Nemko Digital

# AI Governance Roles & Responsibilities:

## Why Is It So Hard?

## From governance principles to clear accountability and execution

AI governance is moving from policy to operations. Many organisations have established AI principles, policies, committees and initial governance frameworks. The harder challenge starts when these need to work in practice: when AI systems are being developed, procured and deployed across different business units, and somebody needs to decide what is acceptable, what controls are required and who remains accountable once the system is in use.

This is particularly difficult because AI does not fit neatly into existing organisational boundaries. Business teams, Data and AI, IT, Information Security, Privacy, Legal, Risk, Compliance, Procurement and Internal Audit can all have legitimate responsibilities. The result is a recurring paradox: **many people are involved in AI governance, while it can remain unclear who is actually accountable for making a decision or taking action**.

The hardest part of AI governance is often not defining the rules, but **translating them into clear accountability and decision rights**. This becomes harder, not easier, as AI adoption scales: responsibilities cross organisational boundaries, change throughout the AI lifecycle and increasingly need to move from central specialists into the business.

## Why roles and responsibilities become difficult so quickly

The first challenge is that too many functions can have a legitimate role. The business understands why an AI system is being used and what outcome it should deliver. Data and AI teams understand the technology. IT and Security manage technical environments and controls. Legal, Privacy, Compliance and Risk bring different regulatory and risk perspectives. Procurement manages third-party relationships, while Internal Audit may ultimately provide independent assurance. All these perspectives can be necessary, but more stakeholders do not automatically create better governance. Without clear accountability, shared involvement can easily become shared uncertainty.

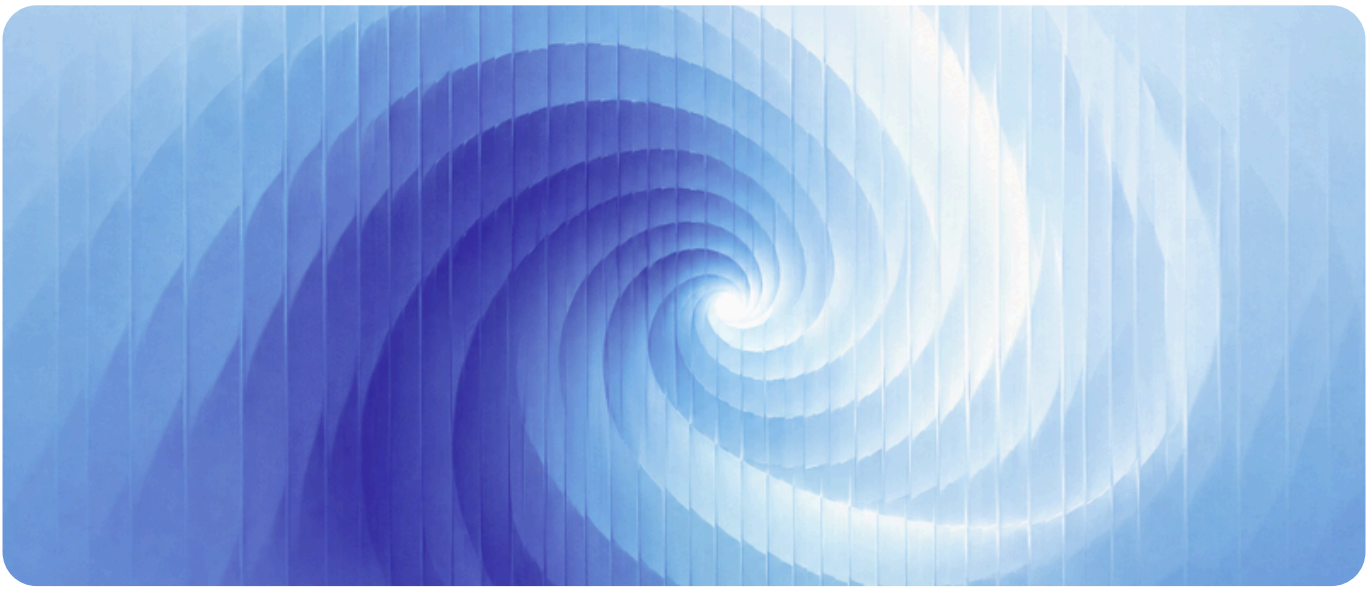
The second challenge is that **“ownership” is not one thing**. Business ownership concerns why an AI system is being used and the outcome it should deliver. System or product ownership concerns its day-to-day management. Technical ownership may sit with the team developing, configuring or maintaining it. Risk ownership concerns who accepts the remaining risk. Control owners perform and evidence specific controls, while Legal or Compliance may have responsibilities for interpreting applicable requirements. Simply stating that “the business owns the AI” therefore does not resolve the problem. Organisations need to be explicit about which responsibilities and decision rights come with each form of ownership.

