Chartered Governance Institute UK & Ireland

Equipping Governance Professionals to Lead AI Conversations

July 2026

This report has been prepared by the Policy Team of the Chartered Governance Institute UK & Ireland. The principal author is Dr. Valentina Dotto, Policy Adviser.

Comments, questions and observations

If you have any feedback on the content of this report, or additional questions that you'd like to discuss, please contact the Chartered Governance Institute UK & Ireland: 020 7580 4741| policy@cgi.org.uk

www.cgi.org.uk

Published: July 2026

Date of next review: July 2029

Contents

Executive summary

Artificial intelligence is rapidly reshaping corporate governance, but adoption is outpacing oversight. Across sectors, organisations are embedding AI into operational processes while governance frameworks lag behind, creating a structural gap between how decisions are made and how accountability is exercised.

Evidence from cross-sector roundtables and survey data shows that AI adoption remains fragmented, operational, and largely bottom-up. Board visibility is often limited, and formal governance structures are not evolving at the same pace as deployment. As a result, AI risk is no longer purely technical, it has become a core governance, assurance, and accountability issue.

A central pressure point is fiduciary duty. Directors remain fully accountable for decisions even where outcomes rely on opaque, probabilistic systems. This is elevating AI literacy from a technical skill to a governance imperative. Governance professionals are correspondingly shifting roles, from administrators to institutional interpreters, validating outputs, challenging machine-generated reasoning, and ensuring that AI-assisted decisions remain explainable, evidence-based, and defensible.

At the operational level, AI is already transforming governance processes, from minute-taking to document review and workflow management. While efficiency gains are clear, these tools introduce risks relating to evidentiary integrity, discoverability, record-keeping, and legal exposure. Practices remain inconsistent, particularly below board level, where informal AI use is widespread and difficult to monitor.

Looking ahead, the emergence of agentic and semi-autonomous systems will further test governance models. As systems begin to initiate actions and operate with reduced human input, organisations will need clearer escalation protocols, tighter control environments, and explicit accountability structures embedded within governance frameworks.

At the same time, a notable blind spot is emerging in sustainability oversight. Despite positioning AI as an efficiency driver, few organisations are integrating its environmental impact, including energy and infrastructure demands, into ESG governance and reporting frameworks.

Importantly, AI is not reducing the need for governance; it is redefining it. As routine tasks become automated, the value of governance professionals increasingly lies in oversight, judgement, assurance, and the preservation of accountable decision-making systems.However, a significant capability gap remains. While awareness of AI is high at board level, practical competence is limited. Confidence, not knowledge, is the primary constraint on effective oversight. This is driving the integration of AI literacy into director recruitment, induction, and continuing professional development, with AI (alongside cyber) set to become as fundamental as financial literacy.

Overall, the key challenge is not whether AI will become embedded in governance systems, but whether governance frameworks can evolve quickly enough to preserve accountability, trust, and defensible human judgement.

Critical Questions for Boards

Boards should be able to answer these clearly, confidently, and on the record:

Governance Framework & Compliance

- Is your AI architecture aligned with ISO/IEC 42001:2023 (AI management systems)?
- Does your AI risk management framework meet the expectations of Provision 29 of the UK Corporate Governance Code?
- If not, can you clearly explain and justify any non-compliance?

Directors' Duties & Accountability

- Has the board explicitly considered how AI use aligns with directors' duties under s174 (care, skill, and diligence)?
- Can you evidence that AI-assisted decisions are subject to appropriate human oversight and challenge?

Oversight & Control

- Do you have visibility of where and how AI is actually being used across the organisation — including below board level?
- Are escalation thresholds, decision rights, and accountability clear for AI-assisted or AI-initiated actions?

Assurance & Evidence

- Can you demonstrate the integrity, traceability, and auditability of AI-supported decisions?
- Are governance records (minutes, papers, outputs) resilient to legal scrutiny and disclosure risks?

Capability & Readiness

- Do directors have sufficient practical competence to interrogate AI outputs and challenge assumptions?
- Is AI literacy embedded in recruitment, induction, and ongoing board development?

Strategic & ESG Oversight

- Have you assessed the environmental and infrastructure implications of your AI strategy?
- Is AI appropriately integrated into your broader risk, strategy, and sustainability governance?

Bottom line: If boards cannot answer these questions with clarity and evidence, AI adoption is likely outpacing governance, and accountability risk is already accumulating.

Introduction

The study focuses on governance transformation. It examines how AI reshapes five core dimensions of governance:

- fiduciary accountability and statutory duties
- delegation structures and decision rights
- internal control and assurance systems
- evidentiary integrity and organisational records
- board oversight capability and governance competence

The central issue is not whether AI can automate governance activity, but how organisations preserve accountability, legal defensibility, and institutional resilience in increasingly AI-enabled operating environments.

The research is structured around seven governance questions, derived from the Institute's initial hypothesis framework and tested through roundtables, survey evidence, and task-level analysis:

1. **Visibility Problem**
   How do boards govern AI systems without full visibility, traceability, or auditability?
2. **Liability Question**
   How is "reasonable care, skill and diligence" redefined where AI informs decision-making?
3. **Evidentiary Pressure Point**
   How does AI affect the creation, retention, and discoverability of governance records?
4. **Autonomy Challenge**
   How should delegation frameworks adapt to increasingly autonomous and agentic systems?
5. **ESG Paradox**
   How can boards align AI adoption with sustainability obligations and reporting coherence?
6. **Governance Value Shift**
   How does AI reshape governance roles as administrative work is automated and oversight becomes interpretive?
7. **Capability Gap**
   How do organisations move from AI awareness to practical competence in oversight, challenge, and assurance?

Together, these questions define the central governance tension examined in this report: the widening gap between automation capability and non-delegable fiduciary accountability.

The analysis integrates two components.

First, the Institute conducted research across the corporate, public, charity, and education sectors between February and May 2026. The study combined 16 cross-sector roundtables involving 61 participants with survey responses from approximately 100 governance professionals. Participants were drawn from across sectors, including 38 from corporate organisations, 16 from the public sector, 11 from the charity and third sector, and 8 from education. Responses were also distributed across career stages, comprising 5 early-career, 37 mid-career, and 31 late-career governance professionals, ensuring that insights reflect a predominantly experienced practitioner base. The questions focused on operational adoption, fiduciary oversight, governance capability, and emerging systemic risks.

Second, the study applies a task governance framework based on the Chartered Governance Institute Company Secretary functional structure (ICSA Company Secretary's Handbook, 12th Edition).

This framework maps governance activity across seven operational domains:

- board and committee management
- statutory compliance and entity governance
- director and officer lifecycle management
- corporate transactions and disclosure processes
- governance frameworks and policy control
- regulatory intelligence and advisory activity
- departmental operations and enterprise data governance

Each task is classified into one of three exposure categories:

- automation: AI can substantially substitute execution
- augmentation: AI supports execution with human validation
- fiduciary judgement: responsibility remains non-delegable

Third, the task framework is calibrated using an International Labour Organization (ILO)-informed occupational exposure methodology. This converts task-level exposure into system-wide implications for workforce design, role evolution, capability requirements, and organisational restructuring.

Governance Risk and System Transformation

How AI is reshaping the architecture of organisational decision-making, accountability, and fiduciary exposure.

Fiduciary Duty Under Algorithmic Reliance

(When "reasonable care" meets machine-mediated judgement)

The roundtable evidence, read alongside emerging scholarship on generative and agentic AI, indicates a structural transition in corporate governance from linear delegation hierarchies to distributed, machine-mediated decision environments. Governance is no longer organised solely around identifiable human actors exercising delegated authority through formal reporting lines. Instead, authority is increasingly embedded within technical infrastructures where AI systems initiate, execute, and coordinate operational tasks across compliance, finance, procurement and governance functions.

Across corporate, charity and public sector contexts, participants consistently described AI systems as shifting from passive tools to active operational agents. These systems no longer merely draft or summarise information; they increasingly monitor regulatory change, trigger workflows, populate governance outputs and coordinate multi-step processes through integrated enterprise platforms.

For example, survey findings indicate that verification practices for virtual AGMs remain limited, with only 22 organisations confirming that they currently verify the authenticity of participants, compared to 42 that do not and 9 that are planning to implement such measures. At the same time, some organisations are beginning to explore the use of AI for identity verification: 7 respondents indicated that AI will be used to manage mandatory director verification, while 28 are considering its adoption and 38 do not currently plan to use it.

Does your organisation verify the authenticity of participants in virtual AGMs?

Yes 22
No 42
Planning to implement 9

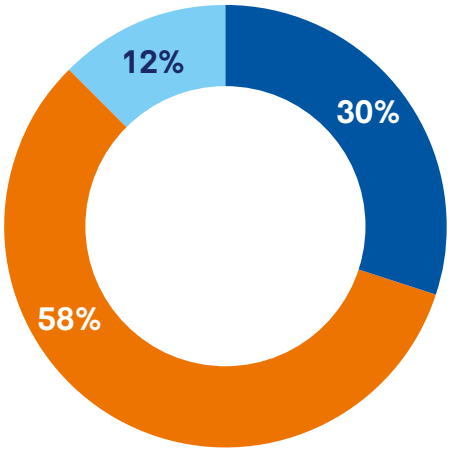

Will AI be used to manage mandatory identity verification for directors?

Yes 7
No 38
Maybe/Under Consideration 28

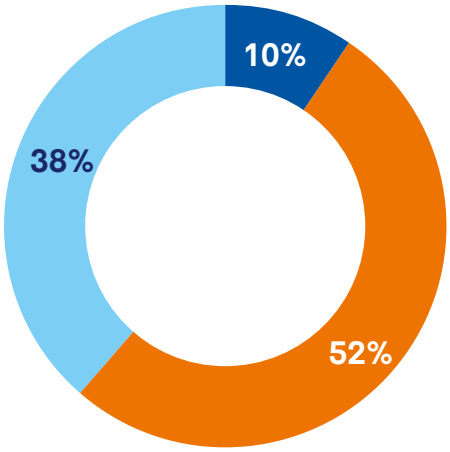

This produces an "action void", in which operational decisions are generated and executed within machine-to-machine environments before human oversight is engaged. While legal doctrine continues to treat AI as an instrument without agency, organisational practice increasingly reflects a form of functional autonomy in which AI systems shape outcomes without formal delegation.

Survey evidence reinforces this structural ambiguity. Respondents show no consistent allocation of oversight responsibility, with responsibility variously located at board, committee, or management level. At the same time, when asked who would be legally accountable in the event of an AI-related error, responsibility remains overwhelmingly anchored in the board rather than the system or vendor.

Taken together, these findings highlight a widening disconnect between where decisions are operationally generated and where accountability is legally assigned, intensifying the governance risks associated with machine-mediated decision environments.

The central governance problem is therefore no longer whether AI supports decision-making, but how decision rights, escalation pathways and fiduciary accountability operate within these distributed infrastructures.

Who is responsible for AI oversight?

A specific board committee (e.g., Audit or Risk) 18
The full board 17
Management only 38

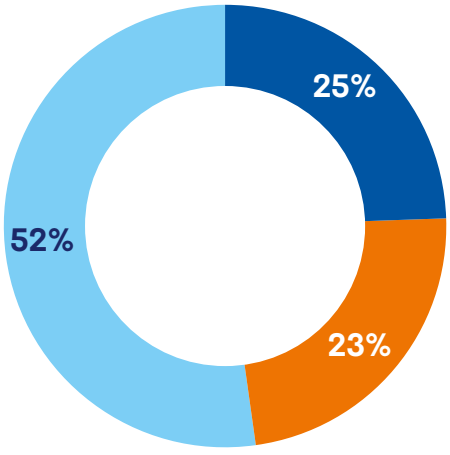

In the event of an AI-related disclosure error, who is legally accountable?

