds on, and how you extend it to reach the systems Copilot can't see on its own. If you're already running Copilot and want to optimize it — and start building an AI tool stack specific to your firm — book a 30-minute working session and we'll look at what your data is actually ready for.