

Before you turn on Microsoft 365 Copilot.

A board-ready governance guide for nonprofits, foundations, and rural hospitals—what Copilot can reach on day one, and how to adopt it safely.



Governed before it's turned on.

Centered Networks · Microsoft Solutions Partner

Turning on Microsoft 365 Copilot looks like flipping a switch. For a nonprofit, foundation, or rural hospital, it isn't—because this switch touches donor data, client records, patient information, and the trust your mission runs on.

The good news: you don't need to slow down your AI plans. You need to sequence them. This guide lays out what Copilot can reach on day one, the questions your board will ask, the controls that make it safe, and a 90-day path to adopt it with confidence. Our standard is simple: **governed before it's turned on.**

1 What Copilot can actually reach on day one

Microsoft 365 Copilot does not create new access. It answers each person through the access that person already has—their files, email, Teams messages, and the SharePoint sites they can already open. Microsoft documents this plainly: Copilot only surfaces organizational data a user already has at least view permission to.

That single fact is the whole risk picture. After years of SharePoint and Teams growth, most organizations' permissions were never designed for a tool that answers open-ended questions in seconds. The classic test prompt—*"How much does our executive director make?"*—will surface a salary file or a set of board minutes the moment they are over-shared, faster than any staff member ever could.

For a hospital, add HIPAA to that picture. For a foundation, add grantee and donor confidentiality. The exposure was already there; Copilot just makes it easy to reach.

Takeaway: the question is not "is Copilot safe?" It's "is our data governed enough for a tool this capable?" That is something you can measure before you turn anything on.

2 Five questions to bring your board

You don't need board approval to enable software. You do need to walk in with answers, not a request for trust. These five hold up in any boardroom:

1. **What can Copilot reach on day one?** It answers through the access staff already have. Turning it on is step one; what it can reach is the real exposure.
2. **Where is AI already in use here—approved or not?** "Nowhere" is an assumption, not an answer. Get visibility before you add a more capable tool.
3. **What evidence of control could we show this board?** Data boundaries, oversight, documentation a board can accept. If nothing comes to mind, that's the gap to close first.
4. **Who is accountable once it's on?** A name and a review cadence—not "IT will handle it."
5. **What do the 90 days after day one look like?** Adoption needs a roadmap with checkpoints, not just a switch.

An executive who can answer these walks into the boardroom with a plan. If you can't answer them yet, that isn't a failure—it's a short, fixed-scope exercise, described in Section 4.

3 The controls that make it safe

Microsoft built a control surface for exactly this. It ships with the platform; the work is deploying and operating it—which most rollouts skip. The pieces that matter for a mission-driven organization (these are Microsoft’s capabilities, documented on Microsoft Learn; CN’s role is to deploy and operate them):

- **Permissions and sensitivity labels**—Copilot respects your existing identity model and permissions, and honors Microsoft Purview sensitivity labels, so people only see what they were already cleared to see. Getting labels and permissions right is what prevents over-sharing.
- **Microsoft Purview**—the layer for classifying sensitive data, applying data loss prevention, and auditing how AI is used across your files.
- **The Copilot Control System**—Microsoft’s admin surface to secure, manage, and measure Copilot and agents: security and governance policy, agent lifecycle management, and adoption and ROI reporting through Copilot Analytics.
- **Data residency and privacy**—your data stays in your tenant, isn’t used to train foundation models, and is covered by Microsoft’s enterprise privacy commitments, including EU Data Boundary support where it applies.

None of this is automatic. Labels have to be designed, permissions cleaned up, policies set, and a review cadence stood up. That operating work is the difference between “Copilot is on” and “Copilot is governed.”

4 A governed path in 90 days

You don’t have to do this all at once, and you shouldn’t. A governed rollout moves in four steps, each a complete outcome on its own:

1. **See your exposure (about 2 weeks).** A shadow-AI audit and a Copilot readiness read that show, in plain terms, where AI is already in use and where your data is over-shared—ending in a briefing written for your board.
2. **Set the guardrails.** Stand up the governance baseline on Microsoft’s controls—sensitivity labels, Purview policies, permission cleanup, and Copilot Control System settings—so the exposure from Section 1 is closed before scale.
3. **Deploy where it pays off.** Put Copilot and a first production agent into one well-defined workflow—a weekly report, a recurring analysis, an intake process—with the exposure contained to what one team can already access.
4. **Operate and report.** Monitor adoption and drift, review governance on a set cadence, and bring the board a quarterly read on impact and risk.

Governance is not the thing that slows AI down. It’s the thing that lets you say yes—because once one department sees a working use case, three more ask for it.

5 A board-ready checklist

Before you turn on Microsoft 365 Copilot, you can show your board:

- ☐ A shadow-AI exposure report—where AI is used today, approved or not
- ☐ A Copilot readiness read—what Copilot could reach across your tenant
- ☐ Sensitivity labels and permissions cleaned up on the data that matters
- ☐ A named owner and a review cadence
- ☐ A 90-day roadmap with checkpoints, not just an on-switch

If you can hand your board this list, the AI conversation stops being about trust and starts being about a plan.

WHAT TO DO NEXT

Start with a two-week AI Quickstart.

Centered Networks helps mission-driven organizations adopt Microsoft 365 Copilot and agents the governed way—built on our Microsoft Solutions Partner depth, including Data & AI and Security. The AI Quickstart is a fixed-scope, roughly two-week exercise that ends with the shadow-AI exposure report, the Copilot readiness read, and a board-ready briefing described above.

[Book a readiness call →](#)

Governed before it's turned on. It's the standard your board would set anyway.