Third, **accountability changes with the system and throughout its lifecycle**. An AI system moves from idea and intake through classification, development or procurement, assessment, approval, deployment, monitoring, change and eventually retirement. Different decisions require different expertise and authority. The person approving deployment may not be responsible for monitoring model performance, and the team performing a control may not have authority to accept the associated risk. Governance therefore needs to define responsibilities throughout the lifecycle rather than concentrate them around a single approval gate.

Third-party AI adds another layer to this challenge. Organisations increasingly use AI they did not build, including SaaS applications, copilots, foundation models and AI embedded in third-party products.

An organisation may have limited visibility into or control over the underlying model, but still control how it is selected, configured and used and how its outputs affect business decisions. **Outsourcing the technology does not remove the need for internal accountability**. In practice, it can make clear ownership, supplier governance, monitoring and escalation even more important.

Finally, roles that work for ten AI systems may fail at one hundred. Early AI governance is often highly centralised. Specialists can review individual systems and a central committee can discuss important use cases. As adoption increases, the same model can become a bottleneck. The roles-and-responsibilities question then changes from “Who could be involved?” to “Who actually needs to be involved in this particular decision?”



# Move beyond RACI: make accountability and decision rights explicit

Many organisations start the roles-and-responsibilities discussion by creating a RACI. This can be useful, but it can also create a false sense of clarity. A function may be marked as "accountable" or "consulted" without resolving the harder questions: **accountable for what outcome, authorised to make which decision, and responsible for what happens after that decision?**

Building on established governance concepts, including decision-rights models and the Three Lines Model, we distinguish five practical accountabilities for operational AI governance: **Own, Do, Decide, Challenge and Assure**. Together, they clarify not only who participates in governance, but how responsibility and authority are distributed.




*Figure 1. The five operational AI accountabilities—Own, Do, Decide, Challenge and Assure—assigned according to risk.*

**Own — accountability for purpose and outcome**. The owner should be accountable for why the AI system is being used, whether that use remains appropriate and whether it delivers the intended business outcome. This accountability should continue after deployment. An AI system without a clear owner can easily become everybody's technical responsibility but nobody's management responsibility.

**Do — responsibility for execution.** Different teams may develop, procure, configure, operate or monitor an AI system and perform its required controls. These responsibilities need to be sufficiently specific to translate governance requirements into day-to-day activities. A policy that says a system "must be monitored", for example, is incomplete unless somebody knows what needs to be monitored, when and by whom.

