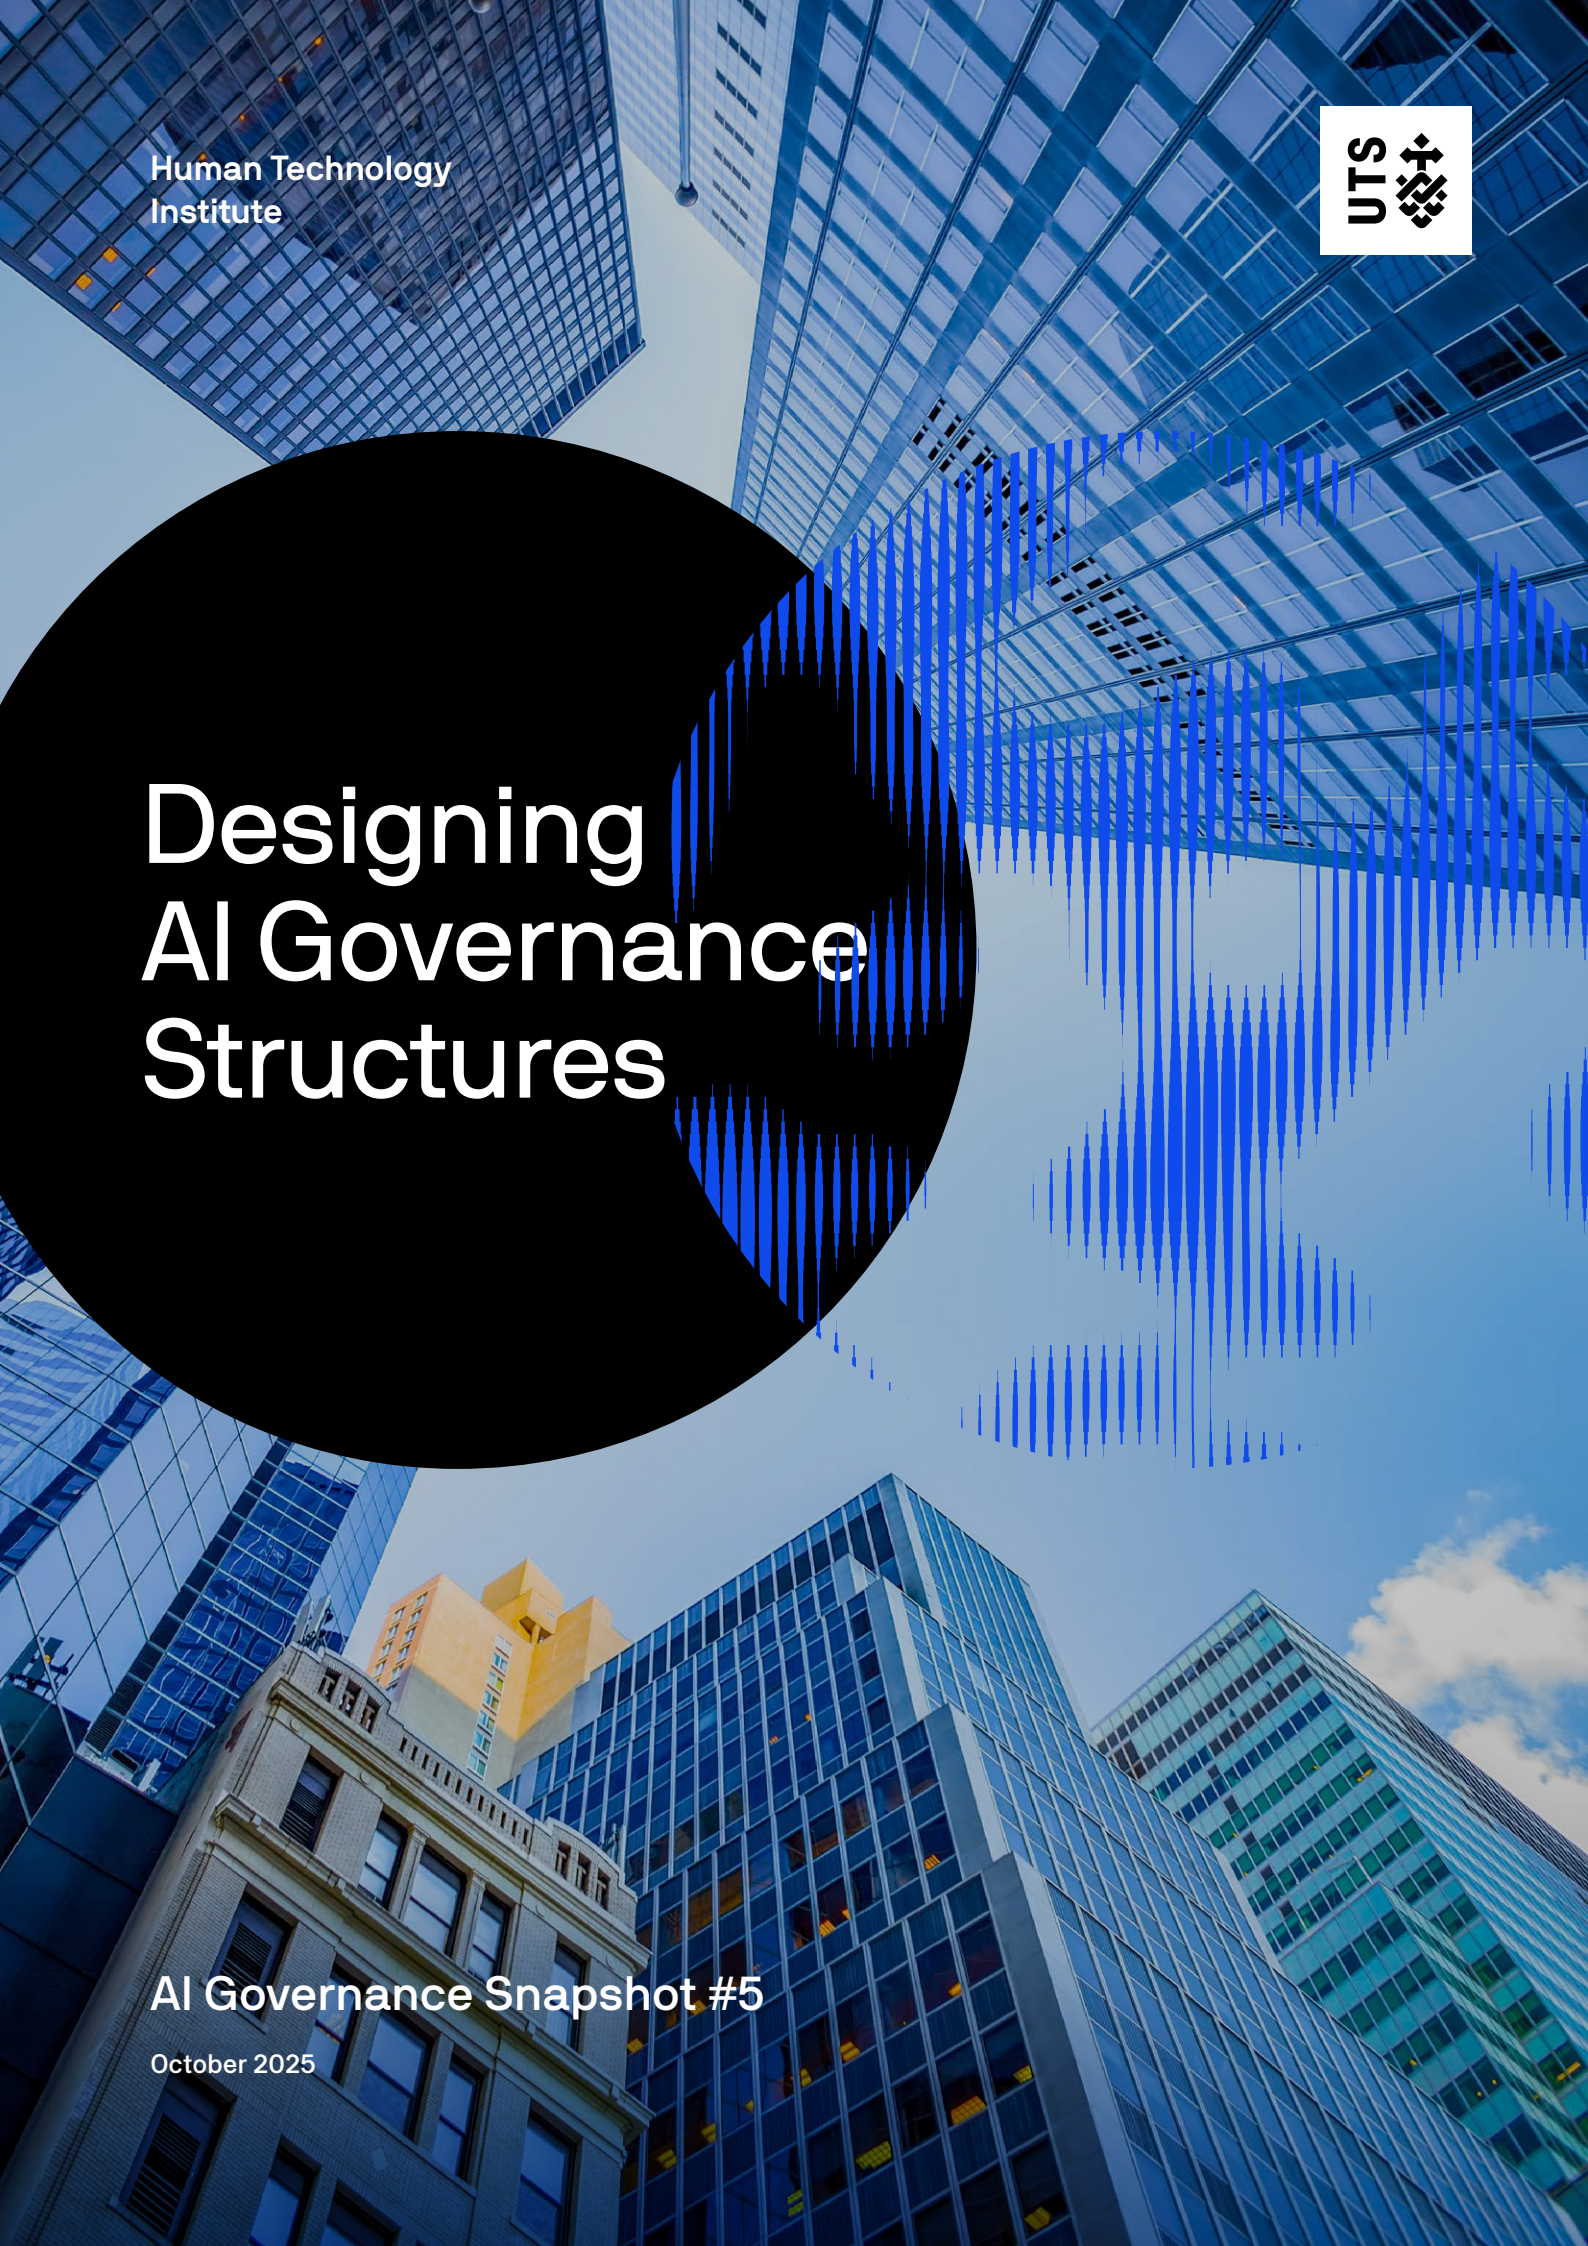
Human Technology
Institute
UTS
Designing
AI Governance
Structures
AI Governance Snapshot #5
October 2025

