THE AI ACCESS AUDIT:

5 CHECKS

EVERY BUSINESS SHOULD RUN BEFORE THEIR DATA IS ALREADY GONE

AI tools are already in your business. *Do you know what data they can see?*

Cased Dimensions

The Problem Hiding in Plain Sight

Your team is already using AI. Copilot, ChatGPT, Gemini and whether it's been formally deployed or not, these tools are in the building and your people are using them to draft emails, summarise meetings, write reports, and pull together information faster.

Every time they do, those tools are potentially accessing business data. Client records. Financial documents. HR files. Board papers. Commercial contracts.

Most organisations have no idea what that actually looks like inside their organisation.

The Challenge

Most leadership teams don't have a clear picture of what AI can actually see inside their organisation, or what happens to data once staff start pasting it into AI tools.
The assumption is that because AI follows existing permissions, the data must already be secure. It isn't. AI exposes the weaknesses that were already in your environment, and makes them visible in ways most businesses aren't prepared for.

AI is already inside your business. The question is what it has access to. This guide gives you five checks to run inside your own environment, and tells you what to do if the answers aren't what you expected.

SCORE YOUR BUSINESS

Before you can protect your business, you need to see where AI can reach.

Most organisations have a general sense that AI tools are being used across the business. Very few have mapped what those tools can actually access, or what staff are doing with them.

These five checks close that gap.

Your AI Access Audit		
1	Are AI tools only surfacing data your staff are genuinely meant to access?	Y/N
2	Is all AI activity happening on governed, auditable platforms?	Y/N
3	Is your sensitive data labelled so AI knows what to avoid?	Y/N
4	Can you produce a full record of AI activity on demand?	Y/N
5	Do your staff know what they can and can't use AI for?	Y/N

These five checks draw on guidance from Gartner's research on AI governance, the UK National Cyber Security Centre's AI security guidance, and Microsoft's own documentation for Microsoft 365 Copilot deployment.

CHECK 1: WHO CAN ACCESS WHAT?

Permissions & Access

Copilot respects existing permissions. That sounds reassuring until you look at what those permissions actually allow.
In most organisations, access has been granted informally over years between shared folders, "everyone" links, leavers never removed. Copilot surfaces whatever sits behind them.

The Problem

Unsecured salary data, Board papers, Client contracts and M&A documents can all be accessed by AI.

One of the most common questions staff ask Copilot is "What is the CEO's salary?"

Try this now

1. Open the Microsoft 365 admin centre
2. SharePoint admin
3. Active sites. Count how many sites show permission level "Everyone" or "Everyone except external users."
4. Then ask a non-IT member of staff to search their OneDrive or Teams for terms like "salary," "confidential," or "board."

What they find is what Copilot can surface too.

CHECK 2: WHERE IS YOUR AI DATA GOING?

Keep AI Inside the Tennant

Consumer AI tools offer no governance and no audit trail. The problem is that unsanctioned AI use is almost certainly happening inside your business right now; and most leadership teams have no visibility into it.

The Risk in 3 Parts

1. **No record:**
   Every time a staff member pastes a contract, client list, or spreadsheet into ChatGPT or Gemini, your data has left the tenant. You can't recall it, audit it, or prove it didn't happen.

2. **No approved alternative:**
   If Microsoft 365 Copilot isn't licensed and available, staff will use whatever they can find. Free tools fill the vacuum.

3. **No safety net:**
   Without Data Loss Prevention policies covering AI endpoints, sensitive data flows out unchecked.

Where to Look

Microsoft Defender portal → Cloud apps → Cloud app catalogue, filtered by "Generative AI."

This shows which AI tools your staff are actually using.

CHECK 3: IS YOUR SENSITIVE DATA LABELLED?

Data Classification

51% of organisations classify less than a quarter of their data.

- Gartner poll of 450 CISOs, 2023

Copilot doesn't know the difference between a sales deck and a set of board papers. Both are just files. Without sensitivity labels, it will pull from either.

Scenario A

Labels in Place:

Sensitive files carry classifications. Data Loss Prevention policies stop confidential content being surfaced in AI-generated outputs. Copilot works with appropriate guardrails.

Scenario B

Labels Missing or Unused:

Copilot pulls freely from board papers, HR files, and financial records into answers, summaries, and generated documents. All is visible to anyone who asks a related question.

Where to Look

Microsoft Purview → Information Protection → Label analytics.

Most organisations discover they're firmly in Scenario B.

CHECK 4: CAN YOU PROVE WHAT AI DID?

Audit & Oversight

Ask Yourself 3 Questions

1. If a regulator asked what Copilot has done with your data in the last 30 days, could you answer?
2. If your auditors asked for a record of every AI interaction involving financial data this quarter, could you produce it?
3. If your board asked for evidence that AI usage is being monitored, what would you show them?

For most organisations, the honest answer to all three is: no

Why This Matters

Under GDPR, FCA rules, and DORA, you need to evidence how personal and regulated data is being processed. That applies to AI the same way it applies to any other business system.

Where to Look

Microsoft Purview compliance portal → Audit. Search for activity type "Copilot."

If that search produces nothing, your AI is running without a ledger.

CHECK 5: DO YOUR STAFF KNOW THE RULES?

Reduce Risk Through Staff Guidance

Most AI risk is accidental rather than malicious. A staff member pastes client data into ChatGPT to summarise a meeting. Someone uses Copilot to draft a board paper using source material they shouldn't have had.

Clear staff guidance is the cheapest line of defence you have; and most organisations don't have one.

The 5-Minute Test

Pick five people across different departments and seniority levels. Ask each of them:

- *Which AI tools do you use at work?*
- *What rules does the business have about using them?*
- *What data are you allowed to put into them?*
- *Where would you find the policy?*

What That Tells You

If you get five different answers, or no answers at all, your people are making individual judgment calls every day about what data goes into AI tools. Those calls will sometimes be wrong, and you won't know until something surfaces.

WHAT HAPPENS NEXT

AI productivity is real but only when your data stays protected, governed, and auditable.

If any of these checks flagged a concern, you don't need to figure it out alone.

The five checks in this guide give you visibility. Fixing what they surface is where most organisations need help, because the work cuts across permissions, data classification, governance, and staff behaviour, and it has to be done in the right order.

Cased Dimensions offers a **free AI Governance & Security Assessment**: a structured conversation to map your exposure against the five checks, identify the priorities, and tell you exactly where you stand.

No commitment. No sales pitch. A practical conversation about what's working, what isn't, and what to do first.

Book your free assessment call with Cased Dimensions and walk away with a clear view of your priorities and the right first step to take.

CONTACT US

+44 2895 900 804

info@caseddimensions.com

www.caseddimensions.com

55-59 Adelaide St, Belfast, BT2 8FE