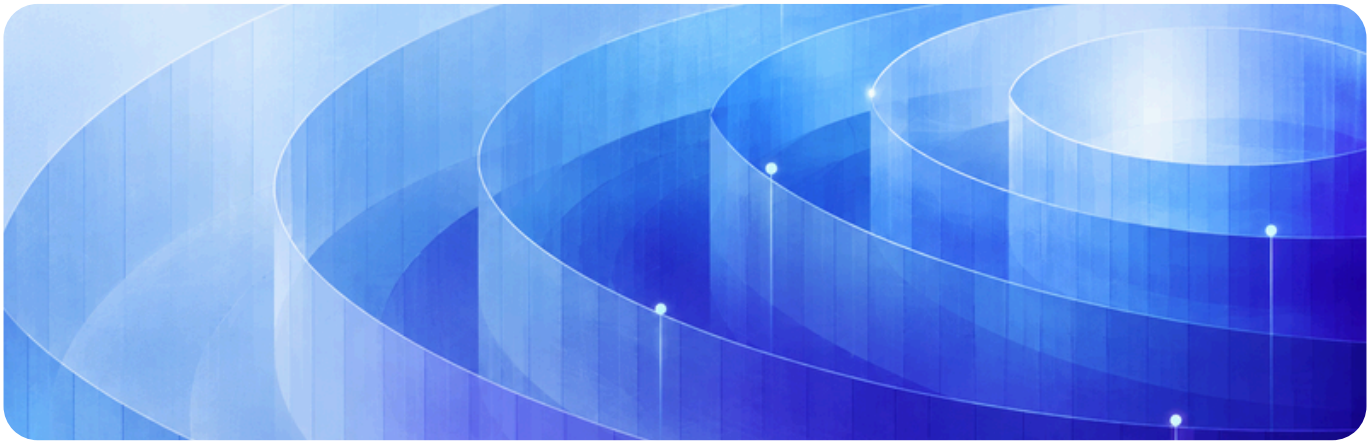
**Decide — authority to make consequential decisions**. This is often where governance models are least explicit. Who can approve deployment? Who can accept residual risk? Who can grant an exception, require remediation or suspend a system? Being involved in an assessment is not the same as having decision authority. For material decisions, organisations need to know where that authority sits and when escalation is required.

**Challenge — sufficiently independent scrutiny**. First-line ownership should not mean that teams assess and approve all of their own risks. Depending on the risk, relevant second-line or specialist functions should be able to challenge assumptions, assessments and controls. The purpose is not to add another approval layer to every use case, but to provide independent challenge where it adds value.

**Assure — evidence that the system works as intended**. For higher-risk or more material AI, management may need independent confidence that governance and controls are operating effectively. Existing third-line functions, independent testing or external assurance can play this role, depending on the nature and significance of the system.

These five accountabilities do not necessarily sit with five different people or functions, nor should every AI system require the same separation. **The appropriate allocation depends on risk**. What matters is that ownership, execution, decision authority, challenge and assurance are deliberately assigned rather than assumed.

This leads to an important design principle: **clarity of accountability should reduce unnecessary involvement, not increase it**. Once an organisation knows who owns the outcome, who executes the controls and who has authority to decide, it becomes much easier to determine when specialist challenge or independent assurance is actually required.



# Let risk determine who needs to be involved

The objective should not be to involve every potentially relevant function in every AI decision. It should be to make good decisions at the right level, with specialist involvement proportionate to risk. Done well, clearer accountability should **accelerate responsible AI adoption rather than slow it down**.

Risk-based governance is therefore not only about applying more controls to higher-risk AI. It is also a mechanism for allocating scarce governance and specialist capacity. A simple green, orange and red model can make this operational.

## The right governance for the right risk

### Business-led governance

· Defined guardrails
· Standard controls
· Clear ownership
· Escalation when needed

### Targeted specialist challenge

Business ownership + relevant Risk, Legal, Privacy or Security involvement

### Formal decisions and stronger assurance

· Documented approvals
· Risk acceptance
· Enhanced monitoring
· Independent challenge and assurance as appropriate

*Governance intensity increases with risk, without involving every function in every decision.*

*Figure 2. Risk-based governance increases governance intensity from business-led guardrails to formal decisions and stronger assurance.*

**For green, lower-risk AI**, accountability should sit predominantly in the first line. Teams operate within predefined guidance and controls, maintain appropriate registration and ownership, and monitor their use. Capabilities should increasingly be embedded within business and technology teams themselves. Central governance does not need to approve every decision; clear guardrails and escalation triggers allow teams to proceed within defined boundaries.

**For orange, moderate-risk AI**, first-line ownership remains, but additional challenge becomes appropriate. Depending on the actual risk, this might involve Privacy, Security, Legal, Risk, Compliance or another specialist function. Additional assessment, monitoring, evidence or approval may be required. Importantly, not every control function needs to review every orange system. **Specialist involvement should follow the risk**.