CEO 13
Board 29
Management 28
Vendor 9

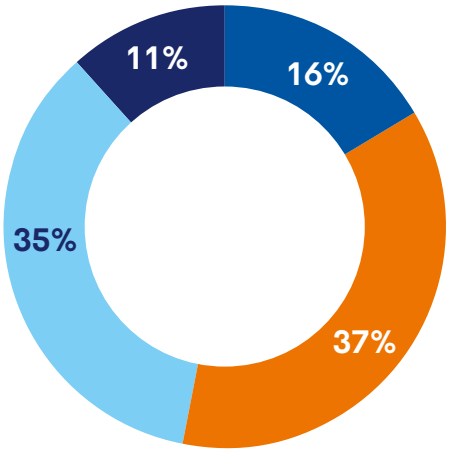

Q1 How are "reasonable care, skill and diligence" evolving under AI-assisted governance?

Areas of Consensus

The roundtable evidence converges on the view that Section 174 Companies Act 2006 is being operationally reinterpreted in practice rather than needing legal amendment.

Participants repeatedly emphasised that AI outputs may "sound great from top to bottom" while still containing hallucinations or fabricated content. This generates a shared concern that directors cannot rely on AI outputs in a passive manner. Instead, they must actively interrogate them. As one participant noted, governance failure increasingly arises where directors are unable to "check for accuracy" or understand when outputs require validation.

Across sectors, there is agreement that AI literacy is becoming integral to the exercise of reasonable care. This includes the ability to recognise the limitations of probabilistic systems, identify when outputs require verification, and challenge machine-generated reasoning rather than accepting it at face value.

The definition of 'reasonable care' is therefore shifting from a procedural standard to a capability-based standard. Directors are increasingly expected to operate within an epistemically complex environment where governance-relevant information is mediated by opaque systems.

A second area of convergence concerns the shift in fiduciary focus. Participants consistently described governance moving away from reviewing individual outputs towards supervising the systems that generate them. AI systems are no longer perceived as discrete tools but as process-generating infrastructures that shape decision pathways.

Survey evidence supports this transition towards system-level oversight. A majority of respondents report operating a human-in-the-loop protocol to audit AI-generated content, indicating that organisations are not treating outputs as self-sufficient but are embedding validation checkpoints within governance processes.

Do you use a human-in-the-loop protocol to audit AI-generated content?

Yes 46
No 27

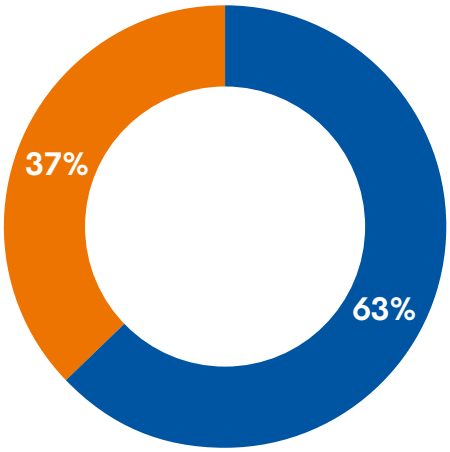

However, this also reinforces a deeper structural shift: human oversight is increasingly positioned as an **ex post control mechanism**, applied after system-generated outputs are produced, rather than as an **ex ante constraint on decision-making itself**. This reorientation further underscores the movement from output verification to system supervision as the central object of governance.

This creates a structural tension between legal doctrine, which assumes delegation through identifiable human actors, and operational reality, where decision production is increasingly system-mediated without formal legal attribution.

A third point of consensus concerns the emergence of a dual-risk environment. Participants recognised that reasonable care now requires balancing under-reliance and over-reliance on AI systems. Failure to use AI may be negligent in complex environments, while uncritical reliance on AI outputs may itself constitute breach of duty. 'Reasonable care' therefore operates as a dynamic balancing exercise rather than a fixed threshold.

Areas of Disagreement

A key divergence concerns whether directors must understand AI systems at a technical level. One position holds that substantive understanding is necessary to interrogate outputs meaningfully. The alternative view argues that governance assurance frameworks and expert intermediaries are sufficient, provided robust oversight structures exist.

While no consensus emerges, both positions reject passive reliance as defensible.

Governance Tensions

The dominant tension is epistemic, but not entirely novel. Directors have long remained legally accountable for systems they do not fully understand, for example, complex financial or accounting infrastructures overseen by trusted specialists. However, AI intensifies this familiar condition. Its ubiquity across organisational functions, combined with the reality that even experts rarely understand it fully, produces a qualitatively broader gap between fixed legal responsibility and increasingly distributed, opaque operational intelligibility.

Q2 What are the implications of reliance on opaque AI systems?

Areas of Consensus

Participants consistently identified opacity as a defining characteristic of AI-mediated governance systems. Outputs may appear correct while embedding errors that are not immediately visible. This creates a governance environment in which reasoning pathways are not easy to reconstruct; they can seem to emerge fully formed from a 'black box'.

This opacity is structurally linked to machine-mediated decision infrastructures, where outputs are produced through layered pipelines involving models, prompts and APIs rather than transparent linear reasoning.

A second area of consensus is that autonomous systems compress decision cycles and bypass traditional governance layers. Participants described systems that trigger workflows automatically, escalate or suppress information internally, and execute multi-step processes without direct prompting. This produces an 'action void', where operational decisions are executed before human oversight intervenes. In governance terms, this does not displace but strains existing oversight assumptions embedded in doctrines such as Caremark and statutory expectations under regimes including the ECCTA, which rely on identifiable allocation of responsibility and effective visibility over material risks. Where AI systems compress decision-making beyond meaningful human oversight, these assumptions become harder to sustain in practice.

A third area of convergence is the shift in governance focus from output review to infrastructure supervision. Participants emphasised the increasing importance of monitoring escalation pathways, system logs, model drift, confidence thresholds and intervention triggers. Governance is therefore shifting from document-based assurance to system-based oversight.

Areas of Disagreement

A key divergence concerns organisational trust in embedded enterprise AI systems. Corporate participants were generally more willing to integrate AI into operational workflows, including transcription and Copilot-based tools. In contrast, charity and public sector participants demonstrated greater caution and a stronger preference for manual verification.

A second disagreement concerns the acceptability of partial autonomy. Some organisations permit threshold-based automation, where AI systems execute tasks autonomously under defined confidence levels. Others restrict AI to assistive drafting only, maintaining strict human approval requirements.

Governance Tensions

Three tensions dominate:

- transparency versus operational efficiency, as increased speed reduces interpretability
- automation versus fiduciary control, where systems generate actions without equivalent legal delegation
- embedded integration versus controllability, as platform-based AI reduces visibility over data flows and system behaviour.

Q3 What discoverability risks arise from AI-generated governance records?

Areas of Consensus

Participants consistently agreed that governance records are expanding beyond formal minutes to include AI-generated artefacts such as transcripts, prompts, system logs and metadata. These materials are increasingly treated as evidential chains rather than informal by-products.

Survey evidence reinforces this expansion: most respondents report that shadow AI usage occurs at least occasionally, indicating that such artefacts are being generated outside formal governance channels. This blurs the boundary between official records and peripheral digital traces, widening the scope, and complexity, of organisational record-keeping.

To what extent is shadow AI (unauthorised use of public tools) occurring?

Not at all 24
Occasionally 32
Frequently 17

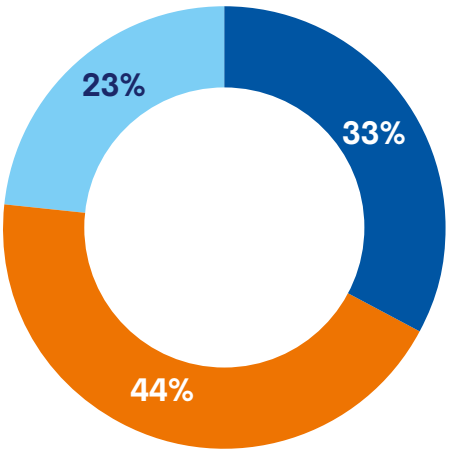

AI transcription systems were widely recognised as capturing complete meeting content. However, participants also noted that downstream summaries often fail to preserve context and nuance. This creates a dual-record environment comprising curated governance minutes alongside machine-generated evidential records.

A further area of consensus is that organisations are actively designing governance processes around litigation exposure. Common approaches include 72-hour deletion windows with certification, immediate post-approval deletion workflows, and external service provider models designed to ensure that no internal records persist.

This reflects a preventative governance logic in which non-retention is used as a deliberate risk mitigation strategy.

Areas of Disagreement

A key disagreement concerns whether deletion eliminates discoverability risk. Some organisations treat deletion as sufficient mitigation, while others recognise the persistence of shadow copies, external system retention, and API-level logging.

There is also no consensus on whether AI transcripts constitute working material or formal governance records. This classification uncertainty directly affects retention obligations.

Governance Tensions

Two tensions dominate. First, evidential completeness increases accuracy and auditability but simultaneously increases litigation exposure. Second, transparency may conflict with legal privilege, particularly where AI systems process sensitive or legally privileged material through external infrastructure.

Q4 What governance obligations arise regarding retention and deletion?

Areas of Consensus

Retention is increasingly treated as a governance control function rather than an operational IT issue, with participants describing growing involvement of legal and compliance functions in decisions on AI deployment, particularly where governance artefacts are generated, processed, or stored.

Survey evidence, however, highlights a clear implementation gap. In relation to AI-generated meeting data, only 22 organisations report having a formal retention and deletion policy in place, compared to 15 relying on informal practices and 36 operating with no policy at all. A similar pattern is evident in third-party assurance, where just 20 respondents have conducted an AI risk assessment, while 53 have not. Board-level control mechanisms remain particularly underdeveloped: only 5 organisations report having a formal AI kill switch policy, compared to 68 that do not.

Taken together, these findings indicate that while AI-related data risks are increasingly recognised, they remain insufficiently codified within formal governance and assurance frameworks.

Data minimisation has therefore emerged as a primary risk mitigation strategy. Organisations are restricting AI access to trained governance professionals, reflecting the need for interpretive expertise in evaluating AI-generated outputs. In parallel, participants reported increased reliance on controlled usage environments and stricter deletion practices, even where formal policies are not yet established.

Participants also noted that AI governance is becoming increasingly jurisdiction-sensitive. Deployment decisions often require legal approval, and regulatory expectations vary materially across regions and sectors.

How does your organisation manage retention and deletion of raw AI-generated meeting data?

Formal policy in place 22
Informal practice 15
No policy 36

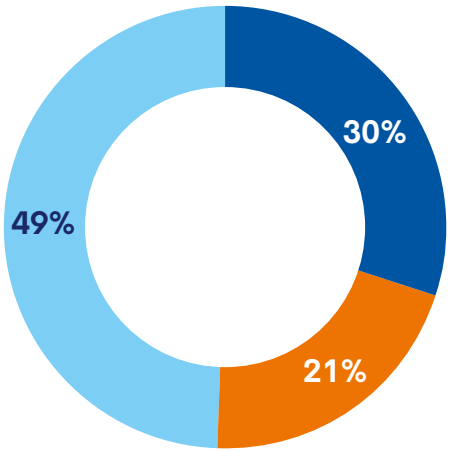

Has your organisation conducted a third-party AI risk assessment?

Yes 20
No 53

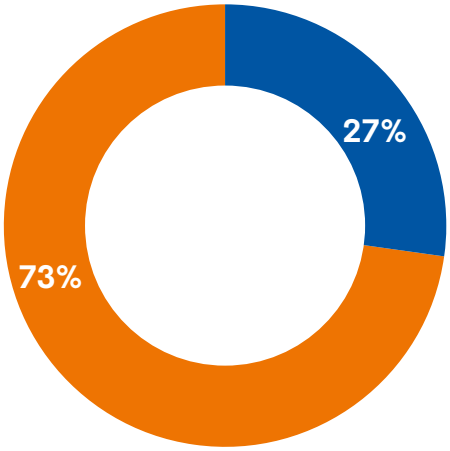

Does your board have a formal AI kill switch policy?