The UTS Human Technology Institute (HTI) is an impact-oriented institute building human values into new technologies. Bringing together policy, legal and technical experts, HTI provides independent expert advice, policy development, capability building, and data science solutions to support government, industry and civil society.

Authors: Llewellyn Spink and Nicholas Davis.
HTI contributors: Gaby Carney, Myfanwy Wallwork. Jack Goldsmith, and Gianluca Pecora.

Acknowledgement of Country

UTS acknowledges the Gadigal people of the Eora Nation, the Boorooberongal people of the Dharug Nation, the Bidiagal people and the Gamaygal people upon whose ancestral lands our university stands. We would also like to pay respect to the Elders both past and present, acknowledging them as the traditional custodians of knowledge for these lands.

Contents

Executive summary

AI governance structures are the architecture of responsible AI

As organisations rapidly adopt AI systems, they face a critical challenge: how to capture transformative benefits while managing new and amplified risks. AI governance structures assist in this endeavour, serving as the formal mechanisms through which organisations exercise oversight of AI. These structures define decision-making processes, establish clear lines of responsibility and accountability, and create reporting and escalation pathways, often operationalised through dedicated bodies like AI councils or committees.

Effective governance structures directly drive business value. They accelerate AI adoption by establishing clear guardrails that give teams confidence to innovate within defined boundaries. They prevent costly mistakes by enabling early evaluation of strategic value and preventing investment in AI use cases that are legally, ethically, or commercially unsound. Most importantly, they transform abstract AI principles into concrete decision-making frameworks, ensuring that AI initiatives align with organisational values while managing risk.

Governance structures are a core component of an organisation's AI governance operating model. They convert organisational directions (informed by AI strategy and risk appetite) into operational movement (through rules and processes for AI use). To function effectively, they require key enablers, including organisational culture, staff capabilities, and technical infrastructure.

The need for robust AI governance structures has reached an inflection point. There are growing regulatory requirements for effective AI governance (such as the EU AI Act where key obligations came into force in February 2025), alongside clear expectations for best practice established through global and national standards frameworks (such as ISO 421001 and NIST AI Risk Management Framework).

As AI deployments shift from isolated pilots to enterprise-wide systems affecting millions of customers and critical business decisions, leading organisations understand that informal AI oversight is not sufficient. They are therefore investing in formal governance structures not just to manage downside risks, but to build competitive advantage through trusted, scalable AI deployment. Those with mature governance structures report faster AI adoption rates, fewer failed projects, and greater stakeholder confidence.

To that end, this snapshot equips corporate leaders and governance professionals with the tools and insights needed to start building AI governance structures suited to their organisation's unique context and needs.

This snapshot does not address all aspects of AI governance structures. Instead, it focuses on two foundational questions: the form that an AI governance structure can take, and the location of authority for the structure.

