September 24, 2026

AI Governance Isn't a Policy, It's an Operating Discipline

TECH SPOTLIGHT

Written by: Tom Craft, Chief Growth Officer

Most companies started their AI journey with experimentation. Today, AI is writing content, generating code, analyzing data, automating workflows, and increasingly becoming part of production systems and customer experiences.

The challenge is that AI doesn't stand still. Models change. Pricing changes. Regulations change. And the ways organizations use AI change. That's why AI governance can't be treated as a one-time policy exercise.

Good governance isn't about creating a list of approved tools. It's about giving the business visibility into how AI is being used, where risk exists, and how decisions get made as the technology evolves.

One area where this becomes immediately visible is cost. Unlike traditional software licensing, AI spending can be spread across applications, APIs, agents, business units, and individual users. Usage often grows faster than governance processes can keep up.

Organizations should be able to answer a few basic questions:

- Where are we using AI?
- Which models are we using?
- What are we spending?
- Are we getting value from that investment?

The goal isn't to use the cheapest model. It's to make intentional decisions instead of discovering costs after they've accumulated.

The technology itself creates another governance challenge. Choosing an AI model is no longer a decision you make once and revisit years later. New models regularly deliver better performance, lower costs, and new capabilities. At the same

time, providers retire older models and require migrations.

Organizations need a process for evaluating new models, testing changes, and understanding the impact those changes may have on existing applications. Even with no code changes, changing the underlying model can change application behavior. That reality needs to be managed, not discovered by accident.

Then there's data.

Before deploying AI, organizations need clear answers about what data is being shared, where it is processed, how it is retained, and what controls exist around it. The details vary across providers, deployment models, and use cases. A cloud-hosted model introduces different considerations than a locally deployed one. Sensitive data requires different controls than public information.

The same is true for regulation. Frameworks such as the EU AI Act are beginning to formalize expectations around AI risk and oversight. As regulations mature, governance will become increasingly important for organizations operating across multiple markets.

The answer, however, isn't bureaucracy. Organizations can create so many approvals, reviews, and governance checkpoints that AI adoption slows to a crawl. That's just as problematic as having no governance at all. The goal is to create enough structure to manage risk without preventing innovation.

For most organizations, that starts with a few practical capabilities:

- Inventory AI use cases.
- Establish clear data policies.
- Monitor usage and spend.
- Evaluate models before production deployment.
- Periodically review vendors, applications, and controls.

Frameworks like NIST's AI Risk Management Framework can provide useful guidance, but the objective isn't compliance for compliance's sake. The objective is adaptability.

Because the AI landscape will continue to evolve. Models will improve. Costs will shift. Regulations will expand. New use cases will emerge. Organizations that treat AI governance as an ongoing operating discipline, rather than a one-time policy exercise, will be in a much stronger position to take advantage of those changes while managing the risks that come with them.