

EXECUTIVE BRIEF

The AI Accountability Illusion

Why responsibility is shared, but accountability is not

Manoj Tavarajoo

MyConsultancy

AI systems are distributed. Accountability is not.


Vendor reliance does not transfer accountability. It concentrates it.

AI SYSTEMS ARE INCREASINGLY

- ⚠ Externally sourced from vendors and platforms
- ⚠ Dependent on foundation models and APIs
- ⚠ Continuously updated by third parties
- ⚠ Operationally opaque to the organisation
- ⚠ Connected across complex dependency chains

ORGANISATIONS REMAIN ACCOUNTABLE FOR

- ✓ The authority granted to AI systems
- ✓ The decisions those systems execute
- ✓ The controls implemented around them
- ✓ The outcomes produced through them
- ✓ The governance visible to regulators



**The system you approved
is not the system running.**

Most organisations cannot fully observe the systems they depend on

MODERN AI SYSTEMS RELY ON

- ↘ Foundation models updated by their providers
- ↘ External APIs and orchestration layers
- ↘ Embedded AI services in SaaS platforms
- ↘ Continuously evolving model ecosystems
- ↘ Third-party training data and pipelines

YET MANY CANNOT DETERMINE

- ✗ How and when underlying models change
- ✗ When system behaviour has shifted materially
- ✗ Where decisions are shaped by hidden layers
- ✗ How downstream actions propagate
- ✗ What operational assumptions remain valid

Dependency without visibility creates governance exposure.

Shared responsibility does not mean shared accountability

VENDORS MAY

- ✓ Operate and maintain AI models and platforms
- ✓ Apply safety controls and access management
- ✓ Provide compliance and audit documentation
- ✓ Manage infrastructure and availability
- ✓ Support incident response and remediation



ORGANISATIONS REMAIN ACCOUNTABLE FOR


- ⚠ Operational outcomes for customers and stakeholders
- ⚠ Regulatory obligations regardless of vendor controls
- ⚠ Delegated authority granted to AI systems
- ⚠ Customer impact of AI-enabled decisions
- ⚠ Governance effectiveness demonstrated to regulators

Accountability concentrates at the point of decision, not the point of development.

Responsibility is distributed across the chain. Accountability is not.



The system can change without the organisation changing anything



Operational behaviour can change without operational approval.

CONTRACTS ARE NOT GOVERNANCE

CONTRACTS MAY

- ✓ Allocate liability between contracting parties
- ✓ Define obligations and service expectations
- ✓ Create rights of audit and data access
- ✓ Establish breach and remediation processes
- ✓ Provide legal recourse after harm occurs

CONTRACTS CANNOT

- ✗ Monitor AI behaviour in production environments
- ✗ Enforce escalation when thresholds are breached
- ✗ Guarantee operational visibility into model changes
- ✗ Prevent governance failure before harm occurs
- ✗ Provide evidence of control to regulators on demand



**Legal protection
does not create
operational control**

Contracts allocate responsibility.
They do not govern behaviour.

Accountability requires evidence, not assurance narratives

ORGANISATIONS TYPICALLY PROVIDE

- ↘ AI governance policies and frameworks
- ↘ Vendor attestations and compliance certificates
- ↘ Responsible AI principles and commitments
- ↘ Governance structures and committee minutes
- ↘ Point-in-time approval documentation

REGULATORS INCREASINGLY EXPECT

- ✓ Evidence that controls operate in production
- ✓ Demonstration of behavioural monitoring capability
- ✓ Escalation records and intervention logs
- ✓ Independent assurance of control effectiveness
- ✓ Accountability traceability for AI-enabled decisions

Governance confidence without operational evidence creates hidden exposure.

The accountability decisions cannot remain implicit

Map AI dependency chains



Identify every material AI system and its full dependency structure. Foundation model providers, integrators, APIs and embedded AI services all require visibility.

Define accountability boundaries




For every AI-enabled decision that affects customers, employees or financial outcomes, establish who is accountable for the outcome. Make it explicit and enforceable.

Establish evidence requirements



Determine what operational evidence of control the organisation must be able to produce. Define the standard before regulators define it for you.

Govern vendor dependencies actively



Require notification of material model changes. Define contractual rights to audit. Monitor behavioural drift. Do not rely on vendor assurance narratives alone.

Design intervention capability



Build the ability to pause, contain and remediate AI-enabled decisions when they move outside acceptable limits. Intervention must be possible before harm occurs.

AI accountability is no longer a legal question. It is an operational governance capability.

Five questions leaders should be uncomfortable answering

- | | |
|---|----|
| → Could we explain who is accountable for every material AI-enabled decision today? | 01 |
| → Do we understand where foundation model dependencies exist across our organisation? | 02 |
| → Would we detect behavioural changes introduced through vendor or model updates? | 03 |
| → Are our contracts creating operational control or only allocating legal responsibility? | 04 |
| → Could we demonstrate evidence of AI control to a regulator on demand? | 05 |

If any of these are difficult to answer, the accountability exposure is already present.

The defining governance challenge is no longer whether organisations use AI.

It is whether they can observe it, govern it, intervene in it, and demonstrate evidence of control.


The organisations that manage AI accountability successfully will not be those with the most vendor controls. They will be those with the clearest accountability boundaries and the strongest operational evidence of governance in practice.



**Vendor reliance does not transfer accountability.
It concentrates it.**

THE AI TRANSFORMATION FRAMEWORK

Lead AI. Organise AI. Govern AI. Execute with discipline.



MANOJ TAVARAJOO

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 leads MyConsultancy, an independent advisory and transformation practice based in Australia, partnering with boards and executive teams on enterprise, digital and AI transformation spanning strategy, operating model design, governance and execution. He also undertakes select senior embedded leadership roles where sustained executive ownership is required to translate strategy into measurable enterprise outcomes.

MyConsultancy

THE AI TRANSFORMATION TRILOGY

Three books, one system: how to lead AI, organise for it, and govern it once it acts.




LEAD AI
AI Strategy and Transformation Leadership



ORGANISE AI
AI Operating Model and Organisational Design




GOVERN AI
AI Governance, Accountability and Assurance



Leading the AI Transformation
Lead AI
How Bold Leaders Unlock Value, Reshape Strategy, and Build the Future

A practitioner-focused guide for executives navigating the shift to AI-enabled organisations. Covers leadership alignment, AI strategy, governance, responsible adoption, and building sustained enterprise AI capability.



The AI Operating Model Playbook
Organise AI
Why Structure, Not Algorithms, Determines AI Outcomes

A structural guide to scaling AI in the enterprise. Examines the operating model decisions that determine whether AI initiatives scale or stall, including portfolio governance, funding logic, decision rights, team design, and process integration.



The AI Governance Imperative
Govern AI
How Boards and Leaders Govern Value, Risk, Accountability and Autonomous AI

A governance guide for boards and leaders on AI behaviour, evidence, accountability, vendor risk and autonomous systems. At its centre is the Behavioural Envelope, converting approval into a continuous evidence obligation.



The AI Accountability Illusion