Yes 5
No 68

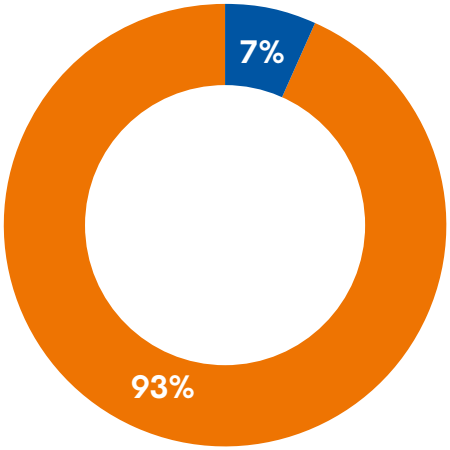

Areas of Disagreement

There is divergence over classification of AI outputs. Some organisations treat outputs as transient operational artefacts, while others classify them as formal governance records subject to audit retention requirements.

A second disagreement concerns automation in retention systems. Some organisations permit automated deletion and retention cycles, while others require manual oversight for all governance-related data handling decisions.

Governance Tensions

Retention creates a structural tension between compliance and evidential risk. Stronger retention improves auditability but increases exposure in litigation. Conversely, broader accessibility improves efficiency but weakens control and oversight.

Integrated Structural Insight: The Emergence of Agentic Governance

Across all questions, the roundtable evidence indicates a consistent structural transformation in governance architecture.

Governance is shifting from linear delegation hierarchies to distributed machine-mediated execution systems. AI systems increasingly initiate actions, coordinate workflows, and escalate or suppress information. This produces what can be conceptualised as synthetic agency: functional decision-making without legal personhood.

This creates a persistent structural mismatch between de facto operational autonomy and de jure legal accountability. AI systems exercise functional authority in practice, but legal systems continue to assign responsibility exclusively to human actors and corporate entities.

As a result, governance shifts from supervising discrete decisions towards supervising decision infrastructures. The object of governance becomes system integrity, including escalation design, auditability, traceability and confidence calibration.

A further consequence is responsibility diffusion. Accountability is distributed across developers, vendors, governance professionals and operational users, yet legal liability remains concentrated in corporations and directors. This produces a structural accountability gap in which responsibility is fragmented in practice but not in law.

Conclusion

The roundtable evidence, together with the survey questions, demonstrates that generative and agentic AI do not alter fiduciary duty in formal terms but fundamentally transform the environment in which it is exercised.

Reasonable care, skill and diligence is becoming technologically mediated, system-dependent and evidentially continuous. Governance is shifting from linear delegation and document-based oversight towards distributed, machine-mediated control environments.

Within this context, fiduciary responsibility remains non-delegable. However, its meaningful exercise increasingly depends on whether organisations can design governance architectures that ensure AI systems remain visible, controllable, auditable and subordinate to human legal authority.

AI Value, Capital Allocation and Systemic Risk Exposure
(When investment logic meets planetary and budget constraints)

This strand examines AI not as a discrete technology deployment issue, but as a capital allocation decision embedded within wider governance, regulatory, environmental, and technological constraints. Roundtable participants consistently framed AI investment as a trade-off between efficiency gains, governance overheads, infrastructure dependency, and long-term strategic resilience rather than as a straightforward productivity enhancement.

The evidence indicates a widening gap between anticipated transformation and realised organisational value. While low-risk automation delivers measurable efficiency benefits, scaling AI into governance-critical functions introduces significant verification, compliance, auditability, and oversight costs. These governance burdens materially reshape the economics of AI adoption and weaken assumptions of linear productivity return. At the same time, organisations face increasing scrutiny over AI-related disclosures, sustainability commitments, and long-horizon technological risk exposure. This positions AI governance within a broader framework of capital discipline and evidential accountability, particularly given the high failure rates of AI initiatives and their significant opportunity costs. Investments in AI must therefore be justified not only in terms of operational efficiency, but also against competing uses of capital and growing ESG sensitivities, including environmental impact and workforce implications.

Q1 What evidence does management provide to substantiate AI-related claims in public disclosures and annual reports?

Areas of Consensus

Participants consistently reported a widening gap between internal AI usage and external disclosure practice. Across sectors, AI is widely embedded in operational workflows, particularly drafting, summarisation and analysis, yet this deployment is rarely matched by structured disclosure in annual reports or ESG statements.

AI is often framed externally as a strategic capability, while internally it remains fragmented, use-case specific and experimental. Correspondingly, both evidential and evidentiary standards for AI-related claims remain underdeveloped: organisations frequently rely on narrative statements of 'ongoing experimentation' rather than verifiable operational data such as deployment logs, performance metrics, or documented governance controls.

Survey evidence reinforces this disconnect. While organisations make outward claims about AI capability, accountability for disclosure errors

remains firmly located between the board and management rather than within the systems themselves. This highlights a persistent asymmetry between narrative signalling and evidential governance, where claims are increasingly decoupled from auditable substantiation, but liability remains unchanged.

Across sectors, there is also recognition that regulators and investors are beginning to expect greater transparency. Participants referenced increasing pressure from reporting frameworks and regulatory guidance requiring clearer articulation of AI usage, particularly where it influences risk management or strategic performance narratives.

What evidence does management provide to substantiate AI claims in public disclosures?

Independent audit 9
Internal validation only 29
None 35

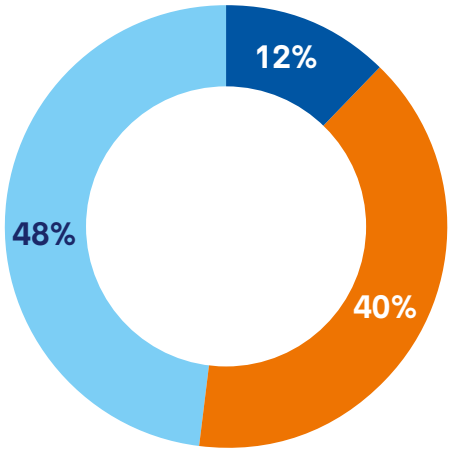

Areas of Disagreement

A clear divergence emerges on whether current disclosure practices already meet adequate governance standards. Some corporate participants argue that narrative disclosure is sufficient at this stage of technological maturity, given the rapid evolution of AI capabilities.

Others take the opposite view, stating that absence of structured evidence creates a material governance weakness. This group emphasises that without documented deployment evidence, AI claims risk becoming aspirational rather than auditable.

A second divergence concerns whether disclosure should be standardised. Some participants favour flexible, narrative-driven reporting aligned to existing annual report structures. Others advocate for formalised AI disclosure frameworks aligned to regulatory expectations and auditability standards.

Governance Tensions

The central tension lies between strategic narrative construction and evidential governance discipline. AI is increasingly used as a signalling mechanism in corporate reporting, yet governance systems for validating those claims remain immature.

This creates a structural disclosure risk where organisations may be perceived as "AI-enabled" without demonstrable operational substantiation. The result is an emerging gap between reported capability and verifiable deployment reality, increasing exposure to regulatory scrutiny and investor scepticism.

Q2 What percentage of AI pilots have scaled to enterprise-wide implementation, and who is accountable for failed capital allocation?

Areas of Consensus

Participants consistently reported low conversion rates from AI pilots to enterprise-scale deployment, with most initiatives remaining confined to bounded use cases such as drafting, transcription, or departmental experimentation. Survey data reinforces this pattern: 45 respondents indicated that only 0–25% of pilots successfully scale, compared to 14 reporting 26–50%, 5 reporting 51–75%, and just 9 reporting scaling rates above 75%.

There is broad agreement that these limitations are not primarily technological but organisational. Common barriers include data fragmentation, governance constraints, cybersecurity requirements, and challenges associated with legacy system integration. At the same time, accountability for these outcomes remains diffuse. While capital allocation decisions are often taken at board or executive level, responsibility for failed AI investments is most frequently attributed to management teams (47 responses), with fewer assigning accountability to boards (18) or CEOs (8).

This ambiguity is compounded by uneven pre-investment discipline. Only 26 respondents report conducting a formal cost-benefit analysis before initiating AI projects, compared to 22 who do not, 17 where this is still in progress, and 11 for whom the question was not applicable.

As a result, AI value realisation remains inconsistent. Lower-risk automation continues to deliver incremental efficiency gains, but more ambitious governance and decision-support applications frequently stall under the weight of control, risk, and validation requirements, reinforcing a persistent gap between experimentation and scalable organisational capability.

What percentage of AI pilots have scaled to enterprise-wide implementation?

0–25%	45
26–50%	14
51–75%	5
76–100%	9

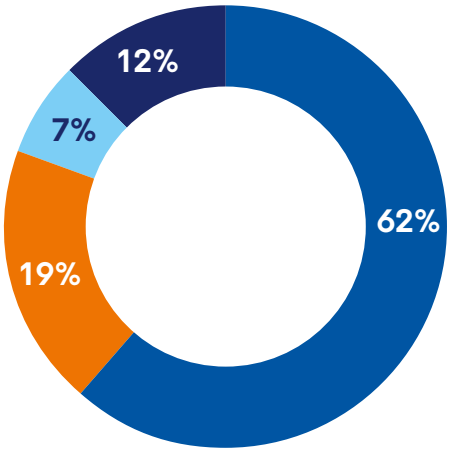

Has your organisation conducted a cost-benefit analysis before initiating any AI projects to determine whether AI is appropriate for your business?

Yes 26
No 22
In progress 17
Not Applicable 11

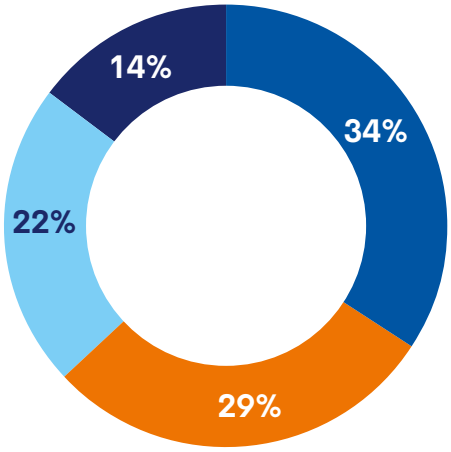

Areas of Disagreement

Some corporate participants argue that scaling failure is an expected feature of innovation cycles and should not be interpreted as governance weakness. From this perspective, pilot attrition reflects necessary experimentation.

Others view the same pattern as evidence of weak capital discipline, suggesting that organisations overstate AI readiness while underinvesting in integration infrastructure and governance design.

A second divergence relates to whether responsibility should sit with boards, CIO functions, or operational teams. There is no shared view on where accountability for failed AI investments should formally reside.

Governance Tensions

The dominant tension is between innovation tolerance and capital discipline. Organisations are investing in AI at scale yet lack consistent frameworks to evaluate whether pilots represent genuine strategic capability development or fragmented experimentation.

This produces what participants implicitly described as a value realisation gap, where investment intensity exceeds demonstrable organisational integration.

Q3 How is your board reconciling the energy and water demands of AI infrastructure with existing Net Zero and sustainability commitments?

Areas of Consensus

Participants consistently acknowledged that AI expansion is creating direct tension with Net Zero commitments, given the energy intensity of cloud infrastructure, data centres, and computational scaling. Across sectors, environmental impact is largely externalised through third-party providers, limiting visibility within ESG reporting frameworks and embedding emissions within broader IT or Scope 3 categories.