**For red, high-risk, regulated or materially impactful AI**, responsibilities become more formal. Stronger approval authorities, segregation of duties or four-eyes principles, documented risk acceptance, enhanced monitoring, independent challenge and appropriate testing or audit may be necessary. Incident and escalation responsibilities should also be explicit.

The objective is not less governance. It is more effective governance: **governance intensity should increase with risk; organisational complexity should not**.

# Roles must evolve as governance moves from start to scale

Getting the roles right for today's AI portfolio is only half the challenge. **The allocation of responsibility itself needs to evolve as the organisation matures**. This can be considered across three stages: start, scale and maintain.

### Start — Centralise enough to create clarity

During the start phase, some centralisation is useful and often necessary. A central AI Governance team or Centre of Excellence can establish the initial process, define roles and decision rights, develop risk classification, support early use cases, provide specialist expertise and introduce templates, workflows and supporting technology. Rather than perfecting the model in isolation, organisations can apply it to a representative set of real AI systems and learn where ownership, decisions and workflows break down.

Business accountability, however, should exist from the beginning. The central team can enable governance, but it should not become the owner of the organisation's AI.

**Centralise enough to establish consistency and capability, not to transfer accountability away from the business.**

### Scale — Decentralise by design

The scale phase is where many governance models face their real test. If every AI system continues to require the CoE, Legal, Risk and a central committee, increasing AI adoption will produce a corresponding increase in governance workload. Eventually, the process becomes a bottleneck or teams begin to work around it.

Responsibilities therefore need to move deliberately into the organisation. First-line business and technology teams should increasingly own AI systems, perform standard controls and manage risk within defined boundaries. Existing second-line functions should monitor and challenge, provide specialist oversight and support continuous improvement. Third-line functions should provide independent assurance. The central AI Governance function increasingly shifts towards standards, methodology, enablement, supporting technology, coordination and specialist expertise rather than individual operational decisions.

Standardised intake, risk-based routing, reusable controls and assessments, workflow automation and integration with existing procurement, development, security, privacy, risk and quality processes make this possible. Over time, AI governance should increasingly use the organisation's existing first-, second- and third-line structures rather than remain a parallel control organisation.

**A Centre of Excellence that remains responsible for every AI decision may have successfully established governance, but it has not successfully scaled it.**

### Maintain — Institutionalise accountability

Approval is not the end of AI governance. Models change, vendors update their products, use cases expand, data changes, regulation develops and organisational responsibilities evolve. Accountability therefore needs to persist after deployment.

This includes monitoring actual model and system performance, but also monitoring whether the governance workflows and controls themselves continue to operate effectively. As the AI portfolio grows, supporting technology and automation can help execute workflows, capture evidence and identify where intervention is required.

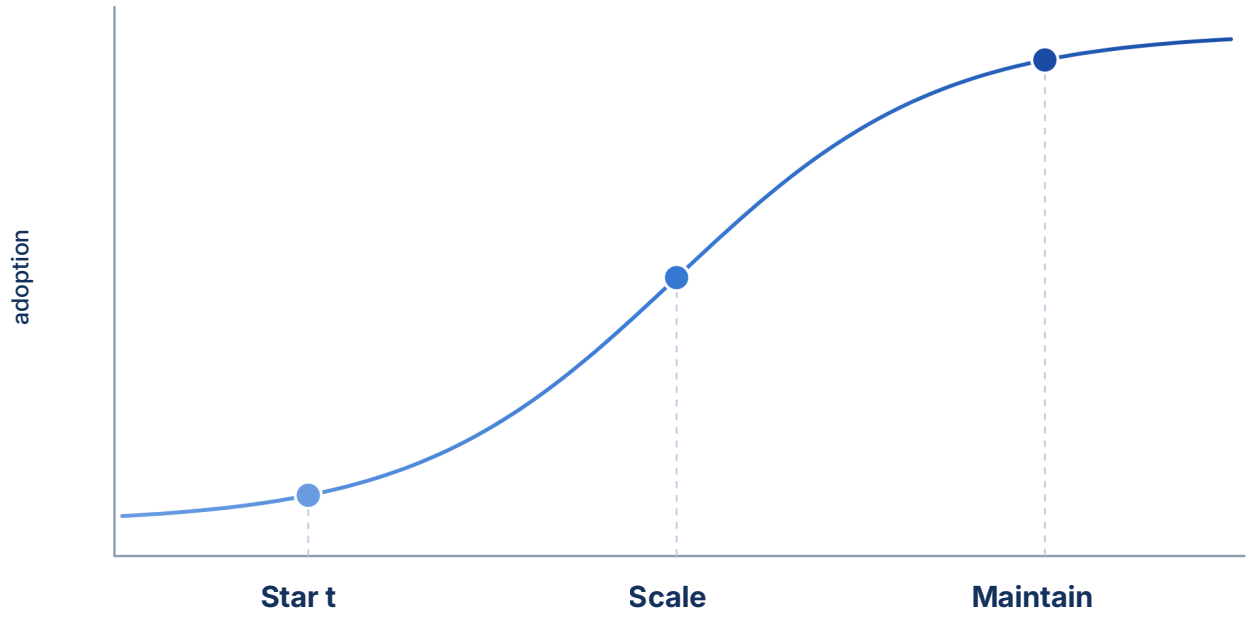
Clear accountability is equally important for incidents and change. Organisations need to know how issues are identified, assessed and escalated; who can intervene; and what changes trigger reassessment. For more material AI, incident arrangements may also need to be periodically exercised rather than merely documented.