Governance Structures define decision-making processes, establish clear lines of responsibility and accountability, and create reporting and escalation pathways, often operationalised through dedicated bodies like AI councils or committees.

Designing AI governance structures

While designing AI governance structures involves numerous decisions, two fundamental choices will define how your organisation governs AI:

1. Form of structure

Should your organisation establish a new, dedicated AI governance structure, or adapt existing structures to include AI oversight?

- Dedicated AI governance structures (particularly when operationalised as governance bodies, such as AI committees) provide focused oversight, consolidate expertise, and establish clear accountability.
- Existing structures can reduce duplication and costs, but should be appropriately adapted to address the distinct risks and opportunities of AI.
- Many organisations start with dedicated structures to build capability, then plan to gradually integrate AI governance into existing frameworks as maturity grows.

2. Location of authority

Should AI governance authority be centralised, decentralised or hybrid?

- Centralised approaches support consistency, accountability, and strategic alignment but can create bottlenecks if all AI-related decisions are made by a single body.
- Decentralised approaches enable responsiveness and domain expertise, but risk fragmentation and uneven governance.
- Hybrid approaches combine centralised oversight (particularly for high-risk AI) with decentralised implementation and assurance.
- The optimal approach often depends on organisation size, with smaller organisations favouring centralised models and larger enterprises adopting hybrid approaches.

Building and maintaining effective AI governance structures

AI governance structures are effective when they shape how AI systems are adopted and managed across the organisation. To do so, they must be built on three foundations: clear accountability, cross-functional integration, and formal authority. Without these, they risk being ignored or ineffectual.

Even well-built structures require active maintenance. That means regular meetings, ongoing monitoring and reporting, and meaningful stakeholder engagement rather than box-ticking exercises. Otherwise, these structures only offer the appearance of governance without real oversight.

Recommendations and conclusions

While there is no one-size-fits-all approach to AI governance, our analysis suggests that many organisations, particularly large, complex ones or those with significant AI investments, will benefit from initially establishing dedicated AI governance structures with centralised decision-making authority. Such structures often serve as valuable transitional arrangements. They build organisational AI literacy and capability during the early stages of AI adoption, creating the foundation for eventually embedding AI governance into business-as-usual frameworks.

Ultimately, however, the right approach will depend on your organisation's size, AI maturity, strategic objectives, risk appetite, set of legal obligations or sector-specific requirements and AI use, existing governance mechanisms. Governance structures are a core component of an organisation's AI governance operating model, and decisions about their design must reflect the organisation's context and needs.

With proper authority, accountability, and integration, these structures transform AI governance from a compliance checkbox into a strategic enabler, providing the practical oversight needed to unlock AI's value while managing its risks.

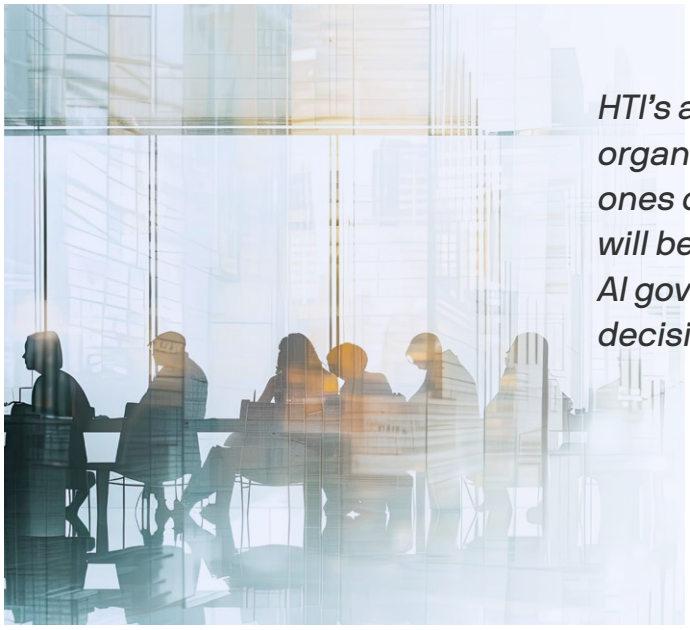

HTI's analysis suggests that many organisations, particularly large, complex ones or those with significant AI investments, will benefit from initially establishing dedicated AI governance structures with centralised decision-making authority

Introduction: what are AI governance structures and why they matter

AI governance structures are the formal mechanisms through which oversight of AI is exercised in an organisation. They determine who makes decisions about AI, who is accountable for outcomes, how concerns are escalated, and how AI activities connect to broader organisational oversight.

As artificial intelligence (AI) systems become deeply embedded in organisational operations, fit-for-purpose AI governance structures are essential to ensure that AI is used safely, responsibly and effectively.

Yet governance arrangements are often slow to evolve. This is creating a governance gap: organisations are adopting AI faster than they are adapting their governance.[1] This gap leaves organisations exposed: deploying customer-facing AI systems without clear accountability, or investing millions in AI initiatives without proper evaluation frameworks.

