

AI Governance Operating Model

Designing Governance as an Operating System for Scalable,
Accountable and Controlled AI



Manoj Tavarajoo



Executive Summary

AI governance is often approached as a set of principles, policies or compliance controls layered onto AI initiatives. In practice, this approach fails. It does not define how decisions are made, how accountability is assigned, or how risk is controlled in real operating environments.



AI does not fail because of models. It fails because no one is clearly accountable for what those models do.

This paper argues that AI governance must be designed as an operating model. It must function as the operating system through which the enterprise runs AI, connecting board oversight, executive accountability, business ownership, technology execution and assurance into a single, integrated system.

Organisations that rely on policy-driven governance will struggle to scale AI safely or effectively. Those that design governance into how AI actually runs will be able to accelerate decision-making, control risk and realise value at scale.

The Governance Problem

In a precedent-setting **2024 ruling**, the Civil Resolution Tribunal of British Columbia held Air Canada liable for the misinformation its AI chatbot provided to a customer. The airline argued that the chatbot was a separate entity responsible for its own actions. The tribunal rejected that argument and held the company fully accountable. The case was not exceptional. It was a signal.

AI governance has moved from theoretical concern to operational reality. Failures involving automated decisions, bias, misinformation and system behaviour now carry direct consequences in terms of liability, regulatory scrutiny and reputational damage.

Board attention has increased significantly. The Harvard Law School Forum on Corporate Governance's 2025 analysis of Fortune 100 disclosures found that explicit **board-level AI oversight had tripled year on year**. Yet governance maturity remains inconsistent.

Most organisations still operate with fragmented governance structures. Policies exist. Committees meet. Frameworks are documented. Yet **governance frequently disappears at the point of execution**.

Teams continue to make decisions without clarity on authority or escalation. Risk and compliance operate alongside delivery rather than within it. **Accountability remains diffuse**. When issues occur, organisations struggle to identify who is responsible for outcomes.



The gap is not one of awareness. It is one of design.

Why It Matters Now

The urgency is measurable.

In EY's 2025 Responsible AI survey, only 14% of CEOs believed their AI systems complied with current regulations.

The gap between organisations with strong governance practices and those without is structural, not marginal, and it widens as AI scales.

Board-level oversight is increasing, yet many boards are operating without clear visibility into how AI decisions are made, how risk is controlled or how accountability is enforced.

AI is becoming embedded in core business processes, yet governance models remain anchored in static approval mechanisms. This mismatch becomes more pronounced as systems evolve over time, interact with changing data and begin to exhibit autonomous characteristics.

Governance has become a precondition for safe and scalable AI, not an addition to it.

Where Current Approaches Fall Short

Most organisations approach AI governance through policies, principles and oversight forums. While necessary, these are insufficient.

Policies define intent but do not determine how decisions are made. Committees provide structure but do not create accountability. Frameworks describe what good looks like but do not define how it operates in practice.

Many organisations attempt to extend existing IT or digital governance models. These models assume predictable system behaviour. AI systems do not behave in this way. They evolve, drift and require continuous oversight.

As a result, governance exists as documentation rather than as a functioning system. It appears complete but fails under real conditions.

This is why governance becomes either a bottleneck or a formality. It slows decision-making without improving control, or it is bypassed entirely.

Core Principles of an AI Governance Operating Model

An AI governance operating model begins with the recognition that governance is not a control layer applied after the fact. It is the mechanism through which the organisation runs AI.


Accountability must be treated as the central organising principle.

In most organisations, ownership is distributed but accountability is not explicitly defined. This creates ambiguity at the point where decisions matter most. Effective governance requires clear accountability for outcomes, not just participation in process.



Governance must operate across the full lifecycle of AI systems.

Approval at the point of deployment is insufficient. Models evolve, data shifts and outputs vary. Governance must extend from intake through to retirement, with continuous oversight and intervention capability.



Decision rights must be explicit. Every AI system requires clarity on who approves its use, who accepts associated risk, who can intervene when thresholds are breached and who is accountable for outcomes. Without this clarity, decisions are either delayed or made informally.



Governance must connect value, risk and execution. It must enable prioritisation and scaling of high-value use cases while ensuring that risk remains within defined tolerance. A model that focuses only on risk will constrain value. A model that ignores risk will create exposure.



Assurance must be embedded into systems and processes.

Governance must produce continuous evidence of control through monitoring, auditability and traceability. Documentation alone does not provide assurance.



Governance must also be designed to remain effective as AI systems become more autonomous.



The AI Governance Operating Model

The AI Governance Operating Model Map describes governance as a connected enterprise system rather than a set of isolated functions.




Figure 1: AI Governance Operating Model Map

The model is structured across five layers. At the top, the board defines risk appetite, strategic direction and oversight expectations. This intent is translated by the executive layer into governance structures, decision rights and enforcement mechanisms. The business layer owns use cases, outcomes and value realisation. Technology and data functions are responsible for system delivery, integration and monitoring. Risk functions are accountable for controls. Audit is accountable for independent assurance.

What differentiates this model is not the layers themselves but how they are connected.

