

BIGID COPILOT ACCESS-RISK PLAYBOOK

The Copilot Access-Risk Playbook

What to audit before you turn AI loose on enterprise data, what's actually at risk, and how to close the gap before it becomes an incident.

WHAT'S INSIDE

- 1 Why Copilot doesn't create new risk, it exposes the access risk already there
- 2 Six things to audit before AI gets a single query against your data
- 3 The exclusion lists most security teams forget to keep current
- 4 Three direct questions to answer before you flip the switch

The risk nobody scoped before launch

Copilot, and every enterprise AI assistant like it, doesn't create new data risk. It exposes the access risk that was already there.

Before Copilot, an over-permissioned SharePoint folder or a stale access grant was a slow-moving problem that someone might eventually stumble onto. Copilot removes the stumbling: it's built to surface exactly what a user has access to, summarized, synthesized, and handed back in seconds, across every file, email, chat, and document that permission technically allows. That makes the assistant less a vulnerability than a magnifier, the thing that finally makes years of access sprawl queryable in plain English.

This is why most Copilot rollouts that go wrong don't go wrong because of the model. They go wrong because the access layer underneath it was never actually audited, just assumed.

What "audit before enabling" actually means

This goes beyond a standard security review. It comes down to six specific things to check, in order, before AI gets a single query against enterprise data.

THE SIX CHECKPOINTS AT A GLANCE

- 1 Map what the AI can actually reach, not what it's supposed to reach
- 2 Find the sensitive data sitting in places nobody classified
- 3 Check what happens at the boundary of permissions
- 4 Decide what AI should never touch, and enforce it technically
- 5 Plan for what the AI generates, not just what it reads
- 6 Build the control loop, not a one-time audit

Each checkpoint builds on the one before it. Skipping ahead just moves the gap downstream, where it's harder to see and more expensive to fix.

1 Map what the AI can actually reach, not what it's supposed to reach

The risk: Most organizations can describe their intended access model. Almost none can describe the actual one, the permissions that accumulated over years of project shares, departing employees whose access never got revoked, and "just give them edit access for now" decisions nobody circled back on.

AUDIT STEPS

- ☐ Inventory every data source Copilot (or the AI tool in question) can query: SharePoint, OneDrive, Exchange, Teams, and any connected third-party repositories
- ☐ For each source, identify actual permission grants, not policy documentation, the real ACLs
- ☐ Flag broad-access grants: "Everyone," "Everyone except external users," and similarly wide groups
- ☐ Identify orphaned permissions tied to inactive or departed accounts
- ☐ Cross-reference access against actual usage: who has access nobody's used in 12+ months

How to fix it: Run access reviews specifically scoped to AI-reachable data before launch, not as part of the standard annual cycle. Treat any broad-access grant on sensitive content as a blocker, not a backlog item.

2 Find the sensitive data sitting in places nobody classified

The risk: Copilot doesn't know the difference between a meeting agenda and a file with customer PII in column G. If sensitive data lives in an unclassified spreadsheet, a forwarded email thread, or an old project folder, the assistant will treat it exactly like every other document it can see.

AUDIT STEPS

- ☐ Run discovery and classification across all AI-reachable repositories before enabling access
- ☐ Prioritize unstructured data: file shares and email are where sensitive data hides longest unclassified
- ☐ Identify regulated data specifically: PII, PHI, financial data, source code, and anything with contractual confidentiality terms
- ☐ Flag sensitive data sitting in broad-access locations identified in step one, this is the actual exposure, not just a theoretical one
- ☐ Establish a recurring classification scan, not a one-time pass, since new content lands daily

How to fix it: Classification has to run before access does. An AI tool layered on top of unclassified data is the same old risk category, just attached to a much faster retrieval engine.

3 Check what happens at the boundary of permissions

The risk: Permission inheritance is rarely clean. A file shared three reorgs ago, a guest account never removed, a group membership that grew without review, these are exactly the conditions where Copilot will return something a user technically can open but never should have been able to find.