At the same time, many appear overconfident. Fifth Quadrant's Responsible AI Index 2025 found that 77% of organisations rated themselves capable of designing and building responsible AI systems, but only 25% had established clearly designated roles for AI responsibility.[2] Allocating responsibility for AI decision-making is a foundational step for AI governance; confidence alone is clearly insufficient. Governance structures are needed to ensure that organisations turn belief into practice.

These structures are more than administrative arrangements. They translate organisational strategy, values, and risk appetite in relation AI into clear practices, supporting the responsible adoption of AI while managing risks and opportunities. In this way, governance structures are a critical component in an organisation's broader AI governance operating model.

Importantly, AI governance structures are also enablers of value creation. Clear accountability, responsibility and decision-making provide the structure within which organisations can experiment and innovate with AI responsibly.

Our research shows that there is no one-size-fits-all model for effective AI governance. Organisations must tailor their governance structures to their context, particularly their size, sector, maturity, and strategic priorities. The success of these structures ultimately depends on how they are designed, built and maintained.

This snapshot examines two of the critical design choices that shape AI governance structures, the benefits and risks the different choices entail, and what is needed to build and maintain them effectively.

A critical component of the AI governance operating model

To understand how AI governance structures create value, it's helpful to see them within the an organisation's broader governance ecosystem.

AI governance structures are one of the four critical components that make up HTI's AI governance operating model (as set out in Figure 1). Each component plays a distinct role but works in concert to ensure AI is governed consistently and effectively.

Within this model, AI governance structures are the critical bridge between strategy and practices. They connect an organisation's overarching strategic direction and risk appetite to day-to-day oversight, by clarifying accountabilities and responsibilities, defining decision-making and escalation pathways, and establishing reporting lines.

Figure 1. HTI's AI Governance Operating Model

An AI governance operating model can support a structured and coordinated approach to adopting AI responsibly. An AI governance operating model has four components, each with critical aspects of effective AI governance

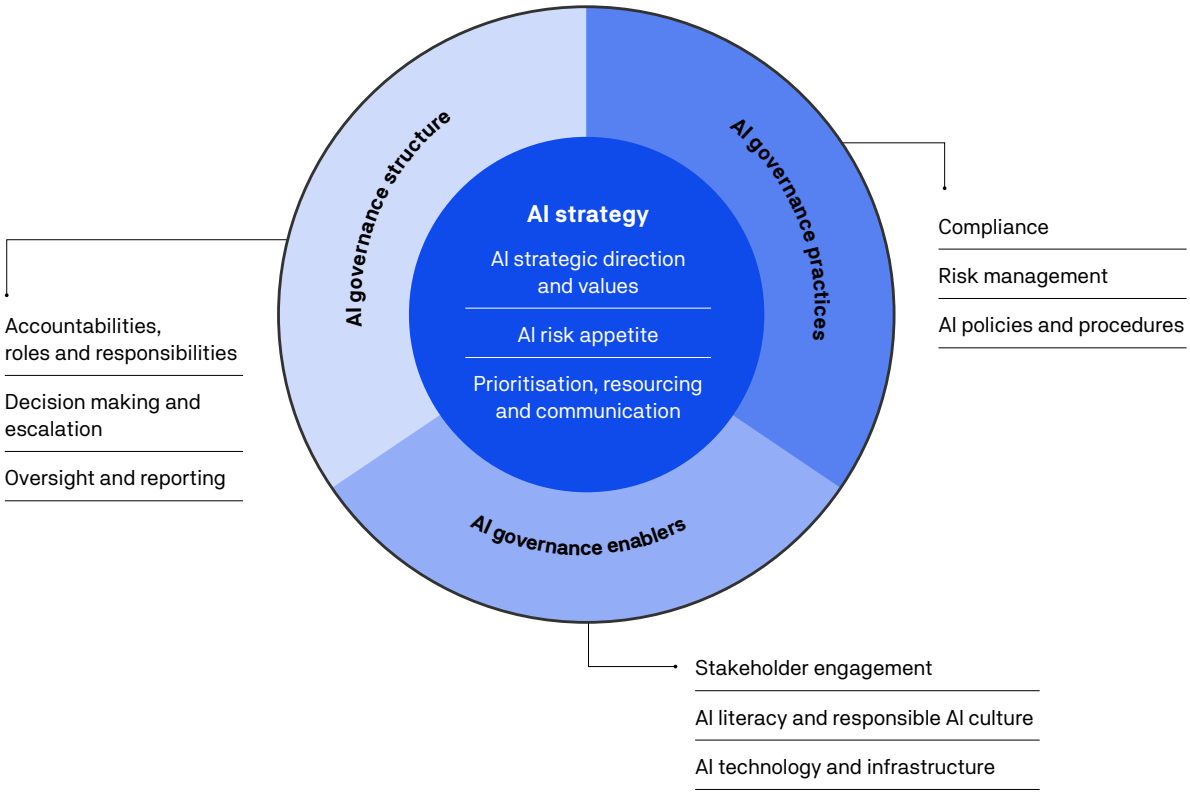

AI governance structures serve multiple critical functions, they:

Strategic alignment

- review high-risk AI use cases and determine whether to approve, escalate or reject, to ensure alignment with the organisation's strategic direction, values and risk appetite
- connect AI initiatives to strategic direction and business objectives.