Decision flow moves across all layers, from intake through approval, deployment, monitoring and intervention. Governance is not a single decision point. It is a sequence of decisions that must be coordinated across the organisation.

At the same time, governance operates across the full lifecycle of AI systems. From initial intake and design through build, deployment, operation, monitoring and eventual retirement, governance remains active.

This operating model is designed to remain effective as AI systems evolve towards greater autonomy. As systems move from generating outputs to taking actions, the underlying requirements do not change. Decision rights, accountability and lifecycle governance remain the core mechanisms through which control is exercised.

A practical example illustrates the model in operation. In a financial services organisation deploying a credit risk model, the board defines acceptable levels of automated decision risk. Executives translate this into approval thresholds and escalation criteria. The business owns lending outcomes. Technology deploys and monitors the model. Risk and compliance define control requirements and evidence standards.

When model performance deteriorates or risk thresholds are breached, monitoring triggers escalation. Decision rights determine whether the issue can be resolved within the business or requires executive intervention. If the impact is significant, escalation moves further up the governance structure. Each step is predefined, traceable and accountable.

The model does not introduce more process. It establishes governance by design, ensuring that decisions, accountability and control operate as a coherent system rather than a fragmented set of responses.

Accountability as the Core Design Element

Accountability is the most consistent point of failure in AI governance and the most reliable indicator of maturity.



McKinsey's 2026 responsible AI maturity research found that organisations with explicit AI ownership scored 2.6 on McKinsey's maturity index, compared with 1.8 for those without it (on a scale of 0 to 4). This was identified as the single largest maturity differentiator in the survey.




Figure 2: Board to Execution Accountability Stack

The Accountability Stack defines a clear chain of accountability from board oversight through to execution. Each layer has a distinct role.

The board is accountable for oversight and risk appetite. Executives are accountable for governance design and enforcement. The business is accountable for outcomes and value. Technology is accountable for system behaviour and reliability. Risk functions are accountable for controls. Audit is accountable for independent assurance.

This structure removes ambiguity. It ensures that accountability is not shared loosely across functions but assigned clearly.

Consider a retail organisation deploying a pricing optimisation model. The business is accountable for pricing outcomes and customer impact. Technology is accountable for model performance. Risk defines acceptable boundaries and monitors compliance. Executives ensure that decision rights are enforced.

If the model produces unintended bias, the business cannot defer responsibility to the model or vendor. Technology cannot claim that outcomes fall outside its remit. Risk cannot remain purely advisory. The accountability chain ensures that responsibility is clear and enforceable.

Governance fails not when something goes wrong, but when no one can immediately identify who is accountable.

Decision-making and Governance Forums

Governance operates through decisions. Those decisions must be structured, not ad hoc.




Figure 3: AI Governance Forum Architecture

The Governance Forum Architecture defines three types of forums.

Strategic forums focus on portfolio prioritisation, funding decisions and major escalations. They ensure alignment with enterprise strategy.

Operational forums manage use case approval, deployment decisions and risk acceptance. They govern how AI systems enter and operate within the organisation.

Assurance forums monitor performance, review incidents and provide audit and compliance oversight. They ensure governance continues after deployment.

These forums are connected through defined escalation pathways. Decisions move downward for execution and upward when thresholds are exceeded.

Consider a telecommunications organisation managing AI across its network operations. A strategic forum prioritises AI use cases aligned with network performance objectives. An operational forum approves specific models for deployment. An assurance forum monitors system performance and identifies issues.

When a deployed model begins to degrade network performance, the assurance forum identifies the issue. The operational forum determines whether adjustments can be made within defined thresholds. If the impact exceeds those thresholds, the issue escalates to the strategic forum for broader intervention.

The structure ensures that decisions are made at the appropriate level and that escalation is systematic rather than discretionary.

Decision Rights and Escalation in Practice

Governance becomes effective only when it operates under pressure.




Figure 4: AI Decision Rights and Escalation Model

The Decision Rights and Escalation Model defines how decisions move through the organisation based on risk and impact.

At the first level, business and delivery teams operate within defined thresholds. They can make standard decisions without escalation.

When thresholds are exceeded, decisions move to executive governance. This includes exception approval and elevated risk acceptance.

When impact becomes critical or systemic, decisions escalate to board or risk committee level.

A healthcare example illustrates this clearly. An AI diagnostic system operates within clinically approved confidence thresholds. If model performance drops below those thresholds, clinical teams cannot proceed without escalation to a medical governance body. If patient harm is involved or systemic risk is identified, escalation moves to board-level oversight within a defined response window.

This ensures that high-risk decisions are handled consistently and with appropriate authority.

Board and Executive Questions



Boards should ask:

- ▶ Do we have a clearly defined risk appetite for AI, and is it translated into operational decision thresholds?
- ▶ Where in the organisation are AI decisions actually made, and how visible are those decisions to us?
- ▶ Who is accountable for the outcomes of each AI system currently in operation?
- ▶ What evidence do we have that governance is functioning in production, not just at approval?



Executives should ask:

- ▶ Are decision rights clearly defined for all AI use cases, including escalation thresholds?
- ▶ Is accountability for outcomes explicitly assigned and understood across business, technology and risk functions?
- ▶ Do our governance forums make decisions, or do they primarily review and defer?
- ▶ Can we trace a decision from approval through to outcome and, if necessary, intervention?