There is also clear agreement that AI sustainability impacts are not yet fully integrated into governance systems, with a disconnect between AI strategy and ESG reporting. Survey evidence supports this, with most organisations indicating that environmental considerations are only partially incorporated into AI-related ESG or sustainability reporting.

This reinforces a structural misalignment between **digital transformation and sustainability governance**, where the environmental implications of AI adoption remain acknowledged but only partially embedded in formal oversight and disclosure practices.

To what extent has your organisation considered the environmental impact of AI projects within its ESG or sustainability reporting obligations?

Fully integrated 3
Partially considered 33
Not considered 26
I don't know 11

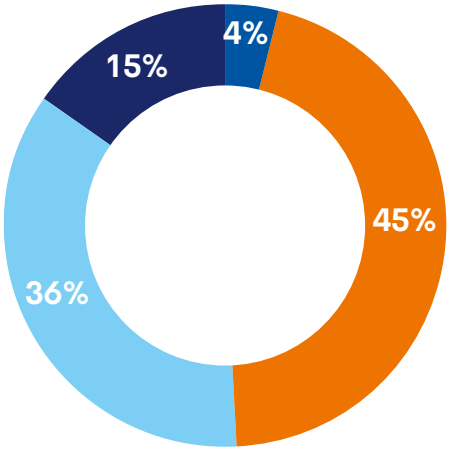

Areas of Disagreement

Corporate participants are more inclined to treat environmental impact as a trade-off within broader efficiency gains, taking the view that the operational benefits of AI are likely to outweigh associated sustainability costs. By contrast, public and charity sector participants express greater caution, highlighting the absence of clear reconciliation mechanisms between AI expansion and Net Zero targets. Public sector discussions emphasised the tension between digital transformation agendas and environmental commitments.

A further divergence exists on whether AI should be treated as a distinct ESG category. Some participants argue it should remain embedded within existing frameworks, while others suggest AI requires standalone environmental accounting treatment.

Governance Tensions

The central tension is between digital transformation ambition and sustainability credibility. Organisations are simultaneously increasing AI adoption while maintaining Net Zero commitments that are not explicitly adjusted for computational demand.

This creates a structural risk of ESG misalignment, where reported sustainability progress may not fully reflect underlying increases in computational intensity and infrastructure dependency.

Q4 Beyond efficiency gains, how can boards measure the strategic value of AI in accelerating decision-making and mitigating tail-end risks?

Areas of Consensus

Participants broadly agree that AI delivers value beyond cost efficiency, particularly in decision support, data analysis, and risk identification. AI is increasingly used to accelerate production of strategy materials, scenario analysis, and internal decision briefs.

There is also consensus that AI enhances organisational responsiveness by identifying risks and patterns that may not be visible through traditional analysis. Participants described use cases involving scenario testing, comparative analysis, and structured challenge generation for board discussions.

Tail-end risk mitigation is recognised implicitly. AI is viewed as improving completeness of analysis and reducing oversight gaps in large datasets, although this is not formally defined within governance frameworks.

Areas of Disagreement

A key divergence concerns whether strategic value can or should be formally measured. Some participants believe AI value is inherently qualitative and should be assessed through judgement-based governance processes.

Others argue that absence of structured metrics represents a governance weakness, particularly where significant capital is allocated without clear performance benchmarks or ROI frameworks.

There is also disagreement on whether AI should be evaluated as a productivity tool or as infrastructure-level capability contributing to organisational resilience.

Governance Tensions

The central tension lies between qualitative strategic judgement and quantitative capital discipline. While AI is widely acknowledged as strategically valuable, organisations lack consistent frameworks to measure its impact on decision quality, risk reduction, or organisational resilience.

This creates a governance gap where AI is assumed to deliver strategic benefit, but outcomes remain difficult to evidence or compare systematically.

Q5 Is quantum risk included in your board's five-year technology roadmap, or is it currently dismissed as a future issue?

Areas of Consensus

There is near-universal recognition that long-term technological risk remains underdeveloped within current governance frameworks, with board attention focused primarily on immediate AI deployment rather than long-horizon infrastructure exposure. While awareness of technological uncertainty exists, quantum computing risk is not embedded in board-level technology roadmaps.

Survey evidence confirms this gap, with most organisations reporting that quantum risk is not currently included in their five-year technology planning. Although participants acknowledge that existing cryptographic and data security systems underpinning AI infrastructure may be vulnerable to future quantum-enabled threats, this risk remains conceptual rather than operational within governance processes.

This reinforces a broader structural bias towards **short-term optimisation over long-term resilience**, where emerging systemic risks are recognised in principle but not yet translated into formal governance planning.

Is quantum risk included in your five-year technology roadmap?

Yes 9
No 40
Maybe/Under Consideration 24

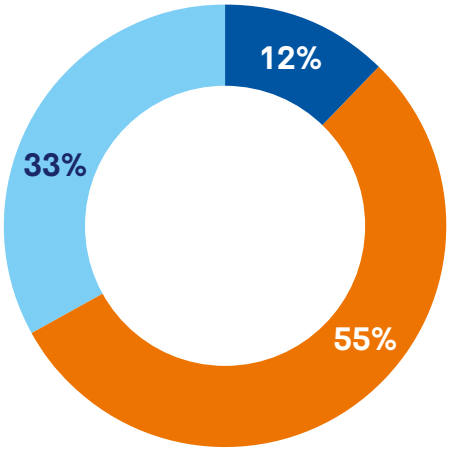

Areas of Disagreement

No meaningful divergence exists on whether quantum risk is actively integrated into current governance structures, as all participants indicate it is not.

The only variation lies in perceived urgency. Some participants treat quantum risk as a distant horizon issue unsuitable for near-term governance prioritisation. Others suggest it should already be incorporated into early-stage resilience planning.

Governance Tensions

The key tension is between short-term operational optimisation and long-term systemic resilience. Boards are optimising for immediate AI deployment benefits while underinvesting in structural technological risk preparedness.

This creates a latent governance exposure where current infrastructure decisions may embed future cryptographic and systemic vulnerabilities that are not yet reflected in strategic oversight frameworks.

Integrated Structural Insight: AI as Capital Allocation and Systemic Governance Risk

Across all questions and survey findings, the evidence converges on a consistent structural finding: AI is increasingly treated as a strategic capital allocation decision rather than a pure technology deployment issue.

However, governance systems remain fragmented across financial reporting, ESG oversight, technology governance, and risk management functions. This fragmentation produces three interrelated gaps.

First, a value realisation gap, where AI investment intensity exceeds demonstrable scaling outcomes and operational integration. Second, a disclosure gap, where AI capability is narratively overstated relative to verifiable governance evidence. Third, a resilience gap, where long-term risks including sustainability constraints and quantum vulnerability are not embedded in board-level planning.

Taken together, these findings reposition AI governance as a question of capital discipline, evidential integrity, sustainability alignment, and long-term systemic resilience.

The central governance challenge is no longer whether organisations should adopt AI, but whether they can do so in a way that produces measurable value, maintains disclosure credibility, preserves environmental commitments, and sustains technological resilience over time.

Governance Capability, Institutional Adaptation and Role Evolution

How governance authority is being redefined through capability, literacy, and institutional adaptation.

AI Literacy and the Evolution of Board Capability (From informed oversight to interpretive competence)

This strand examines the growing requirement for directors and governance professionals to develop interpretive competence in AI systems, particularly large language model (LLM) architectures embedded within governance, compliance, and administrative workflows. AI literacy is no longer a peripheral technical skill, but a core governance capability directly implicated in fiduciary oversight, regulatory defensibility, and the standard of reasonable care.

As LLM systems become embedded across governance architectures, boards can no longer treat them as passive infrastructure. They function as probabilistic decision-support systems, structured through retrieval pipelines, automated drafting, and workflow integration, introducing a fundamental epistemic shift. Governance actors are no longer reviewing linear outputs but supervising machine-mediated inference processes characterised by uncertainty, non-determinism, and latent error.

Within this environment, interpretive competence becomes the defining capability. Directors must understand not only what AI produces, but how outputs are constructed, where they may fail, and how underlying assumptions shape conclusions. Governance failure therefore shifts from procedural omission to the misinterpretation of probabilistic outputs as stable fact.

This reconfigures the governance competency model. Traditional literacy, anchored in financial interpretation and audit trails, is replaced by oversight of system architectures that continuously transform inputs into probabilistic outputs through opaque layers. Accordingly, governance professionals must move beyond awareness towards operational competence: challenging outputs, interrogating assumptions, tracing data provenance, identifying model risk, and assessing legal and strategic defensibility under uncertainty.

In this context, the inability to interrogate AI systems can no longer function as a credible defence. AI literacy becomes a fiduciary threshold, directly linked to standards of reasonableness in governance law and practice.

At the same time, this transformation reshapes the profession itself. Governance functions are moving from administrative execution to the design of decision environments, where outputs are conditioned not only by models but by how context is structured, constrained, and interpreted. This introduces a second-order capability layer centred on system design: on text engineering, prompt construction, workflow orchestration, and control architecture.

Across the literature, a clear convergence emerges: AI literacy is now a foundational governance obligation, and the central challenge is no longer adoption, but the development of interpretive capability to supervise increasingly autonomous, opaque, and probabilistic systems.

1. Previous AI Training and Preparedness

Have you or your board received any formal AI training in the past 12 months?

If so, what areas did it cover: basic awareness, ethical considerations, risk management, or technical skills such as prompt engineering?

If training was provided, how effective was it in preparing you to challenge AI outputs and integrate AI into governance processes?

Across sectors, the roundtable evidence indicates that AI training is present but structurally uneven and predominantly modular, pragmatic, and compliance-oriented rather than governance-theoretic or architecture-driven. Training typically combines internal mandatory modules, external executive education, and self-directed learning, with content focused on operational functionality: understanding probabilistic outputs, recognising hallucination risk, and developing prompting techniques for drafting and summarisation workflows.

Survey evidence reinforces this functional orientation. While most respondents report having received some form of AI training in the past 12 months, the content remains heavily weighted towards basic awareness (49 respondents), with significantly lower coverage of risk management (22), ethical considerations (18), and technical capability such as prompt engineering (17).

Have you or your board received any formal AI training in the past 12 months?

Yes: 49
No: 24

If yes, which areas did the training cover? (Select all that apply)

Basic AI awareness 49
Ethical considerations 18
Risk management 22
Technical skills (e.g., prompt engineering) 17

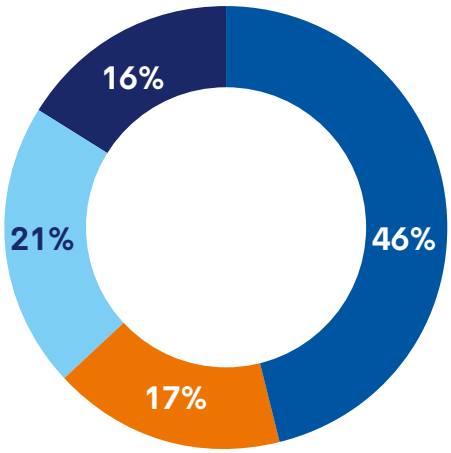

A consistent analytical finding is that training functions less as epistemic governance education and more as expectation calibration — correcting both over-reliance and excessive scepticism. This produces operational confidence, but not deep interpretive oversight capability. This limitation is reflected in training effectiveness scores, with an average rating of 2.4 (Q20), indicating that participants do not feel adequately equipped to challenge AI outputs or integrate them into governance processes.

How effective was the training in preparing you to challenge AI outputs and integrate AI into governance processes (from 1 to 5)?

