JUNE 19, 2026

How to map your AI exposure: every tool, every action, one afternoon

List every AI tool your team uses, mark the ones that can act on their own, and decide where a human stays in the loop. A one-afternoon exercise.

By **Dave Taylor**

To map your AI exposure, list every AI tool your team actually touches, mark which ones can act on their own — send an email, write to a system, move money — and decide where a human stays in the loop. Most Irish SMBs can run the exercise in an afternoon, and most are surprised by what they find. It's worth doing because the map you think you have is already out of date: 78% of employees now bring their own AI tools to work, per Microsoft's 2024 Work Trend Index. This is the exercise the first two parts of this series kept pointing at — here's how to run it.

Why the map you think you have is wrong

Every owner has a mental list of the AI in their business — the chatbot on the site, the Copilot licence, maybe the note-taker on calls. That list is the tip of it. The tools that create real exposure are the ones nobody signed off: the marketing contractor's content generator, the sales rep's personal ChatGPT with the pipeline pasted in, the browser extension that summarises every email it can see. This is shadow AI, and it scales with how useful the tools are.

The numbers say the gap is structural, not sloppy. Microsoft found 78% of workers bring their own AI tools — 80% at small and mid-sized firms, where there's rarely an IT gatekeeper to slow it down. Lenovo's 2026 Work Reborn survey of 6,000 employees put between a fifth and a third of all AI use outside the IT function's governance entirely, with 31% of users getting no employer training at all. For an Irish SMB that's the AI-literacy duty we covered in Part 2 showing up as an operational problem: you can't govern, secure, or train people on tools you've never written down. None of this means the staff are reckless — people reach for AI because it makes the work faster, and most would use a sanctioned option if one existed. The point is narrower: the real list of AI in your business is longer than the one in anyone's head, and the gap is where the exposure hides. The map is the precondition for everything else.

Step one — list every tool that touches your work

Start with the boring sources, because they're complete. Pull the last three months of card statements and expense claims and flag every software line — AI features now hide inside tools you already pay for. Export the list of connected apps from your Google Workspace or Microsoft 365 admin console; both show the OAuth grants, the third-party apps staff have wired into company data. Then check browser extensions across the team, because that's where the quiet ones live. Finally, ask the team directly, with an amnesty: "what are you actually using to get work done?" — phrased as curiosity, not audit, or you'll get the sanitised answer. Write each tool on one row: name, who uses it, what data it sees. Don't rank yet. The goal of step one is a complete list, not a clean one — and complete almost always means longer than the owner expected.

Step two — mark what can act, not just answer

Now split the list in two. For each tool, ask one question: can it only answer, or can it act? A chatbot that drafts a reply for a human to send only answers. An agent that reads an inbox and sends the reply itself acts. The difference is the whole game, because answering produces a sentence you can ignore, and acting produces a consequence you can't.

This is what Part 1 called excessive agency — the 2026 risk the OWASP Foundation now ranks in its Top 10 for LLM applications. The danger isn't that the AI says something wrong; it's that a single manipulated instruction sends the email, updates the record, or moves the money, running under the assistant's own permissions so nobody is asked first. And the exposure is already live: 39.7% of workplace AI interactions expose sensitive data, per Cyberhaven's 2026 telemetry. Mark each acting tool with what it can reach — your CRM, your bank feed, your customer inbox — and how reversible that action is. A tool that drafts a tweet and a tool that issues a refund aren't in the same risk class, even if they run on the same model.

Step three — decide where a human stays in the loop

With the acting tools marked, rank them by blast radius: what's the worst thing this could do before anyone notices? The honest version of that question sorts the list fast. Anything that moves money, sends external messages under your name, deletes data, or makes a decision about a person needs a human checkpoint before the action completes — not a notification after. That last category is also a legal line: under GDPR's Article 22, a decision with legal or significant effect on someone can't be left to the machine alone, the same human-in-the-loop rule we reached on security grounds in Part 1 and legal grounds in Part 2. For everything low-stakes — drafting, summarising, internal search — let it run free. Adding approval gates to harmless tools just trains your team to click through every prompt without reading it, which is exactly how the dangerous one slips past.

What the map usually turns up

Run this in an afternoon and a few things surface almost every time. There's usually one agent or integration nobody remembered wiring up — a Make or Zapier automation a former contractor built that still holds live keys into a real system. There's usually one tool with far more access than its job needs: a note-taker with full calendar and inbox read, a support bot that can see the whole customer database to answer one FAQ. And there's usually a person, not a tool, doing something risky out of helpfulness — pasting client data into a free model to save ten minutes. None of these show up in a vendor's security brochure, because none of them are a product defect. They're the ordinary residue of useful tools adopted one signup at a time, which is precisely why the only way to find them is to look.

What it adds up to

The exposure map is the deliverable this whole series was pointing at. Part 1 showed how AI breaks; Part 2 showed what the EU AI Act and GDPR now require; both ended on the same question, and this is the answer — a single sheet that names every AI tool you run, marks

the ones that can act on their own, and records where a human stays in the loop. It turns a vague worry into a short, ranked list you can act on, and it usually pays for itself the moment it surfaces the forgotten agent with live keys. You don't need a bigger model or a security platform to start; you need an afternoon and an honest list. It's the exercise we run with Irish SMBs every week. If you'd rather not map your AI exposure alone, book a 30-minute working session and we'll walk your stack — every tool, every action — together.