



BURAK GALİBA
Thoughts on AI Transformation

ESSAY · ENTERPRISE AI

Your Employees Already Picked Your AI Vendor — Under Terms You Never Signed

Banning unapproved AI tools pushes the work onto personal phones and consumer terms. Shadow AI is about which contract your data travels under — and the sanctioned route only wins if it's the easiest one.

GOVERNANCE

VENDORS

LEADERSHIP

Burak Galiba · October 7, 2026 · 5 min read

<https://burakgaliba.com/blog/shadow-ai-is-a-contract-problem>



SCAN TO READ ONLINE

Somewhere in your company this morning, someone pasted a customer email into an AI tool to draft a reply faster. Maybe it was the tool you licensed. Quite possibly it was a personal account on their phone, because it was open, it was quicker, and it gave a better answer than the one IT approved.

Most organizations file that under discipline. Someone broke a rule, so the response is a policy reminder, a blocked domain, and a mandatory module about acceptable use. I think that framing is wrong, and it is wrong in a way that makes the actual risk worse.

The ban that moves the problem out of sight

Blocking AI sites on the corporate network does one reliable thing: it moves the work to devices you can't see. The employee who was using a consumer chatbot in a browser tab now uses the same chatbot on a personal phone, and copies the result back by hand. The usage didn't stop. Your visibility into it did.

The reason is not that people are reckless. It's that the work is real and the deadline is real, and the tool helps. When the sanctioned option is slower, weaker or harder to get to, a ban is a request that people do their job worse on purpose. Most of them, reasonably, decline.

So the honest starting assumption is that shadow AI is already happening at scale in your organization, regardless of what the policy says. The useful question isn't how to stop it. It's what it's actually costing you.

It's a question of which contract the data travels under

Here is the part that gets lost in the acceptable-use framing. When your company signs an enterprise agreement with an AI vendor, a large part of what you are paying for is the contract: how long prompts and outputs are retained, whether your content can be used to train models, which sub-processors may touch it, where it is stored, and what happens when you leave.

A personal account comes with none of that. It comes with consumer terms — terms written for an individual, which on many consumer plans allow longer retention and, by default or unless the user opts out, the use of conversations to improve the vendor's models. Those terms also change. The major vendors have all rewritten their data policies in the past two years, some more than once, and an employee's personal account is not a contract anyone in your legal team is tracking.

That's the real exposure. The customer email wasn't leaked by being sent to an AI. It was sent to an AI under terms your organization never agreed to, never reviewed and can't enforce. Shadow AI is a vendor relationship you didn't choose, governed by a contract you didn't sign, holding data you're accountable for.

Seen that way, the goal changes. You are not trying to reduce AI usage. You are trying to move as much of it as possible onto the contract you actually negotiated.

The sanctioned route has to win on convenience

That only happens if the approved option is the easiest one. Not the most compliant, not the most secure on paper — the easiest. Every extra click, every access request, every week waiting for a license is a reason to open the phone instead.

In practice that means a few unglamorous things:

- **Access by default, not by request.** If someone has to justify why they need the approved AI tool, you have built a queue in front of the safe option and left the unsafe one open.
- **A tool that's actually good.** A sanctioned assistant on an older or restricted model, with half the features turned off, loses to the consumer version every time. People can tell the difference within a day.
- **Clear lines, stated once, in plain language.** Which data may go into the approved tool, which may not go into any tool, and who to ask when it's unclear. A one-page rule people remember beats a twenty-page policy they clicked through.
- **A fast way to say yes to new tools.** When a team finds something genuinely better, there should be a route from "we'd like to use this" to "it's on contract" measured in weeks, not quarters. Otherwise the next shadow tool is already chosen.

None of this is a security project dressed up as enablement. It is enablement, and it is the most effective security control you have for this particular risk.

Trained, sanctioned, and measured

The question I'd ask of any organization is whether AI use across the workforce is trained, sanctioned and measured — all three, not one.

Trained means people know how to use the tool well and what not to put into it. A single launch webinar doesn't count; people learn this from examples in their own work, repeated.

Sanctioned means there is an approved tool, on a negotiated contract, that people can actually reach today.

Measured means you know whether they use it. Not survey answers about whether they use AI "responsibly" — usage. How many people signed in this month, which teams, how often.

Most organizations I'd expect to score one out of three. Training happened once; a tool was licensed; and nobody looked at the numbers again. The third leg is the one that tells you whether the first two worked, and it's the one most often missing.

Where to start

Take two numbers: active monthly users of your sanctioned AI tools, and the headcount of the teams who were given access. The gap between them is a rough lower bound on your shadow AI — the people who have a reason to use AI at work and are doing it somewhere else. It isn't precise, but it is honest, and it is far more useful than asking people whether they've ever used an unapproved tool.

Then ask the people in that gap why. The answers will tell you exactly what to fix: the tool is too slow, the access request was never answered, the model is worse, they didn't know they had it. Every one of those is cheaper to fix than the data you are currently handing over on someone else's terms.

And when you've moved the work onto your own contract, [put it on the inventory](#) with everything else. Shadow AI is, by definition, the AI that never made it onto the list.

I built a free [AI Readiness Score](#) that asks this directly — whether AI use across your workforce is trained, sanctioned and measured — alongside 19 other questions across pilots, data, talent and governance. It takes about ten minutes, and you don't need to give an email to see the result.

Read the essay online at <https://burakgaliba.com/blog/shadow-ai-is-a-contract-problem>

I built a free AI Readiness Score — 20 questions across pilots, data, talent and governance, about ten minutes, no email required: <https://burakgaliba.com/readiness>

© 2026 Burak Galiba · <https://burakgaliba.com> · Views are my own.