AUDIT STEPS

- ☐ Test actual retrieval: have a sample of users query Copilot for content adjacent to their role and see what surfaces
- ☐ Check guest and external user access specifically; confirm it's excluded from AI retrieval scope or tightly scoped if not
- ☐ Audit group-based permissions for groups that have grown past their original purpose
- ☐ Confirm departing employees' access is revoked immediately, not on a standard offboarding cycle, while AI retrieval is active
- ☐ Identify any service accounts or bots with standing access broader than their function
- ☐ Maintain a current inventory of every AI agent operating in the environment, including ones that weren't formally provisioned, and track what each one actually did with the data it touched

How to fix it: Run retrieval tests before general rollout, with real users and real queries, not synthetic test accounts. The gap between "permission exists" and "permission should exist" only shows up when you actually query it.

4 Decide what AI should never touch, and enforce it technically

The risk: Some data shouldn't be in scope for AI retrieval regardless of who has access to it: legal hold material, HR investigations, M&A documents, security incident records. Permission to view a file isn't the same as appropriateness for AI summarization and synthesis.

AUDIT STEPS

- ☐ Identify categories of data that should be excluded from AI retrieval regardless of access level
- ☐ Confirm exclusion is enforced at the data or policy layer, not via a written guideline employees are expected to self-police
- ☐ Apply sensitivity labels that AI tools respect (most enterprise AI assistants honor existing sensitivity labeling, if it exists and is applied)
- ☐ Test that excluded categories actually don't surface, don't assume the label works until you've verified retrieval against it

- ☐ Build a process for adding new exclusion categories as new risk arises (active litigation, new investigations, deal activity)
- ☐ Monitor what users actually share with Copilot and agents in prompts, and confirm sensitive data can be blocked at the point of sharing, not just at the point of retrieval

How to fix it: Sensitivity labeling and DLP policy have to exist and be enforced before AI rollout, not configured reactively after the first incident. If labeling doesn't exist yet, this is the forcing function to finally do it.

WHY EXCLUSION LISTS DRIFT OUT OF DATE FAST

Most exclusion policies are written once, against a snapshot of where sensitive data lived at the time. New legal holds open, new deal rooms get created, new investigations start, and none of it automatically updates the exclusion list. The list is only as good as the discovery and classification feeding it.

5 Plan for what the AI generates, not just what it reads

The risk: Governance usually stops at the input side. But Copilot outputs, summaries, drafts, generated insights, become new artifacts that get saved, shared, and forwarded, often without inheriting the sensitivity or access restrictions of the source material.

AUDIT STEPS

- ☐ Determine whether AI-generated content inherits sensitivity labels from source material, or whether that has to be configured separately
- ☐ Check where generated outputs are stored by default, and who can access that location
- ☐ Confirm retention policy covers AI-generated content and conversation logs, not just primary data
- ☐ Identify whether generated summaries of sensitive data are themselves treated as sensitive
- ☐ Audit logging: confirm you can see what was queried, by whom, and what was returned, at least for sensitive data categories

How to fix it: Treat AI output as a new data type requiring the same lifecycle (classification, retention, access control) as the source it was generated from. Most rollouts miss this entirely.

6 Build the control loop, not a one-time audit

The risk: Access, data, and usage all drift continuously. A clean audit on launch day says nothing about month six, when new files have landed, new employees have joined, new groups have formed, and nobody's re-checked any of it.

AUDIT STEPS

- ☐ Establish recurring re-audits of access, classification, and retrieval testing, not a single pre-launch pass
- ☐ Set alerting for anomalous AI retrieval patterns: unusual volume, access to categories a user hasn't queried before, off-hours activity
- ☐ Assign clear ownership for responding to flagged AI access incidents
- ☐ Report on AI access risk to leadership on the same cadence as other security metrics, so it doesn't quietly fall off the radar after launch
- ☐ Revisit exclusion lists and sensitivity labels as new risk categories emerge

How to fix it: The teams that get this right don't treat the audit as a launch gate they clear once. They treat it as a standing control that runs continuously, because the AI tool didn't change, the data underneath it did, and it keeps changing.

6 things to audit before launch, then re-audit on a standing cadence. The teams that get hurt are the ones who treated this as a one-time gate instead of a continuous control.

Before you flip the switch

Three questions worth a direct answer before enabling AI over enterprise data, not after:

If you ran a Copilot query right now across your riskiest data category, what would come back, and would you be comfortable with the answer?

Can you name every piece of sensitive, regulated, or legally sensitive content currently reachable by the AI tool, and confirm whether it should be?

