AI Governance Is an Operating Model, Not a Policy

The challenge is not simply deciding which AI tools employees can use. It is establishing how AI should be used, who remains accountable, and how organizations know it is working.

By Ravinder Tulsiani, DBA

November 16, 2026

An organization approves an AI assistant to help employees draft documents. The security review is complete, the tool is authorized, and employees receive guidance on acceptable use.

Over time, the same tool starts being used for more consequential work. Employees ask it to interpret policies, summarize sensitive information, evaluate options, and recommend decisions.

The technology may be the same. The organizational risk is not.

Who approved those additional uses? Who checks whether the recommendations are reliable? When must an employee challenge the output? Who has the authority to stop its use when something goes wrong?

These are not questions that can be settled through a technology policy alone.

AI governance becomes meaningful when an organization can translate its principles into operational decisions and responsibilities, with controls and evidence to show they work.

That makes it an enterprise leadership responsibility involving technology, legal, privacy, risk, operations, human resources, and workforce capability.

Start with the work, not the technology

The first governance decision should not be which AI platform to approve.

It should be what organizational problem AI is expected to solve, whether AI is appropriate, and what consequences its use could create.

Consider two applications.

One team uses AI to draft an internal meeting summary that an employee reviews before circulation. Another proposes using AI to recommend decisions affecting people's employment, services, or access to resources.

Both involve AI, but their consequences, oversight requirements, and acceptable levels of risk are fundamentally different.

An organization therefore needs visibility into its AI use cases, not merely an inventory of its AI applications.

For each material use case, leadership should understand the intended outcome, information involved, people affected, decisions influenced, and consequences if the system produces an unreliable result.

This also creates an important opportunity to challenge the investment itself.

If the underlying process is poorly designed, adding AI may simply accelerate a process that should have been changed or eliminated.

Governance should establish whether the proposed use deserves investment before determining how to control it.

Make accountability explicit

AI governance committees can coordinate decisions and provide oversight. They cannot substitute for accountability within the business.

Every material AI use case needs an accountable business owner who understands the intended outcome and accepts responsibility for how the technology is incorporated into operations.

That does not mean business leaders independently approve security, privacy, or legal risks. Those responsibilities remain with the appropriate specialists and decision authorities.

The operating model must distinguish among several responsibilities:

- Enterprise leadership establishes risk appetite, strategic priorities, and escalation expectations.
- Technology and security functions assess their technical requirements. Privacy and legal functions assess theirs.
- Business owners are accountable for the use case, operational decisions, and intended performance outcomes.
- Managers ensure that employees follow approved procedures and that problems are identified and escalated.
- Workforce and organizational development functions help employees build the judgment and practical capability they need.

These responsibilities must connect through defined approval and escalation processes.

When an AI system produces a questionable recommendation, the organization should not have to assemble a new committee to determine who can intervene.

Accountability must exist before the incident, not be negotiated afterward.

Apply governance in proportion to the risk

A single approval process for every AI application creates two problems.

If the process is too restrictive, routine low-risk improvements become unnecessarily difficult. If it is too permissive, consequential applications may proceed without sufficient scrutiny.

A better approach is to classify use cases based on the risks they introduce.

I saw this in the education technology portfolio I chaired. Requests to try a new AI tool had built into a backlog of about three months. We replaced a single approval route with tiered assessment by risk. Low-risk requests were those used to develop assets with no company-sensitive data exposed. Articulate 360 and Adobe AI features fell in that group, and cleared in days. Requests that involved sensitive data went for deeper review. Salesforce fell in that group in this case, because it involved patient data. IT and legal led that review to protect sensitive data. The point was not to approve faster for its own sake. It was to match scrutiny to risk, so routine work could proceed and consequential uses received proper attention.

The relevant considerations include the sensitivity of the information, potential harm to individuals, reversibility of the outcome, reliability of the system, and degree of autonomy it is permitted to exercise.

For decisions affecting employment or services, fairness to the people affected must be part of that assessment.

An AI assistant that drafts non-sensitive internal communications may need relatively straightforward controls and human review.

An AI system that recommends consequential decisions may require more extensive testing and safeguards. Some proposed applications should not proceed at all.

AI agents introduce additional considerations because they may do more than generate information. An agent that can access records, send communications, change information, or initiate transactions needs clearly defined permissions, operational boundaries, and intervention mechanisms.

The governing principle is straightforward:

Approval of an AI tool does not constitute approval of every way that tool might be used.

Organizations need a process for reconsidering risk when an approved technology is used for a new purpose or gains new capabilities. Approval is not permanent.

Put controls into the workflow

Policies tell employees what is expected. Operational controls determine whether those expectations can actually be met.

Suppose an organization requires a human to review AI-generated recommendations before a decision is finalized.

That sounds reasonable. But what does the reviewer actually receive?

Can the person inspect the evidence supporting the recommendation? Do they have sufficient expertise to identify a mistake? Are they given enough time to review the output? Can they reject the recommendation without being pressured to accept it?

If the answer to those questions is no, human review may exist on paper without providing meaningful oversight.

This is where AI governance intersects with organizational design.