Operational oversight

- allocate key roles and responsibilities in relation to the oversight, design, development and use of AI across the organisation
- oversee the practices, policies and procedures that support responsible AI use.

Accountability and trust

- define clear decision-making and escalation pathways
- report on AI use, risks and opportunities to senior executives or the board
- build stakeholder trust through transparency, engagement and accountability.[3]

Foundations for effective AI governance structures

To ensure the organisation's strategic direction is reflected in its practices, organisations need governance structures that are not only well-designed but also properly empowered, integrated into decision-making processes, and connected to broader organisational governance. Without this, they risk becoming symbolic rather than substantive, adding bureaucracy without meaningfully strengthening responsible AI practice.

HTI's analysis into leading Australian and international practices reveals three foundations for effective AI governance structures:[4]

- Clear accountability, roles, responsibilities, and reporting lines
- Meaningful integration of key teams across the organisation, supported by senior champions
- Formal decision-making authority supported by suitable frameworks and processes.

These foundations support the oversight, coordination, and risk management for safe and responsible AI use. However, governance structures must also be actively maintained through:

- Regular meetings of relevant governance bodies
- Monitoring, and reporting processes
- Meaningful stakeholder engagement.

To translate these foundations into practice, organisations should define who does what, by when, with what evidence, and how effectiveness will be measured.

To translate these foundations into practice, organisations should define who does what, by when, with what evidence, and how effectiveness will be measured.

The rise of AI governance structures

The rapid scaling of AI across organisations is driving a fundamental shift in AI governance: from informal, ad-hoc AI oversight to formal governance structures designed to make information flows and decision-making more systematic, consistent and reliable.[5] Two distinct approaches are emerging. Larger organisations with significant AI investment or more mature AI governance, like Telstra or Microsoft, have established dedicated governance bodies, such AI boards or committees.[6] Research suggests that organisations that have a *larger AI governance program* are more likely to have an AI governance committee.[7]

Alternatively, organisations that are smaller or less AI-mature often integrate AI oversight into existing structures, at least initially. This evolution was seen at the Australian Tax Office: it first relied on existing governance structures, including its Risk and Strategy Committees, to support its adoption of AI. However, in recognition that stronger governance arrangements were needed, it then also established a dedicated Data and Analytics Governance Committee to improve AI oversight.[8]

Across Australia and internationally, investment in governance structures is increasing but implementation remains uneven and more concentrated in leading organisations. Survey data indicate that one-third to one-half of respondents reported having some form of dedicated AI governance body, with adoption much higher among leaders (see Figure 2).

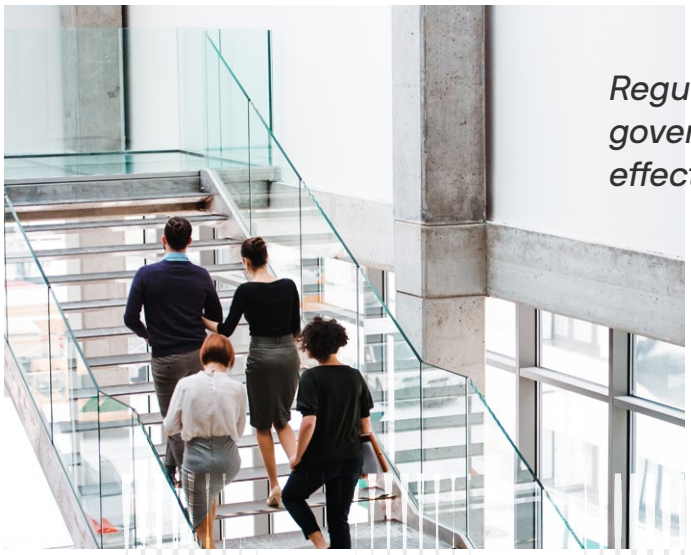

Regulators are highlighting that formal governance structures are essential for effective AI oversight.

Figure 2: Adoption of AI Governance Structures

55% of 1808 global respondents

said their organisation had implemented an AI board (Gartner, 2024).[9]

39% of 670 global respondents

reported having an AI governance committee,[10] rising to 49% where organisations were working on AI governance (IAPP, 2025).

33% of 418 Australian respondents

reported having an AI risk/governance committee, this jumps to 90% among the 50 leadings organisation in responsible AI (Responsible AI Index, 2025).

