Labor · Essay

Program Management Is Becoming an AI Governance Discipline, Whether We Call It That or Not

AI governance keeps landing on program managers' desks whether the org chart says so or not. A case for treating the discipline as architecture, not administration.

Jennifer McKinney · August 8, 2026 · 6 min read

AI governance is often described as a new function. That is partly right. New laws, new tools, and new model risks have created work that did not exist in its current form a few years ago.

But much of the work is landing somewhere more familiar: program management.

Someone has to turn broad principles into a sequence of decisions. Someone has to determine who owns the model inventory, who signs off on a vendor, what happens when a product team wants an exception, how legal guidance reaches an engineering backlog, and whether the organization can demonstrate what it decided six months later. Those are governance questions. They are also core technical-program-management questions.

I have spent my career in that space between strategy and execution. A good TPM is not simply a project tracker. The job is to make complex work legible across teams with different incentives, vocabulary, time horizons, and risk tolerance. AI systems amplify the need for exactly that capability.

The reason is simple. AI governance is not a department-shaped problem.

Legal understands the obligations and contract terms. Privacy understands data use and purpose limitation. Security understands attack surfaces and access. Product understands the feature and user experience. Data science and engineering understand model behavior and deployment. Procurement understands the vendor boundary. HR may own a system whose outputs affect workers. Communications may be responsible when it fails publicly. No single team can govern the whole lifecycle from inside its own silo.

That is why the work so often becomes a program.

The IAPP's 2025 AI Governance Profession Report found responsibility distributed across privacy, legal and compliance, IT, data governance, ethics and compliance, and security, rather than concentrated in a single function. It also found that only 10 of 671 respondents said they would not need additional staff in the next 12 months. I do not treat that as a precise headcount forecast. I treat it as a clear organizational signal: the work is broad, and organizations do not believe their current capacity is enough. IAPP

The bad response is to create an "AI governance" label and assume coordination will appear. Titles do not create operating mechanisms. Teams need a shared decision model.

In practice, that means building a small number of durable artifacts. A system inventory that identifies the use case, owner, data categories, deployment context, vendor dependencies, and material changes. A risk register that distinguishes concerns requiring action from concerns requiring monitoring. Review thresholds that define when human approval is necessary. An exception process that identifies who can accept what risk and for how long. Evidence that connects a decision to the policy, test, contract term, or design control behind it.

None of these artifacts is glamorous. That is precisely why they matter. They turn governance from a periodic workshop into a system of work.

NIST's Generative AI Profile recommends that organizations proactively incorporate trustworthy characteristics into system requirements. A TPM can translate that instruction into calendar events, backlog acceptance criteria, design-review gates, and accountable owners. The work is not to paraphrase the guidance. The work is to make it executable before a launch makes the architecture difficult to change. NIST

This changes the craft of program management.

Traditional programs often use milestones as a proxy for progress. Requirements complete. Build complete. Testing complete. Launch complete. That model can fail with AI because a launch is not the final state of the system. Models change. Vendors change their terms. Inputs drift. People discover uses that were never described in the original business case. A program manager working on AI needs to think in feedback loops, not just delivery phases.

The governance cadence must survive launch. What is monitored? How do users report a problem? What level of behavior change triggers reassessment? What evidence is kept after an exception expires? Who notices when a vendor's "minor update" materially changes how a system behaves? These are operations questions with real ethical and labor implications.

Hiring systems make the stakes obvious, but the same logic applies to internal productivity tools. When an assistant prioritizes work, summarizes performance information, drafts a recommendation, or automates part of a workflow, it changes the distribution of authority inside

the company. It can make knowledge workers faster. It can also make their judgment less visible, their work harder to contest, or their errors easier to scale.

That is why I resist the idea that governance is only about stopping AI. Good governance is a way to decide where automation is appropriate and what conditions preserve human agency. It should let an organization move faster on low-consequence, well-bounded work while demanding more review when systems affect employment, access, health, finance, or rights.

The labor implication is significant. AI does not eliminate the need for program management. It changes the work that organizations need program managers to do. The role becomes more architectural, more risk-literate, and more concerned with decision rights. A TPM who can facilitate a dependency review but cannot ask where the data came from, who has authority to override the model, or how an affected person gets recourse will be missing part of the system.

Organizations should respond with structure, not vague demands for everyone to “be responsible.” They need named owners for use cases, a cross-functional review mechanism with actual authority, and TPM capacity dedicated to the operating model rather than borrowed only at launch. They need to treat documentation as evidence infrastructure, not as compliance theater. And they need to fund the connective work: intake, review coordination, training, change management, monitoring, and incident learning.

There is evidence that organizations are already looking outside their existing structures for this capacity. In the 2026 ACC Chief Legal Officers Survey, 47% of CLOs said CEOs wanted them to develop technology and AI proficiency; 48% reported using outside counsel and 27% consultants, while 63% expected stable internal headcount. The data is not specific to TPM staffing, but it describes a familiar pattern: new responsibility arrives faster than organizations can redesign their permanent teams. ACC Chief Legal Officers Survey

The long-term answer cannot be an endless layer of external review. Governance has to become part of how an organization runs programs. That does not mean every TPM becomes a lawyer, a data scientist, or a security architect. It means the discipline must become fluent enough to convene those experts around the right decisions, early enough to matter.

This is where the best program managers have always added value. They make ambiguity actionable. They identify a decision hiding inside a dependency. They force a vague concern into an owner, a threshold, a date, and an evidence trail. In AI, those skills are no longer just delivery hygiene. They are part of how an organization decides whether its systems deserve trust.

Program management is becoming an AI governance discipline whether we call it that or not. The organizations that recognize it will build the role accordingly. The ones that do not will keep discovering, too late, that “someone should have owned this” was the governance plan all along.

Sources

- IAPP
- NIST
- ACC Chief Legal Officers Survey