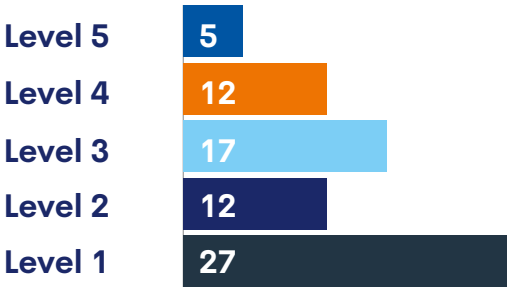

2.40

Average rating

Across organisational contexts, capability formation therefore remains uneven and often informal. In some environments, training is structured but narrow; in others, it is minimal or replaced by peer exchange and iterative experimentation. The result is a common structural condition: AI capability is functionally present but institutionally underdeveloped, with no consistent mechanism for assessing interpretive competence or embedding escalation discipline.

The governance tension that emerges is consistent across the dataset: **operational fluency in AI use is increasing faster than interpretive competence in oversight**, raising the risk that familiarity with systems is misrecognised as governance capability.

2. AI Training and Board Competence

How should boards structure AI training programmes to move beyond basic awareness and develop practical competence in interpreting AI outputs and associated risks?

Should AI literacy become a mandatory component of director induction and ongoing professional development, similar to financial literacy requirements?

Across the roundtable evidence, a clear normative shift is emerging: AI literacy is increasingly understood as a core governance competence, analogous to financial literacy in post-reform governance frameworks. Training is expected to move beyond awareness towards applied, experiential competence embedded in live governance workflows, including the interpretation of probabilistic outputs and integration into board-level decision-making. Survey responses strongly reinforce this expectation. A significant majority of respondents consider it extremely or highly important for boards to develop practical competence beyond basic awareness, and most support making AI literacy a mandatory component of director induction and ongoing professional development.

How important is it for boards to structure AI training programmes that go beyond basic awareness and develop practical competence?

Extremely important 46
Somewhat important 15
Neutral 7
Somewhat unimportant 4
Extremely unimportant 1

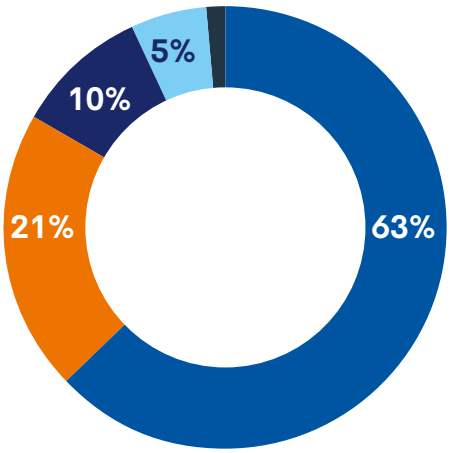

Should AI literacy become a mandatory component of director induction and ongoing professional development, similar to financial literacy requirements?

Yes 53
No 3
Not sure 17

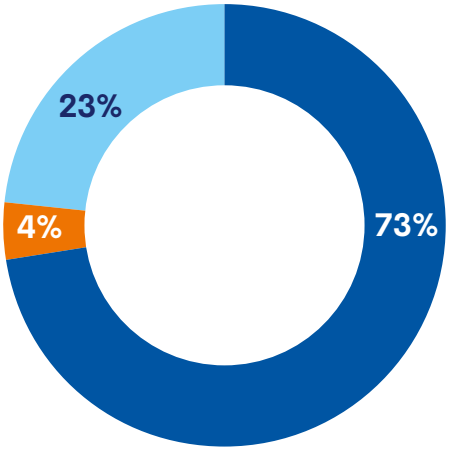

However, this expectation is not matched by institutional design. Across sectors, there is no standardised board-level curriculum, no codified induction framework, and no agreed competency threshold for directors. The result is a structural mismatch between **expected capability and formalised training provision**, particularly acute where boards rely on outputs they are not systematically trained to interrogate.

This produces a governance tension between assumed competence and unassessed interpretive capacity, reinforcing the broader capability gap identified across the study.

3. Prompting and Corrective Oversight

As boards increasingly rely on AI-generated insights, how can directors develop the ability to "prompt correctly" and challenge outputs effectively? And should governance professionals be trained in context engineering to ensure that prompts reflect organisational priorities and ethical standards?

Across the roundtable evidence, prompting is increasingly conceptualised not as a technical interface skill, but as a governance function embedded within system architecture. Outputs are understood as co-produced artefacts, shaped by both model design and input framing, with human cognitive responsibility remaining irreducible.

However, survey results reveal a clear gap between this conceptual recognition and practical capability. Most respondents report limited confidence in their ability to prompt effectively and to challenge outputs, indicating that operational reliance on AI systems is not matched by equivalent interpretive control.

At the same time, there is strong convergence on the need for formalisation. An overwhelming majority of respondents support structured training in context engineering with most highlighting demand for capability development that goes beyond ad hoc practice.

As boards increasingly rely on AI-generated insights, how confident are you in your ability to "prompt correctly" and challenge outputs effectively?

Promoters 10
Passives 25
Detractors 38

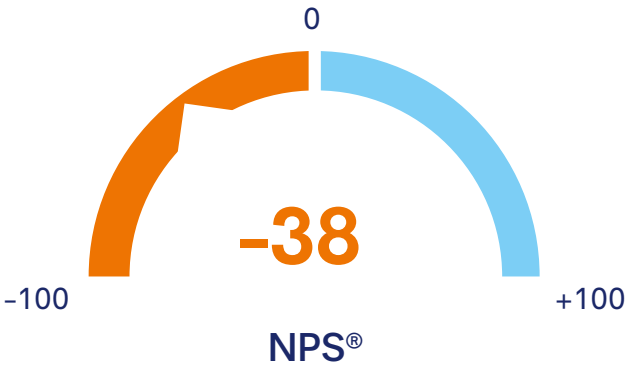

Should governance professionals receive formal training in context engineering to ensure prompts reflect organisational priorities and ethical standards?

Yes 68
No 2
Unsure 3

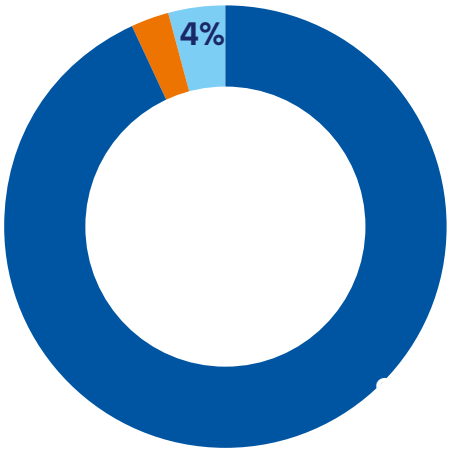

Despite this, prompting remains largely informal and inconsistently governed across organisations. There is little formal recognition of context engineering as a governance discipline, and limited integration into training or oversight frameworks. As a result, prompting practices are often decentralised, reliant on individual judgement, and uneven in quality, dependent largely on individual ability, aptitude and personal interest.

This produces a consistent vulnerability across sectors: automation bias combined with underdeveloped interpretive scaffolding, where outputs may be accepted or rejected without sufficiently rigorous interrogation of how they were produced.

Conclusion

Across the combined evidence, a single structural condition emerges: AI systems are already embedded in governance workflows, but interpretive governance competence has not been institutionally scaled to match.

Survey evidence consistently reinforces this gap. Training is widespread but shallow; competence expectations are high but unstandardised, and practical oversight capabilities—particularly in prompting, interpretation, and system interrogation—remain uneven and underdeveloped.

AI literacy is therefore evolving into a de facto fiduciary requirement but remains fragmented and insufficiently embedded within governance systems. The central challenge is no longer technological adoption, but the absence of structured capability frameworks capable of sustaining interpretive control under conditions of algorithmic uncertainty, system opacity, and distributed decision-making.

Governance Roles and the Reconfiguration of Institutional Authority

(When governance functions move from control to orchestration)

This theme examines how AI integration is reshaping governance work, shifting it from procedural execution toward interpretive, supervisory, and system-design functions. The central issue is no longer whether governance tasks can be automated, but how governance authority is reconstituted when large language models and automated workflows increasingly mediate the production of organisational knowledge.

Across corporate, charity, and public sector contexts, AI is already embedded in drafting, synthesis, reconciliation, and reporting processes. These systems do not replace governance judgement; they displace its location. Authority shifts from producing governance artefacts to validating, correcting, and contextualising machine-generated outputs. Governance therefore becomes less about generating outputs and more about governing the conditions under which they are produced and legitimised.

Within this environment, governance operates as continuous interpretation under probabilistic conditions. The central challenge becomes ensuring that machine-generated outputs remain institutionally coherent, legally defensible, and aligned with fiduciary intent in the absence of direct human authorship. It is within this tension that the governance professional is being redefined.

1. Re-skilling Governance Professionals as "Context Engineers"

If AI automates routine governance tasks, how should governance professionals re-skill to act as context engineers who ensure AI outputs reflect institutional memory and organisational culture?

The roundtable evidence indicates that this re-skilling is already emerging in practice, albeit unevenly and without formal recognition. Governance professionals are beginning to transition into implicit "context engineers" as AI systems take over drafting, transcription, and summarisation tasks. However, these systems operate without access to institutional memory, board nuance, or fiduciary interpretation.

As a result, re-skilling is being driven less by formal training and more by necessity. Governance authority is shifting from production to constraint: machines generate artefacts, while humans determine whether those artefacts are institutionally valid. Context engineering is not a technical specialism, but a fiduciary control function over probabilistic outputs.

In practice, this means that governance professionals are no longer primarily producing records, but reconstructing meaning. AI-generated minutes and summaries, while syntactically coherent, are frequently described as "flattened" or contextually incomplete, requiring active human correction. Survey evidence reinforces this point strongly: confidence in AI-generated board summaries preserving the unique voice of leadership is low, with the overwhelming majority of respondents classified as detractors (60), compared to very few promoters (4). This indicates that institutional voice, nuance, and intent are not being reliably captured by AI systems and must be actively reinserted through human judgement.

How confident are you that AI-generated board summaries maintain the unique voice of leadership?

Promoters 4
Passives 9
Detractors 60

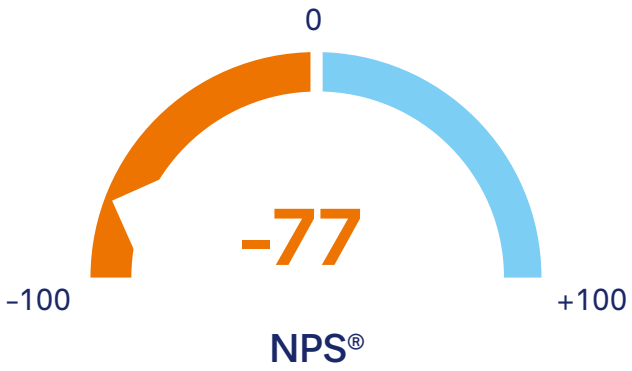

Across organisational contexts, approaches remain uneven. Some rely heavily on manual correction of tone and meaning, while others treat AI outputs as inherently unsafe unless interpreted through institutional context. What is consistent, however, is that the critical capability no longer lies in producing outputs, but in interpreting, constraining, and validating them against governance standards.

This reveals what re-skilling actually requires. It is not primarily about technical fluency, but about developing three interrelated capabilities.

- First, interpretive competence: the ability to interrogate outputs, identify omissions, and assess whether they are contextually and fiduciarily sound.
- Second, contextual encoding: the ability to embed institutional memory, organisational culture, and board-level nuance into AI-generated workflows and outputs.
- Third, system design: the ability to structure prompts, workflows, and controls so that outputs are shaped within governance boundaries rather than corrected after the fact.

Seen in this way, the governance professional is being redefined as the point at which organisational meaning is stabilised. The role of the Company Secretary increasingly centres on curating institutional memory, validating machine-generated outputs, and designing the constraints within which those outputs are produced.