These adoption rates must be viewed in context. With 98% of Australian businesses being SMEs – who typically lack resources for dedicated AI committees – the economy-wide adoption rate is likely significantly lower than these enterprise-focused surveys suggest.[11]

HTI's engagement with industry leaders reflects the focus on AI governance structures. At HTI's *Shaping our Future Symposium* (January 2024), attendees identified governance structures as the element of AI governance receiving the second most attention.[12] In late 2024, HTI held a series of workshops with over 200 corporate leaders, which revealed that many organisations were struggling with whether to create new structures or adapt existing ones, having concerns about duplication, resource allocation, and organisational change management. Now, many organisations we talk to do have them.

Regulators are highlighting that formal governance structures are essential for effective AI oversight. In October 2024, the Australian Securities and Investments Commission (ASIC) published its review of 624 AI use cases across 23 Australian financial services (AFS) licensees (ASIC Report 798). ASIC considered those AFS licensees designated as 'most mature' as having demonstrated clear ownership that 'clear ownership and accountability for AI at an organisational level, including an AI-specific committee or council' and took a 'strategic, centralised approach' to AI governance.

ASIC's position represents a clear shift in regulatory expectations. For AFS licensees, AI governance is not a 'nice-to-have', with ASIC actively identifying dedicated, centralised AI governance structures as better practice and markers of maturity. As regulators increase their focus on AI systems and their governance, organisations without such structures may be subject to greater scrutiny.

Design choices for AI governance structures

Organisations face two fundamental choices when designing AI governance structures:

- **Form:** What form will the oversight mechanism take? More specifically, should a ***new, dedicated*** AI governance body be created or will ***existing*** governance structures be given oversight of AI?

- **Location:** Where in the organisation should authority for AI governance reside? Should governance decisions be ***centralised***, or ***distributed*** across different units and domains?

These choices are fundamental because they determine whether governance exists in practice, not just on paper. Form creates the *mechanism for oversight*, while location determines *where and who wields authority* and also influences *information flows and coordination patterns.* Together, these choices affect every major AI decision in your organisation. Without clarity on both choices, organisations risk falling into default or convenience-based structures that more closely resemble 'AI governance theatre' rather than genuine oversight.

There is not a single correct answer to these questions. Nor are the choices binary, as many organisations are adopting hybrid approaches. The right choice will depend on an organisation's AI maturity, the scale and complexity of AI use, and existing governance structures and capabilities.

Furthermore, structural decisions should not be set in stone. AI governance cannot be static: AI technologies are continually evolving, new use cases are emerging daily, and many commercial applications can be considered experimental, rather than mature. As an organisation matures and adoption deepens, its governance structures will inevitably need to adapt. Governance must therefore be an adaptive capability.

Rather than aiming for the 'perfect' governance structure, organisations should prioritise what is fit-for-purpose today, while planning for iterative improvement as technology and governance practices evolve. Organisations that lock in rigid structures risk being overtaken both by the pace of AI evolution and by more agile competitors. Leading organisations are building governance structures with change mechanisms built in from day one.

1. What form of AI governance structure should you choose?

Create a new, dedicated AI governance structure

Many organisations have established dedicated AI governance structures, often in the form of governance bodies, such as AI councils or committees, to provide specialised oversight and accountability for AI systems. These structures consolidate expertise, create defined decision-making pathways, and signal a visible commitment to responsible AI use to stakeholders.[15]

Dedicated governance bodies tend to have:

- **Formal mandate:** a written charter defining scope, terms of reference, decision rights, and escalation pathways.
- **Clear, unified entry points:** all AI use cases must go through a single intake and registration process to prevent systems eluding oversight.
- **Cross-functional composition:** membership drawn from across the business, so decisions are well-rounded.
- **Independent oversight backed by senior leadership:** chaired or sponsored by a C-suite executive, with separation from delivery teams and external expertise where needed.
- **Risk-based approval thresholds with standardised review:** a standard process is used to review and approve systems in line with their potential impact, with high-risk cases often escalated to senior decision-makers or the board.

There are clear benefits to such structures. The IAPP *AI Governance Profession Report 2025* reports that organisations with an AI governance committee encountered fewer issues when developing AI systems and fewer challenges in AI governance reporting.[14] By convening the right disciplines in one forum (e.g. business owners, risk, legal, procurement, IT/data, people & culture etc), dedicated bodies accelerate capability building and shorten the feedback loop between policy and practice. They also clarify information requirements for oversight, such as what risks and impacts to consider, what evidence is needed, and how proposed uses align to values and risk appetite.

It's important to note that 'dedicated' does not have to mean complex. A lean structure, such as appointing a single accountable executive supported by a small cross-functional working group or external advisors working with templated processes, can deliver dedicated governance benefits even in lower-resource contexts. This ensures clear accountability and visible commitment to responsible AI, while avoiding unnecessary overhead. It also makes this approach relevant for SMEs, most of which will not have dedicated risk, legal and procurement teams.

