GUIDE

Guide to GenAI Risks and How to Mitigate Them

ChatGPT, Gemini and Copilot

AI Assistants Are Great, Until They Leak Your Data

By now, your teams are already using AI tools like Copilot, Gemini, and ChatGPT, whether you sanctioned it or not. They're writing emails and summarizing docs with it, of course. But they're also tapping into your most sensitive data, remixing it on demand, and sharing it in places you never expected (or wanted).

Here's the tradeoff: these tools take that data and amplify it. Sensitive information that used to sit quietly in a document or spreadsheet is now only a prompt away from being surfaced in a meeting, pasted into a slide deck, or summarized into a chat. In other words, be afraid, be very afraid.

If your data isn't labeled, locked down, and monitored, AI assistants can become high-speed leak machines.

What Makes it All Even Worse

- AI doesn't recognize what's sensitive unless you explicitly specify it
- Default permissions = risk amplification
- AI hallucinations can fabricate facts using real data
- Logs, memory, and summaries create long-tail exposure

This guide breaks down the risks of each major platform and explains what you can do to keep your data safe — while still enjoying the productivity boost you're using them for in the first place.

The Risk Matrix — Copilot vs. Gemini vs. ChatGPT

Risk Factor	Microsoft Copilot	Google Gemini	ChatGPT (OpenAI)
Data Ingestion Scope	Full Microsoft 365 portfolio (Teams, Outlook, SharePoint, OneDrive)	Workspace content (Docs, Gmail, Sheets, etc.)	User-submitted prompts; API integrations vary
Default Access Risks	Inherits existing M365 permissions where misconfigurations can surface sensitive data	Depends on workspace permissions — defaults can be broad	Public tool has no access control; Enterprise offers admin control
Data Retention and Training	Customer data not used for training; session-based memory	Enterprise content not used for training	Free/Plus may train on inputs; Enterprise: no training by default
Auditability and Visibility	Strong via Microsoft Purview	Admin console logging; improving but limited	Chat history for users; Enterprise has admin tools
Governance and Controls	Deep integration with Microsoft security stack	Policy control evolving; fewer granular features	Varies widely by deployment; Enterprise plans offer controls

Big Picture Risks

Copilot

Secure by design, but risky if your Microsoft permissions or labeling are sloppy

Gemini

Strong potential, but audit and access control features are still maturing

ChatGPT

Most flexible but also the most risky if you use the public version; safest when deployed in Enterprise mode

How to Safely Embrace GenAI Without Ignoring the Risk

AI tools aren't optional today, but data security isn't either. Here's how to make sure they can coexist in (almost) perfect harmony:

1. **Map usage before something breaks**
   Inventory what tools are in use, and what types of data they touch.

2. **Fix your labeling and access gaps**
   Apply semantic classification to detect real sensitivity and automate labeling.

3. **Monitor usage, not just permissions**
   See how data is being accessed and summarized, not just where it's stored.

4. **Choose enterprise-grade or nothing**
   Public tools are great for experiments, but not for regulated data. Use ChatGPT Enterprise, as opposed to free accounts.

5. **Build AI into your DSPM stack**
   Govern data posture with tools that understand AI-powered workflows.

Pro Tip

Watch for hallucinated content that appears trustworthy. AI-generated summaries can be confidently wrong, using confidential inputs you didn't know were exposed.

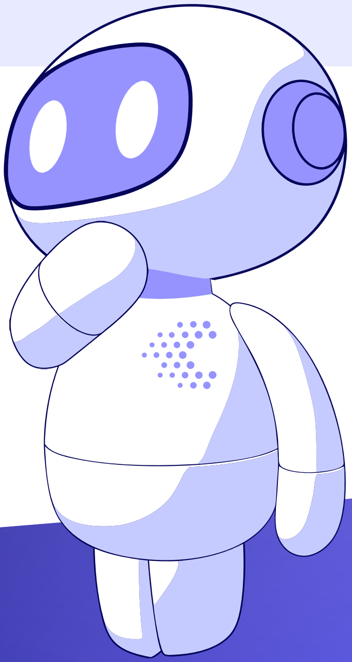

AI tools move fast. Data risks move faster.
You need visibility and control that can keep up with both.

Want to know how your AI deployments stack up?

Book a demo and we'll talk about your GenAI risk surface and see where you need to be.

Book a Demo ↘