Re-skilling therefore reflects a shift from executing governance processes to designing and supervising the conditions under which governance meaning is produced. AI does not remove the need for interpretation, it makes it structurally unavoidable

2. The Future Company Secretary

How should governance functions prepare future professionals?

This question exposes a deeper structural disruption: the breakdown of apprenticeship-based governance formation. Historically, governance competence was developed through progressive exposure to foundational tasks: minute-taking, statutory administration, register maintenance, drafting, and iterative correction. These activities were not merely administrative; they were the mechanisms through which judgement and institutional understanding were formed.

AI now automates many of these same tasks, creating a structural paradox: the processes that once produced governance judgement are the very processes being removed. As a result, preparing future professionals is no longer about optimising training pathways within existing models, but about redesigning how governance competence is formed in the absence of traditional experiential grounding.

Across the roundtable evidence, a consistent concern emerges that junior development pipelines are weakening. Reduced exposure to governance craft is eroding "learning-by-doing" pathways that historically anchored judgement formation. While AI improves efficiency in drafting and reporting, it simultaneously reduces structured engagement with the underlying discipline of governance work. In some settings, this is already prompting experimentation with alternative approaches, such as delaying automation exposure or prioritising shadowing, to preserve early-stage judgement development.

However, these responses remain fragmented and transitional. In many contexts, capability formation is becoming increasingly informal and uneven, with greater reliance on peer learning, ad hoc correction, and individual judgement. The result is a growing risk of competence without craft, where professionals can operate systems but lack the underlying interpretive foundation to evaluate outputs critically.

Within this context, the preparation of future governance professionals requires a fundamental shift in focus. If foundational exposure to procedural tasks is declining, then competency must be developed through alternative mechanisms that deliberately simulate and accelerate judgement formation. Emerging approaches across the evidence base point to three directions.

First, a shift towards scenario-based learning, where professionals are trained through exposure to ambiguous, high-risk governance situations rather than procedural repetition. Second, increased use of simulation environments, allowing individuals to practice interpreting and correcting AI-generated outputs under controlled conditions. Third, the development of supervised AI interaction frameworks, where early-career professionals engage with systems under structured oversight to build interpretive discipline rather than passive reliance; learning not just what is right, but why it is right.

What is changing, therefore, is not simply the content of training, but its underlying logic. Governance is transitioning from an apprenticeship model based on gradual experiential accumulation to an interpretive acceleration model, in which competence must be developed earlier, more intentionally, and with less reliance on embedded process exposure.

This transition creates a central tension between productivity gains and capability erosion. While AI compresses administrative workload, it simultaneously weakens the mechanisms through which governance judgement has historically been produced. The risk is not a shortage of technical capability, but a gradual erosion of the interpretive depth and understanding required for sound decision-making across fiduciary, strategic, and oversight functions.

Preparing future professionals therefore requires a deliberate reconfiguration of governance training systems. The focus must shift from procedural execution to interpretive capability, with structured mechanisms to replace the apprenticeship functions previously embedded in administrative work.

Governance functions will need to prioritise:

- earlier development of interpretive judgement
- structured exposure to system failure and ambiguity
- supervised interaction with AI-generated outputs
- explicit cultivation of institutional reasoning, not just process execution

In this configuration, governance capability is no longer built through accumulated administrative experience, but through designed exposure to complexity, uncertainty, and system-mediated decision environments.

Does the rise of Chief AI Officers risk creating a new silo, or can the Company Secretary position AI as a cross-functional governance priority?

This question concerns institutional control over AI governance systems. While the emergence of Chief AI Officer roles remains uneven and largely concentrated in larger organisations, it highlights a broader risk of fragmentation across IT, risk, data governance, and corporate secretariat functions. This fragmentation is not new, but reflects structural divisions that first emerged during earlier waves of digitalisation and are now being amplified by AI.

Across organisational contexts, AI is already embedded across enterprise platforms and workflows, making governance inherently cross-functional in practice. Yet formal ownership structures remain unclear. Governance is often distributed across functions, with no single point of accountability for how AI systems are integrated, monitored, and constrained.

In this environment, the Company Secretary is increasingly operating as an informal integration point, bridging legal, governance, risk, and operational domains. However, this role is not yet formally codified, and governance architecture has not been redesigned to reflect this integrative function.

At the same time, the introduction of dedicated AI roles risks reinforcing functional silos if responsibility becomes technically centralised but institutionally disconnected. Where ownership is unclear, accountability can become diffused, particularly in environments where AI outputs cut across multiple governance domains.

What emerges is not a settled model, but a structural tension between two competing organisational logics. One favours technical centralisation, with AI oversight concentrated in specialist roles. The other favours institutional integration, where governance functions coordinate AI as part of broader accountability frameworks.

The central tension lies between technical control and institutional coherence. Centralisation may improve technical oversight, but risks disconnecting AI governance from fiduciary and organisational context. Conversely, decentralised integration preserves contextual understanding but may weaken consistency and control.

Conclusion

Across the three roundtable questions, a consistent structural shift emerges. AI is not simply digitising governance processes; it is reshaping how governance authority is exercised: moving from procedural execution towards interpretive oversight of machine-mediated systems.

This shift is occurring through layering rather than redesign. New capabilities and roles are being added onto existing structures, leaving responsibility distributed across technical, risk, and governance domains. This fragmentation is not new, but reflects structural divisions established during earlier waves of digitalisation, now intensified by AI.

In this environment, the Company Secretary plays an increasingly important integrative role: not as a substitute for other functions, but as a coordinator across them. This includes supporting the board in maintaining visibility over AI-enabled decisions, aligning governance processes with technical systems, and ensuring that outputs remain robust, accountable, and decision-useful.

However, these integrative functions remain only partially reflected in formal governance frameworks. As a result, AI governance remains structurally incomplete. The priority is not to constrain organisational evolution, but to move towards a more coherent governance architecture in which technical capability, decision-making, and accountability are explicitly aligned.

Taken together, these developments indicate that AI does not reduce governance work but redefines it. Governance becomes less about producing outputs and more about validating, interpreting, and constraining outputs generated through probabilistic systems.

Accordingly, the governance professional is defined not only by technical understanding, but by the ability to exercise judgement across domains: integrating expertise, maintaining accountability, and ensuring that organisational decisions remain intelligible, defensible, and aligned with institutional objectives.

Modernising Governance Roles: AI-Enabled Context Engineering and the Transformation of Governance Work

The evidence indicates a shift in governance from process-based execution towards oversight of machine-mediated systems. This reflects not simply an efficiency gain, but a change in how governance authority is exercised as AI increasingly shapes the production and organisation of information.

Where governance was historically grounded in process and control, its emphasis is now shifting towards interpretation: assessing the reliability, context, and acceptability of AI-generated outputs. Governance professionals are therefore increasingly defined by their role in validating and contextualising system outputs rather than executing them.

This shift has three immediate implications: authority moves towards oversight and judgement as automation absorbs routine tasks; traditional pathways for developing judgement are weakened; and responsibility becomes more distributed across governance, risk, and technical domains.

The result is not a reduction in governance work, but its reconfiguration, creating the conditions for a more supervisory and system-oriented model of governance.

Structural Recomposition of Governance Work

Within this transformation, governance functions are moving from participation in workflows to supervision of them. As automation absorbs drafting, filing, and coordination tasks, Company Secretaries and governance teams are increasingly positioned above execution layers, responsible for validating outputs and ensuring they remain legally and institutionally robust.

The defining limitation of AI systems is not execution, but interpretation. While they can generate coherent outputs, they do not reliably capture organisational context, board dynamics, or decision intent. As a result, governance professionals are increasingly required to embed context, apply judgement, and assess whether outputs accurately reflect organisational reality and risk exposure.

At the same time, automation disrupts the internal development of governance capability. As routine and entry-level tasks decline, the pathways through which judgement was historically formed are weakened. This creates a structural tension between rising expectations of oversight and reduced experiential grounding. In response, capability must be developed earlier and more deliberately, with greater emphasis on system oversight, prompt structuring, and interpretive judgement.

Taken together, governance is evolving into an orchestration layer between institutional authority and machine systems. Routine production declines, while demand increases for professionals capable of supervising automated processes and translating governance judgement into system-level constraints and design parameters.

Within this configuration, governance increasingly centres on what may be described as "context engineering": structuring inputs, institutional memory, and constraints so that AI-generated outputs remain aligned with organisational purpose, legal obligations,

and decision-making requirements. AI may generate the "what", but governance defines the "why" and the limits within which outcomes are acceptable.

This reconfiguration reinforces the role of the Company Secretary as connective tissue between technical systems and board-level judgement. The role is not displaced, but repositioned: supporting the board in maintaining visibility over AI-enabled decision-making, aligning fragmented governance processes, and ensuring that outputs remain coherent, defensible, and institutionally intelligible.

However, this shift also introduces a capability challenge. As automation reduces exposure to foundational tasks, the development of interpretive competence can no longer be assumed through experience alone. It must be actively reconstructed through training, supervised interaction, and deliberate capability design.

The implication is clear. Governance competence is no longer defined primarily by procedural mastery, but by the ability to interpret, validate, and constrain AI-mediated decision processes under conditions of uncertainty and partial visibility.

At a structural level, analysis of governance tasks across seven domains indicates that more than half of activities are exposed to automation or augmentation. However, this exposure does not translate into uniform displacement. Instead, it produces a differentiated restructuring of governance work, in which process-driven activities are compressed while judgement-based, advisory, and accountability functions are reinforced.

For the purposes of this report, the company secretarial function was analysed across seven interdependent domains: Board and Committee Management; Entity and Statutory Administration; Director and Officer Lifecycle Processes; Share Capital and Corporate Transactions; Governance Frameworks and Policies; Advisory and Horizon Scanning; and Departmental Operations. Across these domains, a consistent pattern emerges. Routine governance processes are increasingly absorbed by GenAI-enabled workflows, while higher-liability and judgement-intensive responsibilities remain structurally resistant due to legal obligations and the persistence of individual accountability.

This structural analysis draws on established task-based exposure methodologies, including GLA Economics Working Paper 103, International Labour Organization frameworks developed by Gmyrek et al., and occupational exposure models advanced by Eloundou et al. Its contribution lies in applying these approaches at a granular occupational level, demonstrating that AI does not displace governance functions in aggregate, but reorganises them, shifting emphasis from execution to oversight, from production to interpretation, and from process compliance to institutional control.

Governance Exposure Framework

Exposure Category	Operational Characteristics	Primary Technical Mechanism
Programmatic Automation	Deterministic administrative workflows	APIs, workflow engines
Cognitive Augmentation	Drafting, synthesis, interpretation	LLMs / Copilot systems
Hybridised Orchestration	Integrated AI + workflow systems	RAG + orchestration pipelines
Insulated Fiduciary Core	High-liability judgement and accountability	Human legal authority

Quantitative Exposure Profile

Exposure Category	Task Volume	Percentage
Insulated Fiduciary Core	143	45.25%
Hybridised Orchestration	91	28.80%
Programmatic Automation	56	17.72%
Cognitive Augmentation	26	8.23%
Total	**316**	**100%**

E_total =

54.75%

Vector Dynamics of Governance Transformation

Within the 54.75 per cent exposure envelope, transformation is not uniform but distributed across three operational vectors.

Programmatic automation (17.72 per cent) consists primarily of deterministic governance processes such as statutory filings, confirmation statements, and administrative renewals. These are increasingly executed through workflow automation and API-based systems, with minimal reliance on generative AI. The effect is primarily the removal of repetitive administrative burden rather than the transformation of governance judgement.