Roles may need to change. Approval processes may need to be redesigned. Managers may require different reporting information. Employees may need practice recognizing unreliable outputs, identifying exceptions, and exercising judgment.

Controls should also be tested against realistic failure scenarios.

What happens when AI produces a plausible but incorrect answer? When sensitive information is submitted to an unauthorized system? When an agent attempts an action outside its approved scope?

The organization needs to know whether its people can respond appropriately when a system fails.

A control that has never been tested is an assumption, not evidence of protection.

Develop the capability to govern AI in practice

AI literacy matters, but general awareness training is insufficient.

Different responsibilities require different capabilities.

Employees using an approved AI assistant need to understand appropriate use, information-handling restrictions, verification requirements, and escalation procedures.

Managers need to know how to supervise AI-assisted work, recognize emerging risks, assess performance, and intervene when necessary.

Executives need enough understanding of AI's capabilities and limitations to challenge investment proposals, make informed trade-offs, and establish appropriate accountability.

The goal is not to turn every manager into an AI specialist.

It is to ensure that individuals can exercise the responsibilities assigned to them.

This also requires attention to the workforce consequences of automation. As AI takes over parts of established roles, organizations must consider which activities develop the expertise employees will still need to exercise judgment and oversight.

Responsible implementation is therefore not just about introducing new capabilities. It is also about preserving the human capabilities that remain essential.

Measure performance and control effectiveness together

Organizations often evaluate AI adoption through licence utilization, employee participation, tasks completed, or time reportedly saved.

Those measures can help describe activity. They do not establish organizational value.

If AI reduces document preparation time but increases verification effort, the net benefit may be smaller than expected. If it accelerates decisions but increases errors, appeals, or rework, the organization may be moving faster without performing better.

Each significant AI use case should have an agreed performance baseline and measures of both value and control effectiveness.

Depending on the work, these could include turnaround time, service quality, error rates, rework, cost, employee capacity, customer outcomes, detected incidents, and effectiveness of required human reviews.

Management also needs to distinguish capacity released from actual financial savings. Time saved does not automatically translate into reduced expenditure.

The decision to expand an AI application should be based on evidence that its intended benefits are materializing and its risks remain acceptable.

A technically successful deployment is not necessarily a successful organizational change.

What leadership should do first

An organization does not need to complete an enterprise-wide AI transformation before establishing useful governance.

It can begin with a small number of real use cases.

Over an initial 30-day period, I would focus leadership attention on four practical actions.

First, establish visibility. Identify AI uses already underway, including approved systems, informal experimentation, and AI functionality embedded in existing applications. Determine which uses involve

consequential decisions or sensitive information.

Second, clarify decision rights. Establish who owns each selected use case, who can authorize its operation, which specialists must review it, and who can suspend or restrict its use.

Third, examine controls in the actual work. Follow a small number of workflows from beginning to end. Test whether the stated protections, human review responsibilities, and escalation mechanisms function as intended.

Fourth, agree on performance evidence. Identify the baseline, expected improvement, potential harms, and monitoring requirements before making decisions about wider deployment. At the end of 30 days, leadership should have a named owner for each selected use case and a written record of the control gaps found.

Thirty days will not establish a mature AI governance system. It can, however, expose gaps between existing policy and operational practice, assign responsibility for closing them, and establish a disciplined basis for further implementation.

From there, organizations can strengthen their enterprise governance model and expand controls proportionately as their uses and risks change.

Governance should make responsible adoption possible

The choice facing leaders is not between unrestricted innovation and restrictive governance.

Effective governance makes responsible innovation possible by establishing where AI can be used, what conditions must be met, who has authority, and how outcomes will be evaluated.

Technology and risk specialists play indispensable roles. But they cannot independently determine how an organization's work should change, what employees must remain capable of doing, or whether an investment has delivered meaningful value.

Those are enterprise management decisions.

The real test of AI governance is not whether an organization has published a policy or established a committee.

It is whether leaders can explain what AI is permitted to do, who remains responsible for the consequences, how the controls work in practice, and what evidence justifies continued use.

That is how AI governance moves from policy to organizational capability.

Supporting references

The Canadian government documents below are public-sector guidance, not universal requirements for private-sector organizations.

- "NIST AI Risk Management Framework and Playbook" (https://www.nist.gov/itl/ai-risk-management-framework/nist-ai-rmf-playbook). Voluntary guidance for using the AI Risk Management Framework.
- "ISO/IEC 42001:2023, Artificial Intelligence Management Systems" (https://www.iso.org/standard/42001). International standard establishing requirements for an

organizational AI management system.

- "Government of Canada, Guide on Departmental AI Responsibilities" (https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/guide-departmental-ai-responsibilities.html). Published January 2026. Guidance on multidisciplinary governance and business value.
- "Government of Canada, Guide on the Use of Agentic Artificial Intelligence" (https://www.publications.gc.ca/site/eng/9.966278/publication.html). Federal guidance on agentic AI systems.

The proposed implementation sequence and executive recommendations are the author's synthesis, not prescribed steps or universal regulatory requirements.