Failure Modes

Five recurring patterns appear in organisations where AI governance has been documented but not designed. They are not isolated issues. They are predictable outcomes of poor design.

- **The Ethics Committee Trap**
A forum exists to review and discuss AI initiatives but does not make decisions.
Symptom: meetings produce minutes rather than outcomes.
- **The Policy Illusion**
Governance is defined in documents but not embedded in execution.
Symptom: teams proceed without reference to governance artefacts.
- **The Diffused Accountability Problem**
Multiple stakeholders are involved, but no single point of accountability is defined.
Symptom: escalation leads to discussion rather than ownership.
- **The Escalation Gap**
Decision thresholds are undefined or inconsistently applied.
Symptom: similar incidents result in different responses.
- **The Monitoring Blind Spot**
Systems are approved but not continuously governed in production.
Symptom: issues are detected after impact rather than before.

What Leaders Should Do Now

Design governance as an operating model, not a policy. Review your current AI governance architecture and ask whether it functions as a connected operating system or as a set of documents, committees and principles. If governance is not embedded in how decisions are made, how accountability is assigned and how risk is controlled in practice, it is not yet an operating model. Design it as one.



Assign explicit accountability for AI outcomes. Identify every material AI system currently in production and confirm that a named individual is accountable for its outcomes — not its process, and not its technology performance. Accountability that is shared across functions is accountability that belongs to no one. Make it explicit, document it and test it.



Define decision rights across the governance stack. Establish who approves AI use cases, who accepts associated risk, who can intervene when thresholds are breached and who is accountable for outcomes at each level of the organisation. Decision rights that are undefined or informally understood will not hold under pressure.



Establish governance forums that make decisions. Review your existing AI governance forums and assess whether they are decision-making bodies or information-sharing forums. A forum that produces minutes rather than outcomes is not governing. Each forum should have a defined scope, clear authority and escalation pathways to the level above.



Extend governance across the full lifecycle. Governance that ends at the point of deployment is incomplete. For each material AI system, confirm that monitoring, intervention and retirement criteria exist. Systems approved and then left unattended are not governed. They are abandoned.



Connect governance to value and risk simultaneously. Ensure your governance model does not operate purely as a risk control function. It must also enable prioritisation, scaling and value realisation. A governance model that can only say no will be bypassed. One that connects value, risk and execution will earn the confidence of the business.



Closing Insight

AI governance is not something organisations add to AI.

It is how the organisation runs AI.

The organisations that scale AI are not those with better models or more data, but those that have made governance a deliberate part of how decisions are made, risks are controlled and accountability is enforced.



The difference is not policy. It is design.

References

Air Canada v. Moffatt, 2024 BCCRT 149. Civil Resolution Tribunal of British Columbia, February 2024.

EY. Responsible AI Pulse Survey 2025. EY Global, June 2025.

Harvard Law School Forum on Corporate Governance. Cyber and AI Oversight Disclosures: What Companies Shared in 2025. October 2025.

McKinsey & Company. State of AI Trust in 2026: Shifting to the Agentic Era. Based on the McKinsey 2026 AI Trust Maturity Survey. March 2026.

McKinsey & Company. The State of AI 2025: Adoption, Investment and Value Realisation. McKinsey & Company, 2025.

National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST, 2023.

National Institute of Standards and Technology. Generative Artificial Intelligence Profile (AI 600-1). NIST, 2024.

International Organization for Standardization. ISO/IEC 42001 Artificial Intelligence Management System. ISO, 2023.

Organisation for Economic Co-operation and Development. OECD AI Principles. Updated 2024.

European Union. Artificial Intelligence Act (EU) 2024/1689. 2024.

Deloitte. State of Generative AI in the Enterprise. Deloitte AI Institute, 2025.

KPMG. Deploying Trustworthy AI. KPMG, 2025.

Australian Institute of Company Directors and Human Technology Institute. A Director's Guide to AI Governance. 2024.

About the Author

Manoj Tavarajoo has spent over two decades working with boards, executives and senior leaders across enterprise, digital and AI transformation. His work sits at the intersection of strategy, operating model design, governance and execution, where transformation ambition becomes measurable enterprise capability.

He is the author of *Leading the AI Transformation*, *The AI Operating Model Playbook* and *The AI Governance Imperative*, together forming a practical leadership trilogy on how organisations lead AI, organise for AI and govern AI at enterprise scale.

Manoj works through MyConsultancy, an independent advisory and transformation practice based in Australia, and is open to senior embedded leadership opportunities partnering with boards and executive teams navigating AI strategy, operating model design and governance at enterprise scale.



@[manojtavarajoo](https://www.linkedin.com/company/manojtavarajoo)

About MyConsultancy

MyConsultancy works with boards and executive teams navigating the distance between AI ambition and operating reality. The practice specialises in AI strategy, operating model design, governance and enterprise transformation, helping organisations build the structures, decision-making disciplines and execution capability needed to scale AI responsibly.



www.myconsultancy.com.au

MyConsultancy