Cognitive augmentation (8.23 per cent) represents cognitively intensive drafting and synthesis work, where LLMs as chatbots (e.g. Claude, Co-pilot, Gemini) are used as co-pilot systems for governance documentation. Here, the dominant shift is from first-draft production to validation and interpretive oversight. AI addresses the "blank-page problem", but governance professionals retain responsibility for correction, contextualisation, and assurance.

Hybridised orchestration (28.80 per cent) represents the most significant structural transformation. In this layer, governance workflows integrate structured data systems, AI drafting engines, and automated approval pipelines. Human intervention shifts away from production and towards system design, exception handling, and output validation. Governance professionals increasingly supervise systems rather than participate directly in execution.

The insulated fiduciary core (45.25 per cent) remains structurally resistant to automation. This is not due to technical limitations but due to legal and institutional constraints. Fiduciary accountability cannot be delegated to automated systems, and governance authority remains legally anchored in identifiable human actors. This layer also includes relational and behavioural dimensions of governance work, stakeholder management, board dynamics, strategic advisory judgement, which remain fundamentally resistant to codification or automation.

Strategic Implications for Governance Work

The cumulative implication of these shifts is that GenAI does not eliminate governance roles but fundamentally restructures their internal composition. Governance work is migrating from a production model
to a verification and assurance model. As organisations automate administrative workflows and accelerate drafting processes, governance professionals increasingly operate as system auditors, oversight designers, and institutional risk managers.

This transformation aligns with wider labour-market evidence indicating that highly educated knowledge workers experience high task-level exposure to AI but relatively low occupational displacement. Routine administrative functions compress, while verification, oversight, and interpretive assurance functions expand in strategic importance.

The resulting organisational requirement is not workforce reduction but capability reconfiguration. Governance professionals must increasingly combine traditional fiduciary judgement with technical fluency in system oversight, workflow architecture, and AI-mediated decision validation. In this configuration, governance authority is no longer grounded in production capacity, but in interpretive legitimacy under conditions of machine-generated uncertainty.

Conclusion

Across all dimensions of the analysis, the central transformation is consistent and irreversible in direction, even if uneven in implementation.

Governance is shifting from procedural administration to interpretive orchestration.

This produces a governance architecture defined by four structural properties: interpretive authority displacing procedural authority, the collapse of traditional apprenticeship pathways, the engineering of institutional memory through system design, and the distribution of governance authority across interconnected organisational functions.

AI does not reduce governance complexity. It intensifies it. It converts governance into a continuous process of interpretive supervision over machine-mediated systems operating under probabilistic conditions.

The governance professional is therefore no longer defined by what they produce, but by what they are able to validate, constrain, interpret, and prevent from being incorrectly institutionalised.

Integrated Governance Synthesis and Sectoral Implications

What the themes are telling us

Across all themes, the evidence converges on a consistent finding: AI is operationally embedded in governance work but strategically under-structured in how it is governed, evaluated, and scaled.

AI is not functioning as a transformational governance technology in a system-wide sense. Instead, it operates as a situational productivity layer, improving drafting, summarisation, and reporting, without materially restructuring governance architecture. The result is a central paradox: AI is widely used in practice but rarely governed as a system.

Core transformation pattern

Across the dataset, a consistent structural pattern emerges:

AI adoption is **high at the task level** but **immature at the system level**.

Operational use is widespread, particularly in drafting, minutes, and reporting, while structural governance remains underdeveloped, including accountability frameworks, measurement systems, and ESG integration.

The dominant shift is therefore not transformation, but **assisted interpretation**: AI absorbs procedural workload, while governance authority remains human and is exercised through validation, correction, and contextual judgement.

AI is not systematically governed

The most consistent finding is the absence of formal governance infrastructure. There is limited evidence of:

- ROI measurement
- pilot evaluation frameworks
- scaling metrics
- structured accountability models

Survey evidence reinforces this gap. A majority of organisations report having no board-approved AI strategy, while shadow AI usage occurs at least occasionally in most cases. At the same time, cost–benefit analysis is inconsistent, with a significant proportion of organisations either not conducting it or still in progress.

Taken together, this indicates that AI adoption is **activity-rich but evaluation-poor**, with organisations scaling usage without corresponding governance discipline.

AI value is informal, not institutionalised

Across themes, AI value is understood in practice but not formalised in governance systems. Benefits are consistently identified, for example efficiency gains, time savings, and reduced administrative burden, but are offset by validation workload, risk exposure, and output inconsistency.

Cost–benefit reasoning exists, but remains implicit, fragmented, and reliant on practitioner judgement rather than structured frameworks. As a result, value is experienced locally but not aggregated or measured at system level.

ESG and sustainability integration is absent

A critical systemic gap is the absence of ESG integration in AI governance. Environmental impacts, particularly energy usage and infrastructure demands, are rarely incorporated into governance or reporting frameworks.

Survey evidence confirms this disconnect: most organisations report that environmental considerations are only partially integrated, or not considered at all, within AI-related ESG reporting.

This places ESG structurally outside the governance framing of AI, despite its growing strategic relevance.

Core structural tensions

Four persistent tensions define the current governance system:

- Adoption vs governance gap: AI usage is widespread, but governance structures lag behind
- Efficiency vs control: productivity gains are offset by validation burden and risk exposure
- Innovation vs measurement deficit: experimentation is high, but evaluation frameworks are absent
- Automation vs judgement dependency: procedural workload declines, but reliance on human interpretation increases

System-level implications

Taken together, the findings indicate that AI is augmentative, not transformational at system level. It improves drafting speed, summarisation, and information processing, but does not replace governance judgement or restructure accountability.

Governance maturity remains low to moderate, characterised by high experimentation but weak formalisation and limited oversight infrastructure.

The most significant risks are not technical, but behavioural: over-trust in outputs, under-validation, uneven prompting capability, and growing dependency, particularly among less experienced users.

Survey evidence reinforces a structural capability gap. Board AI literacy remains uneven, with many respondents indicating reliance on management or external expertise rather than independent understanding. At the same time, demand for support is concentrated on practical tools such as prompting guidance and technical checklists, rather than high-level strategy.

This produces a clear stratification: operational use spreads across organisations, but interpretive authority remains concentrated among more experienced governance actors.

Value realisation pattern

Across themes, AI value follows a consistent pattern:

- Value generation is concentrated in low-complexity tasks such as drafting and summarisation
- Value leakage occurs through validation overhead, correction cycles, and output inconsistency

The net effect is partial efficiency gain, structurally diluted by governance requirements.

Final synthesis

Across all themes, the conclusion is consistent: AI in governance is operationally active but institutionally immature. Organisations are actively using AI, but are not systematically governing it, measuring its value, integrating ESG considerations, or standardising accountability.

The central governance problem has therefore shifted. It is no longer about adoption, but about how governance systems are designed to evaluate, constrain, and validate machine-mediated outputs.

Crucially, the limiting factor is not technological capability, but governance capacity. The binding constraints are:

- capability (skills and judgement)
- governance design (lack of frameworks)
- behavioural risk (over-trust and misuse)
- measurement absence (no consistent ROI or scaling logic)

AI remains operationally useful but strategically fragmented. Without structured frameworks for accountability, performance measurement, training, and ESG integration, it risks becoming embedded without being governed.

Sectoral Implications

Drawing on roundtable discussions and shifting the lens from thematic analysis to sectoral perspective, a consistent structural pattern emerges across corporate, charity/NFP, and public sector contexts. Generative AI does not displace governance labour; rather, it compresses administrative activity while redistributing effort towards higher-value functions, including oversight, verification, and interpretive control.

Despite differing institutional logics, commercial efficiency, mission delivery, and procedural accountability, all sectors converge on a shared outcome: automation absorbs routine governance production while intensifying reliance on human judgement and institutional accountability.

Shared structural conditions of adoption

Across sectors, AI adoption is shaped by a stable set of constraints: limited budgets and staffing, high regulatory pressure, administrative overload, fragmented legacy systems, and uneven digital capability.

These conditions ensure that AI is deployed primarily as an efficiency and continuity mechanism rather than a fully transformative governance technology. The result is not redesign, but **incremental integration layered onto existing governance systems.**

Cross-sector transformation pattern

Across all three sectors, governance is shifting from manual coordination to AI-mediated workflow orchestration. Drafting, reporting, and coordination processes are increasingly automated, reducing production time and administrative burden.

However, this does not reduce governance responsibility. Instead, it redistributes it. Humans move from producing artefacts to validating and supervising machine-generated outputs.

The underlying shift can be summarised as:

- production → verification
- coordination → oversight
- administration → assurance

Corporate sector: reconfiguration under constraint

The corporate sector represents the most advanced but also most constrained environment for AI-enabled governance. Adoption is driven by regulatory pressure, competitive dynamics, and scaling requirements, but bounded by legal accountability.

AI is widely used for drafting, reporting, and regulatory analysis. However, directors retain personal liability for disclosure accuracy, creating a clear accountability boundary. This produces an "accountability bottleneck": AI supports preparation but cannot participate in final judgement.

As a result, governance functions are reoriented towards **verification, evidential control, and disclosure assurance**, with the corporate secretariat evolving into a risk and systems oversight role.

Charity and NFP sector: capacity expansion

In the charity sector, AI functions primarily as a capacity-enhancing tool within resource-constrained environments. Efficiency gains are not used for workforce reduction but redirected into mission delivery and service provision.

AI is most applied to drafting, reporting, and administrative coordination, reducing operational burden and expanding delivery capacity.

However, this creates a structural vulnerability. Lower levels of technical capability increase the risk of over-reliance on AI-generated outputs, with reduced scrutiny and weaker challenge mechanisms. This produces a **vigilance gap**, where oversight capacity does not scale with automation.

Public sector: compression under accountability pressure

In the public sector, AI adoption is driven by administrative overload and resource constraints but shaped by stronger accountability and transparency requirements. AI is used to stabilise operations: supporting records management, compliance processes, and workflow coordination.

Unlike the charity sector, efficiency gains are not redirected towards expansion, but towards maintaining continuity under pressure.

However, a longer-term structural risk emerges. As automation absorbs entry-level tasks, the pathways through which governance capability is developed begin to erode. This creates a training bypass, where fewer opportunities exist to build procedural judgement and institutional knowledge.

Cross-sector risks and convergence

Despite sectoral differences, three shared risks emerge:

- a vigilance gap, were oversight capacity lags behind automation
- a training bypass, where capability formation pathways are weakened
- an accountability asymmetry, where legal responsibility remains human despite machine-mediated production.

Across all contexts, governance risk is therefore not driven by technical failure, but by misalignment between automation and oversight capacity.

System-level transformation

Taken together, GenAI produces a consistent institutional shift: governance moves from manual coordination toward supervised, AI-mediated administration.

Automation compresses production, but elevates the importance of interpretation, auditability, and control. Governance becomes less about generating outputs and more about ensuring that those outputs remain accurate, defensible, and aligned with institutional intent.

Conclusion

Across all three sectors, GenAI does not reduce governance demand; it reconfigures its internal structure. Administrative work declines, while responsibility is redistributed upward into supervision, verification, and interpretive control.

Survey evidence reinforces where organisations feel least equipped to respond. Boards prioritise practical support, particularly prompting guidance and technical assurance tools over abstract strategy, indicating that the immediate need is not conceptual direction, but operational capability.

Which resource would most assist your board?