If something surfaces that shouldn't have, do you have a way to know it happened, and a way to stop it from happening again?

A soft answer to any of these means the access layer isn't ready, regardless of how ready the AI tool is.

This playbook maps directly to how BigID secures AI rollouts in practice. It discovers and classifies data across every AI-reachable source, surfaces the access and permission gaps that turn assistants into exposure engines, and enforces the exclusions and policies that keep AI retrieval inside the lines you intended, continuously rather than only at launch.

NEXT STEP**Talk to our AI security experts**

Before your next Copilot rollout, expansion, or renewal, our team will walk your environment through this exact playbook and show you precisely what's reachable today, what shouldn't be, and what it takes to close the gap before it becomes an incident.

- A map of what your AI tools can actually reach today, versus what they're supposed to reach
- A prioritized list of the access and exclusion gaps to close first, ranked by exposure
- A plan for the re-audit cadence that keeps the access layer current as your environment changes

CUSTOMER VALIDATION

Proof in Action

"BigID gives me better visibility into sensitive data, helps prioritize security protections, reduces attack surfaces, strengthens compliance, and increases operational efficiency overall — a strategic pillar of our AI-First Cybersecurity Transformation"

- CISO, Global Healthcare Company

"We needed to automate processes and data management across systems for the data of a few million customers across a lot of systems and data. BigID was the one solution that did this in the most efficient and sophisticated way - and had more use cases we could add on moving forward."

- Chief Privacy Officer, Global Telecoms Company

ANALYST RECOGNITION

Validated by Industry Leaders

- 2025 Company of the Year for AI Governance - Frost & Sullivan
- Named a Leader across all DSPM research including Omdia DSPM Universe, CB Insights for DSPM, TAG, GigaOm, Frost & Sullivan
- Named a Leader in all Privacy Management evals including Forrester Privacy Management Wave & IDC Privacy Compliance MarketScape
- BigID named a Leader across 4 GigaOm evals: DSPM, Data Security Platforms, Unstructured Data Management, Data Access Governance
- Named in Gartner's 2025 DLP Market Guide
- Named in Gartner's 2025 AI TRiSM Market Guide
- Represented 30+ times in Gartner's 2025 Hype Cycles for Security, Privacy, AISPM

FREE RISK ASSESSMENT

Discover Your Data & AI Risk

An agentless, cloud-native assessment to enable security, privacy, and compliance: pinpoint where sensitive data and AI risks live - and how to fix them.

What You Get:

- Complementary 2 week engagement on real data to gain visibility on risk across data and AI
- Executive-ready report identifying high-risk data and AI across cloud environments with next-step recommendations

NEXT STEPS

Ready to Connect the Dots in Data & AI?

Schedule your free risk assessment or ask for a deeper demo across DSPM, DLP, AI SPM, Privacy Automation, AI Risk, and Data Lifecycle Management

► www.bigid.com/demo

► www.bigid.com/risk-assessment

► info@bigid.com



BigID Next

"The features they offer are the moat that they have built around the product. This is what makes them the market leader in Data Privacy and Protection."



IT Security & Risk Management Associate - Banking

Gartner
Peer Insights.

Gartner® and Peer Insights™ are trademarks of Gartner, Inc. and/or its affiliates. All rights reserved. Gartner Peer Insights content consists of the opinions of individual end users based on their own experiences, and should not be construed as statements of fact, nor do they represent the views of Gartner or its affiliates. Gartner does not endorse any vendor, product or service depicted in this content nor makes any warranties, expressed or implied, with respect to this content, about its accuracy or completeness, including any warranties of merchantability or fitness for a particular purpose.



BigID Next

"Exceptionally powerful and easy to use tool that integrated smoothly into our environment and delivered value from day one. The BigID team is fabulous to work with."



Senior Vice President & Chief Data Officer - Banking

Gartner
Peer Insights.

Gartner® and Peer Insights™ are trademarks of Gartner, Inc. and/or its affiliates. All rights reserved. Gartner Peer Insights content consists of the opinions of individual end users based on their own experiences, and should not be construed as statements of fact, nor do they represent the views of Gartner or its affiliates. Gartner does not endorse any vendor, product or service depicted in this content nor makes any warranties, expressed or implied, with respect to this content, about its accuracy or completeness, including any warranties of merchantability or fitness for a particular purpose.