

WHITEPAPER • LAYER 03

Usage Governance vs. Model Governance for Mid-Market AI

Why most companies need to govern how employees use AI tools before they ever govern a model.

Tesatori Technologies • AI Governance, Layer 03 • 2026

Executive summary

Ask most companies what their "AI governance" plan is, and you'll get one of two answers. Either there isn't one, or there's a policy document written for a problem they don't have — model risk, algorithmic bias audits, training-data provenance — borrowed from frameworks built for the handful of labs that actually train frontier models. Almost nobody in the mid-market is building models. Almost everybody in the mid-market has employees pasting customer data into ChatGPT.

This paper covers the distinction that matters for nearly every mid-market company: usage governance versus model governance, why the industry defaults to writing policy for the wrong one, and the risk-tiered approach we use to get a real, board-ready framework in place — usually in twelve to sixteen weeks.

The symptom: AI policy that doesn't match reality

Two things are usually true at once inside a growing company. First, AI is already in daily use — in sales, in support, in finance, in engineering — largely through commercial tools like ChatGPT, Copilot, and a growing list of AI features quietly turned on inside software the company already pays for. Second, if there's a governance document at all, it reads like it was written for OpenAI or Anthropic, not for a 200-person company that buys AI rather than builds it.

That mismatch isn't cosmetic. A policy about model bias audits and training-data lineage gives a board nothing useful to sign off on, because it doesn't answer the actual question in front of them:

what is stopping an employee from putting a client's financial records into a public AI tool this afternoon? Most companies don't have an answer, because the policy they wrote doesn't cover it.

Why the model-governance playbook doesn't fit

Model governance is the right frame if you're training, fine-tuning, or deploying your own models — weighing training-data provenance, bias testing, model cards, and the kind of technical risk assessment that comes with owning the model itself. It's the discipline built by and for frontier AI labs and the handful of companies embedding custom models deep in their product.

It's the wrong frame for almost everyone else, for a simple reason: most mid-market companies are AI buyers, not AI builders. Employees are using ChatGPT, Copilot, and a growing number of AI features embedded in existing SaaS tools — not training anything. The risk in that world isn't in the model's weights. It's in what an employee pastes into a chat box, what a vendor's AI feature does with the data flowing through it, and whether anyone in the company would know if either went wrong.

Applying model-governance frameworks here produces exactly the mismatch described above: technically sophisticated policy that answers questions nobody is asking, while the actual exposure — data leakage through everyday tool use — goes completely unaddressed.

What usage governance actually means

Usage governance starts from the buyer's vantage point, not the builder's. It asks a narrower, more practical set of questions: which AI tools are actually in use across the company, sanctioned or not; what data is flowing into each of them; which use cases carry enough risk to need review before they scale; and who is accountable for catching it when a new tool or use case shows up unsanctioned.

It doesn't require anyone in the company to understand how a transformer model works. It requires an inventory, a risk tier, and a review process — the same operational discipline a company already applies to vendor risk or data access, pointed at a new category of tool.

What it costs to skip this

The absence of usage governance rarely shows up as a single incident. It accumulates:

- Shadow AI use grows invisibly, because employees adopt tools that make their job easier long before anyone in leadership is aware they exist.
- Sensitive data leaves the company's control through tools with no contractual or technical guarantee about how that data is used, retained, or trained on.
- Legal and security get bypassed by default, not through malice but through speed — a tool gets adopted team by team faster than anyone can review it.
- The board gets asked a question nobody can answer, usually by an investor, acquirer, or regulator: what's your policy on employee AI use? "We don't really have one" is a worse answer than most executives realize until they're the one giving it.

The risk-tiered approach

A usage-governance engagement runs in three phases, and no policy gets written until the inventory and risk-tiering are complete.

Phase 1 — Audit (Weeks 1–4). We inventory every AI tool and use case already in play, sanctioned or not — the commercial tools employees use directly, and the AI features embedded inside existing SaaS platforms. Each use case gets risk-tiered against data sensitivity and business impact, and we flag every place legal, security, or compliance currently have zero visibility into how AI is being used.