Prompting guide for context engineering 33

Library of failure-mode case studies 10

Technical checklist for AI vendor due diligence 26

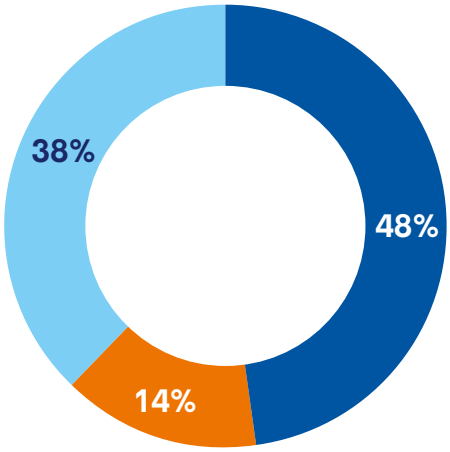

The sectoral divergence lies primarily in where efficiency gains are absorbed:

- compliance and assurance in the corporate sector
- mission delivery in charities
- operational continuity in the public sector

However, the shared outcome is clear. Governance is no longer centred on producing records, but on sustaining defensible oversight within machine-mediated systems.

AI does not replace governance labour. It intensifies reliance on human judgement under conditions of increasing system complexity, distributed accountability, and continuous automated production.

Collection of Good Practices

Drawing on the research findings, this section sets out a practical collection of good practices to guide organisations in governing their use of artificial intelligence.

It is not a prescriptive or purely technical model, but a pragmatic guide to help organisations assess their position in the AI journey and identify the governance structures required next. It translates cross-sector insights into clear, actionable steps, enabling boards and governance professionals to move from fragmented experimentation to structured, controlled, and accountable deployment.

The framework is directly aligned with the questions highlighted in the Executive Summary, providing the mechanisms needed to answer them with confidence. This includes strengthening oversight structures, clarifying accountability, embedding robust control environments, and ensuring that AI-enabled decisions remain explainable and defensible.

In particular, it supports organisations in:

- aligning governance practices with directors' fiduciary duties and statutory obligations,
- embedding AI within risk management and internal control frameworks, and
- ensuring board-level visibility and assurance over AI use across the organisation.

As a technical reference point, organisations should use ISO/IEC 42001:2023 to inform the design of AI management systems and control architectures. This framework complements that standard by focusing on the governance, oversight, and accountability structures required to operationalise it effectively at board level.

Designed for application across listed companies, financial institutions, public bodies, and charities, it can be applied proportionately based on organisational scale, regulatory exposure, and AI maturity. It integrates four core governance components:

- Accountability architecture (who is responsible)
- Control architecture (how AI is governed)
- Capability architecture (who is equipped to oversee it)
- Assurance architecture (how governance is evidenced)

AI systems may generate, analyse, and execute workflows, but they do not hold authority. All fiduciary responsibility remains with the board and legally accountable officers.

In essence, this collection of good practices is a starting point for organisation to start thinking about developing an organisational framework to maintaining trust, accountability, and effective governance in AI-enabled organisations.

Foundational governance principles

All AI-enabled governance activity operates under five binding principles.

1. First, **non-delegable accountability**: the Board retains ultimate responsibility for all decisions, outputs, and disclosures, whether or not derived from AI.
2. Second, **liability follows control**: accountability attaches to those who design, deploy, approve, or materially influence system behaviour.
3. Third, **systems-based governance:** AI must be governed as a socio-technical system, not as a standalone tool or model.
4. Fourth, **lifecycle assurance**: governance applies continuously across procurement, deployment, operation, monitoring, and retirement.
5. Fifth, **auditability by design:** AI outputs must be traceable to data sources, system versions, prompts or rules, and human decision points.

AI governance operating model

AI systems operate across three execution tiers.

Deterministic systems automate rule-based administrative or compliance-driven tasks, such as filings, workflow routing, and rule-triggered notifications. They operate within predefined parameters, require strict validation rules, and produce structured, auditable outputs.

Generative and co-pilot systems support drafting, summarisation, synthesis, retrieval, and analytical assistance. They enhance information processing and document production, but require human verification of accuracy, context, and appropriateness.

Agentic systems initiate or coordinate multi-step workflows under defined conditions. They may support execution but cannot undertake high-impact actions without human approval and do not exercise fiduciary judgement.

The core rule is constant across all three tiers: AI systems may support decision-making and execution, but they do not exercise authority or fiduciary discretion.

Governance domains affected by AI

AI transforms governance across seven substantive domains. It reshapes **fiduciary duty**, requiring boards to interrogate AI-informed outputs under evolving standards of reasonable care, skill, and diligence. It affects **internal control and audit integrity**, requiring evidential traceability and defensible oversight. It intensifies **risk oversight and discoverability exposure**, including record contamination, discoverability expansion, and loss of privilege.
It raises **delegation and agentic threshold** questions, demanding clear escalation rules for when AI may analyse, recommend, or act. It creates new demands for **ESG and sustainability coherence**, especially where infrastructure intensity is not yet reflected in governance or reporting frameworks.
It drives **role reconfiguration**, moving governance professionals from clerical execution towards interpretive and orchestration functions. And it elevates **board capability and AI literacy**, requiring operational competence rather than conceptual awareness alone.

Liability and accountability architecture

Responsibility must be allocated across the governance chain.

The **Board of Directors** retains ultimate fiduciary accountability.

Executive management holds operational accountability for implementation.

A **Chief AI Officer** or equivalent, where such a role exists, may hold responsibility for lifecycle governance and coordination.

Model or product owners are responsible for performance, testing, and monitoring.

Vendors and third parties hold contractual and technical responsibility within defined service boundaries.

Four liability principles follow:

1. Control determines accountability.
2. Foreseeability increases liability exposure.
3. Delegation does not dilute fiduciary duty.
4. Lack of auditability increases legal and regulatory risk.

AI governance control architecture across the lifecycle

AI governance must apply across the full system lifecycle.

At **procurement**, organisations should conduct vendor due diligence, assess data provenance, secure contractual audit rights, and evaluate external dependency risk.

At **deployment**, they should classify risk, define approved use cases, establish escalation thresholds, and require board sign-off for high-risk implementations.

During **operation**, they should maintain monitoring, drift detection, incident escalation, and human override pathways.

At **retirement**, they should manage data retention and deletion, model decommissioning, audit closure, and transition controls. This ensures governance continuity across system evolution rather than point-in-time compliance.

Internal control and assurance model

AI assurance should operate through an adapted three-lines structure.

The **first line** comprises operational users and business owners managing AI outputs within approved environments.

The **second line** comprises risk, compliance, legal, and governance functions overseeing adherence to policies, classifications, escalation rules, and usage controls.

The **third line** comprises internal audit or equivalent independent assurance functions reviewing control effectiveness, system performance, and governance integrity.

Mandatory assurance mechanisms should include model traceability logs, dataset provenance tracking, prompt and rule version control, human override documentation, and incident reporting and escalation trails.

Sectoral application and proportionality

The framework is designed for proportional application across sectors. Automation scale varies, but fiduciary accountability does not.

In **listed and corporate entities**, AI typically operates within a risk assurance model characterised by modular automation, strong disclosure controls, and strict verification.

In **financial and infrastructure organisations**, adoption follows a controlled integration model with higher friction, stronger regulatory scrutiny, and tighter model governance.

In the **public sector**, AI is used through an operational pressure model, supporting continuity under transparency and accountability constraints.

In **charities and NFPs**, AI functions within a capacity expansion model, where efficiency gains are redirected into mission delivery rather than workforce reduction.

AI governance maturity and readiness model

Organisational readiness should be assessed across five maturity levels.

Level 1: Ad hoc — shadow AI and no meaningful governance visibility.

Level 2: Repeatable — informal policies and inconsistent controls.

Level 3: Defined — approved environments, clearer accountability, and structured governance.

Level 4: Managed — measurable controls, monitoring, risk metrics, and established assurance mechanisms.

Level 5: Optimised — continuous governance loops, integrated assurance, and adaptive controls. Most organisations currently remain at Levels 1–2, where AI adoption outpaces governance capability.

Operational control allocation

To translate principles into execution, responsibility should be allocated across key operational control domains: **strategy**, including board oversight and risk appetite; risk, including classification and escalation thresholds; **models and systems**, including validation, testing, and drift monitoring; data, including provenance, integrity, and access control; **vendors**, including contractual governance and audit rights; **operations**, including human-in-the-loop supervision and process compliance; and **assurance**, including independent review, audit, and control validation.

Key systemic failure modes

The framework anticipates recurring governance failure modes, including AI-generated misstatement or "AI-washing", loss of explainability at scale, vendor dependency and concentration risk, automation bias in decision pathways, fragmented accountability chains, silent model drift, and shadow AI operating outside formal governance visibility.

Implementation reality and governance constraints

Effective deployment will be constrained by shortages in AI governance capability, rapid model iteration cycles, vendor opacity, and governance fatigue. Mitigation requires mandatory AI literacy for directors and governance professionals, standardised governance templates and controls, external assurance or independent model audits where appropriate, and the embedding of governance requirements into system design rather than reporting layers alone.

Final governance principle

This framework establishes a single operating constraint across all AI-enabled governance systems: **AI may generate, support, and execute defined governance processes, but it cannot assume fiduciary authority**. The future of governance depends not on limiting AI capability, but on ensuring that accountability, auditability, and institutional judgement remain structurally non-automatable.

Conclusion

The reconfiguration of corporate governance by artificial intelligence is not a temporary adjustment but a structural shift in organisational decision-making. As boards embed probabilistic and non-deterministic systems into core processes, the statutory standard of reasonable care under Section 174 of the Companies Act 2006 increasingly requires critical interrogation rather than passive reliance. Technological ignorance is becoming less and less defensible where AI materially informs governance, disclosure, or strategic judgement.

The findings show that AI is not displacing the governance function but reorganising it. Our analysis of the task governance architecture indicates a total technological exposure of more than 50%, with administrative and drafting tasks most exposed to automation and augmentation, while fiduciary judgement remains non-delegable. The central challenge is therefore not whether AI can be adopted, but whether governance systems are evolving quickly enough to supervise it. A clear implementation gap persists while 76 per cent of organisations report management-level oversight, only 41 per cent have made AI policies accessible to daily users, leaving many organisations operating with formal intent but weak operational control. Across sectors, this gap takes different forms: more advanced but tightly constrained in the corporate sector, fragmented in the public sector, and informal and capacity-driven in charities, but the underlying problem is consistent: AI use is outpacing governance design.

Looking ahead, the most significant risks are not only immediate, but structural. Organisations will need to address the erosion of apprenticeship-based capability as routine governance tasks are automated, or risk weakening the professional judgement on which future governance leadership depends. They will also need to reconcile AI adoption with wider sustainability commitments, particularly where computational intensity is not yet reflected in ESG oversight and begin incorporating longer-horizon threats such as quantum risk into technology planning.

Taken together, the evidence points to a clear conclusion: the future of governance will depend less on whether organisations use AI, and more on whether they can preserve accountability, interpretive control, and institutional judgement within increasingly machine-mediated systems. In that environment, AI may assume a growing share of execution, but responsibility remains firmly human.

cgi.org.uk

Chartered Governance Institute UK & Ireland

The Chartered Governance Institute is the professional body for governance and the qualifying and membership body for governance professionals across all sectors. Its purpose under Royal Charter is to lead effective governance and efficient administration of commerce, industry, and public affairs working with regulators and policymakers to champion high standards of governance and providing qualifications, training, and guidance. As a lifelong learning partner, the Institute helps governance professionals achieve their professional goals, providing recognition, community, and the voice of its membership.

One of nine divisions of the global Chartered Governance Institute, which was established 135 years ago, the Chartered Governance Institute UK & Ireland represents members working and studying in the UK and Ireland and many other countries and regions including the Caribbean, parts of Africa and the Middle East.

cgi.org.uk

The Chartered Governance Institute
Saffron House
6–10 Kirby Street
London EC1N 8TS

Phone: 020 7580 4741
Email: enquiries@cgi.org.uk
Web: cgi.org.uk

X: @CGIUKI
LinkedIn: CGIUKI
Facebook: CGIUKI

The Chartered Governance Institute

2026_165_0626