Maintenance should also include continuous improvement and training, proportionate periodic or annual review, consideration of regulatory developments and targeted internal or external assurance where appropriate. And governance should not focus exclusively on downside risk. Organisations should continue to track whether AI systems are producing their intended outcomes and value. Risk, performance and value should ultimately be managed together.

AI governance is therefore not an approval gate that a system passes once. **It is an ongoing management capability**.

## From central governance to accountability at scale

Illustrative evolution as the number of AI systems increases




**Start**

Central AI Governance team establishes standards, roles and processes, and supports initial use cases.

**Scale**

Business and technology teams take on routine ownership and controls. The central team enables them through guidance, workflows and tools.

**Maintain**

The central team maintains the framework, monitors governance effectiveness and coordinates improvement and specialist oversight.

*Figure 3. Illustrative evolution of AI governance from centralised setup to distributed accountability as AI adoption grows.*

## Keep the complete operating model in view

Focusing on roles and responsibilities does not mean treating them in isolation. Effective accountability sits within a broader operating model connecting **governance and oversight, roles and decision rights, lifecycle processes, risk and controls, supporting technology and evidence, and monitoring and assurance**.

The practical challenge is to keep this complete system in view without trying to perfect every component before starting. Establish enough structure to govern real systems, learn from implementation, standardise what works and progressively embed it into existing organisational processes. Technology can then enable rather than substitute governance: automating workflows, maintaining evidence, supporting monitoring and helping management understand whether the governance process itself is operating effectively.

A useful management test is to select five AI systems already in operation and trace their accountability from business ownership through approval, monitoring and assurance. **If this requires several meetings and emails to reconstruct, the organisation may have an AI Governance framework, but accountability is probably not yet operating effectively at scale**.

## From governance function to organisational capability

The objective is not to build a larger AI Governance function. It is to make **AI accountability part of how the organisation operates**. Start with enough central structure to create clarity. Use risk to determine the level of involvement. Then deliberately move capability and accountability into existing first-, second- and third-line structures as AI scales.

This transition does not happen through governance design alone. It requires implementation: processes that people can actually follow, supporting technology and workflows, capabilities in the first and second lines, ongoing monitoring, and assurance that the system continues to work as AI adoption grows.

Ultimately, mature AI governance is not defined by the number of policies, committees or controls an organisation has. **It is defined by whether the right people can make the right decisions, at the right level, with the evidence to demonstrate that those decisions are being managed over time**.

# About the author



## Bas Overtoom

Bas Overtoom is the Global Business Development Director where he leads global efforts to promote responsible AI adoption at Nemko Digital, working with organizations to operationalize trust, transparency, and compliance in their AI systems. With a strong background in business-IT transformation and AI governance, he brings a pragmatic approach to building AI readiness across sectors.

Nemko Digital