Phase 2 — Design & build (Weeks 5–11). We draft model, vendor, and data-use policy scoped specifically to the highest-risk tiers — not a blanket ban, and not a document nobody will read. We build the intake process so a new AI tool or use case has somewhere specific to go for review before it scales unsanctioned, and pilot the policy with one business unit before it rolls out company-wide.

Phase 3 — Hand off (Weeks 12–16). We present the board-ready oversight framework, train reviewers on the risk-tiering process so it keeps running without us, and set the recurring cadence for policy review as new tools and use cases inevitably show up.

This is the same discipline behind every Tesatori engagement, sized to AI: audit what exists, design and build the fix directly with your team, then hand off so the oversight keeps working without us in the room.

What a working usage governance program looks like

It's worth being precise about what "governed" means here, because it's easy to either over- or under-build this.

A working program doesn't ban AI tools by default — that just pushes usage further into the shadows. It has a simple, visible intake path for a new tool or use case to get reviewed before it scales. It has risk tiers that match the actual stakes: a marketing team drafting blog copy with a public AI tool is a different risk than an analyst pasting unredacted customer financials into the same tool, and the policy should treat them differently. It names who's accountable for catching unsanctioned use, rather than assuming everyone will self-report. And it gives the board a document that actually answers the question they're going to be asked: what's stopping this from going wrong, and who would know if it did.

That's a governance outcome sized to how the company actually uses AI today — not a theoretical framework borrowed from a lab building something entirely different.

What it looks like in practice

One of our clients — a private equity firm overseeing a fast-growing portfolio company — came to us because AI tools were being adopted across the business faster than legal could review any of them. Every team had its own favorite tool, nobody had a clear picture of what data was flowing where, and the board had started asking questions nobody could fully answer.

The fix wasn't a moratorium on AI use. It was an inventory of what was actually in play, a risk tier for each use case, and a policy scoped to the handful of genuinely high-risk ones — the tools touching client financials or personal data, not the ones drafting internal meeting notes. The result gave the board something concrete to review, rather than a vague assurance that things were "under control."

Signals you're already living this problem

A leadership team usually doesn't need a formal audit to know something is off. A few reliable signals:

- Employees are using AI tools that were never formally approved, and nobody's tracking which ones.
- Nobody could produce, on request, a list of every AI tool currently touching company or customer data.
- A vendor's AI feature got turned on inside existing software without anyone evaluating what data it processes.
- The board or an investor has asked about AI governance and the honest answer was closer to "we're working on it" than a document.
- Legal or security found out about a new AI use case after it was already in production, not before.

If two or more of these sound familiar, the underlying cause is very likely structural, not a compliance problem — and worth treating as one.

Getting started

Tesatori's AI Governance engagement is a fixed-scope, 12–16 week program built around exactly this method: a full AI use-case inventory and risk tiering, model, vendor, and data-use policy scoped to your actual exposure, and a board-ready oversight framework the company can run against going forward. It's fixed price at \$240,000, with no hourly clock and no ambiguity about what's delivered.

It's also deliberately the third of three layers, and rarely the place to start. Reporting exposes where the data can't be trusted; data governance fixes the ownership and access model that caused it; AI governance keeps both in place as AI tools start touching the same data. Companies are usually ready for this layer once the first two are in place — governing AI use on top of ungoverned data just moves the same underlying problem somewhere new.

About Tesatori Technologies

Tesatori Technologies is a fractional CDO/CAO advisory that builds the reporting, governance, and AI oversight layers growing companies need before the next board meeting, audit, or AI rollout forces the issue. Engagements are fixed-scope and senior-led, sized specifically for mid-market companies that have outgrown informal data practices but haven't yet built a dedicated data function.

Ready to see what's actually running through your AI tools? Book a 30-minute fit call, no obligation, at calendly.com/devon-tesatoritechnologies/30min — or reach us at contact@tesatoritechnologies.com.