Shadow AI is already in your company. Here is how to govern it before an audit asks.

Your team is using more AI tools than anyone has written down, each one quietly fed company data under no contract. The fix is two artifacts: an honest register and a usage policy. Here is what goes in each.

Published 2026-06-30 · A Mue guide

Count the AI tools your team uses. Now find the written list of them. For almost every company, the second number is zero.

That gap is shadow AI: AI tooling people adopt because it is faster than asking, each one quietly receiving company or client data with no one tracking which tool got what under which contract. None of it is malicious. It is just unmanaged, and unmanaged is exactly what becomes a problem the moment someone official asks.

When the gap turns into a blocker

A client sends a security questionnaire asking which AI subprocessors touch their data. A SOC 2 or ISO 27001 project opens and the auditor wants your AI inventory. A regulator points at the EU AI Act. In each case the first thing asked for is a list you do not have, and assembling it under a deadline is far more expensive than keeping it from the start.

Artifact one: the register

A register is one honest table, a row per tool. Four fields carry the weight:

- Vendor: who runs the tool and therefore who holds the data.
- Data sent: your shorthand for what actually leaves the building. Marketing drafts are low stakes; customer records or source code are not.
- DPA signed: whether a Data Processing Agreement is in place, which is what makes it lawful to send personal data to a vendor under the GDPR.
- EU AI Act risk tier: your own classification of the tool against the four bands the Act defines.

The four EU AI Act risk tiers in plain terms

- Prohibited (Article 5): uses the Act bans outright, such as social scoring or untargeted scraping of faces. The fix is not paperwork, it is to stop.
- High-risk (Annex III): AI used for hiring, credit, biometric identification or critical infrastructure. These carry the full regime: risk management, logging, human oversight and documentation.

- Limited or transparency risk (Article 50): chatbots and content generators. The core duty is to tell people they are dealing with AI or looking at AI-generated output, and that transparency obligation starts applying on 2 August 2026.
- Minimal: everything else, which is most productivity tooling. No system-specific obligations beyond the AI-literacy duty that applies to everyone.

Artifact two: the policy

The register tells you what is happening; the policy sets the rules so the right things keep happening. A useful one is short and specific to your answers: what data may go into which tools, whether personal data or source code is allowed and under what conditions, and whether a new tool needs sign-off before anyone uses it. Generic boilerplate gets ignored. A policy that bans what you said to ban and permits what you said to permit gets followed.

Why the two go together

A register with no rules records bad habits without changing them. A policy with no register governs tools it cannot see. Build both and you have the two documents almost every AI-governance conversation begins with, ready before a client, an auditor or a regulator asks rather than scrambled together after.

One honest caveat

This is a starting template and an organisational aid, not legal advice. Have the policy reviewed for your jurisdiction before you adopt it. The point of doing it now is that the cheap version, written calmly, beats the expensive version written under a deadline.

Frequently asked questions

What is shadow AI?
Shadow AI is the AI tooling people use for work without anyone formally signing off on it: a chatbot to rewrite an email, an image generator for a deck, a coding assistant in an editor. It is usually well intentioned and faster than asking for approval, but each tool quietly receives company or client data and nobody keeps a single list of which tools get which data under what contract. That gap becomes a blocker the moment a client sends a security questionnaire or you start a SOC 2 or ISO 27001 project.

How do I find the shadow AI in my company?
Start by asking, not auditing. Most people will tell you which AI tools help them if the question is framed as governance rather than blame. Pair that with the obvious signals: browser extensions, SaaS expense lines, and the AI features now baked into tools you already pay for. Write each one into a register with the vendor, the data it receives, whether a Data Processing Agreement is signed, and its EU AI Act risk tier. The list is rarely complete on the first pass, so treat the register as a living document, not a one-time inventory.

Do small companies need an AI usage policy?
Yes, and a short one is enough. Without written rules, staff have to guess where the line sits, which is how confidential data ends up in the wrong tool. A one-page policy that says what data

may go into which tools, when personal data or source code is allowed, and whether a new tool needs sign-off first removes that guesswork. It is also the document a client or auditor expects to see, so writing it early is cheaper than scrambling for it under a deadline.

What is the difference between the register and the policy?
The register is a record of what is actually in use; the policy is the set of rules for what is allowed. The register answers an auditor asking what AI you run; the policy answers staff asking what they may do. You need both, because a register with no rules does not change behaviour, and a policy with no register has no idea what it is governing.

Build your shadow AI register and policy

Free, no signup: agent.mue.app/tools/shadow-ai-register-generator

agent.mue.app/articles/find-and-govern-shadow-ai-before-an-audit-asks