Importantly, these structures need not be permanent. They can serve as powerful transitional mechanisms, rapidly building capability and maturity during early AI adoption, then gradually transferring their functions into mainstream governance as the organisation matures.[15]

HTI recommends that most organisations, particularly those investing heavily in AI, establish dedicated AI governance bodies, even if only as a transitional arrangement. This approach provides fit-for-purpose oversight during early adoption while laying the groundwork to embed AI governance into mainstream risk and decision frameworks over time.

Building a dedicated AI governance body

The following questions can help guide organisations as they design an AI-focused governance body. Those questions marked with a star (*) should ideally be decided from the outset, but otherwise not every question needs to be answered on day one. The answers should change over time as the organisation learns and adapts.

Scope and function

- What are the objectives of the body? *
- Is this intended as a permanent or transitional, time-bound arrangement? *
- What is the scope of its oversight (e.g. what kinds of systems, where in the enterprise)? *

Accountability and decision rights

- How does this body connect with any existing governance frameworks (e.g. compliance teams, data governance, IT procurement, privacy reviews etc)? *
- Who is accountable for AI systems and what are the reporting lines to the board and executive? *
- What authority and decision rights does the body have (e.g. request information, approve or veto projects, set standards, escalate issues etc)? *

Composition and roles

- Who will chair the body (e.g. Lead AI Executive), and who are the core members? *
- How are members appointed, how many will there be and what are their terms? *
- Which key capabilities should be represented (e.g. procurement, legal, risk and compliance advisers, IT and data teams, people and culture staff, and stakeholder representatives)?
- How will diversity of perspectives be ensured (e.g. from staff, customers, community)?
- Should independent advisors or experts be included to provide challenge or external perspectives?
- How will the body be supported and resourced?

Policies and procedures

- Is there a formal mandate or charter outlining its scope, authority, reporting and escalation responsibilities? *
- Are review and approval criteria defined (privacy, security, fairness, transparency, performance, workforce/community impact, etc.)? *
- Is there an AI inventory that registers and tracks all relevant AI projects for oversight?
- What methods are used by the body to practically make decisions about projects?

Monitoring, evaluation, and improvement

- What is the cadence for project oversight and monitoring? *
- How will the body respond to and learn from AI incidents if they emerge?
- How will staff training and capability gaps be identified and addressed?
- What mechanisms enable adaptation to new risks, regulations, and organisational needs?
- When and how will the governance body be reviewed for effectiveness?

Rely on or upgrade existing governance structures

Rather than create new, dedicated AI governance bodies, some organisations deliberately choose to allocate accountability and responsibility for AI oversight to allocate within existing structures. This can be a strategic choice when:

- AI use is limited in scale and riskiness.
- Existing governance frameworks function well and are adaptable.
- Executive ownership is explicit and uncontested.
- Speed or resource constraints favour leveraging what already works.

In HTI's experience, a large proportion of organisations rely on existing structures by default, simply because no deliberate decisions have been made about how best to govern AI. This can create significant gaps if assessments ignore how AI can underperform or create novel, material risks.

However, choosing to strengthen and enhance current governance structures, can be a suitable and proportionate response.[16] Particularly for small and medium-sized organisations, limited staff, time or expertise makes creating a dedicated AI governance structure unrealistic. In these circumstances, responsibility and accountability for AI may simply rest directly with the CEO or another senior leader, and strengthening the processes that already go into helping them make decisions will be the most appropriate way to ensure AI governance is structured, consistent and practically achievable.

When adopted by design, this approach offers several advantages. It avoids duplication, leverages established reporting lines and embeds AI oversight directly into mainstream governance, ensuring that AI risks are managed alongside other enterprise risks rather than in isolation.

But it only works if existing structures are upgraded to deal with the 'AI delta' – the ways in which AI systems operate differently to traditional IT systems in ways that can amplify existing risks and introduce new ones. Processes must be reviewed and adapted accordingly, and relevant individuals and teams will need appropriate training and support.

A word of warning: if you are considering introducing AI systems that could have a significant impact on customers, employees or other stakeholders, if you operate in a regulated environment, or if AI will be deployed across core business functions, relying solely on existing structures is probably not enough. In these cases, you should establish at least a dedicated, transitional AI governance body to provide the focus, expertise, and independence for good decision-making.

Which existing governance structure should provide oversight for AI?

Where organisations rely on existing governance structures, primary responsibility for AI governance is typically assigned to a particular business function. According to the IAPP AI Governance Profession Report 2025, the functions most often tasked with this responsibility are: privacy (22%); legal and compliance (22%); IT (17%); and data governance (10%). Allocating responsibility for AI governance to these functions has different benefits and limitations:

Privacy

AI systems often amplify privacy risks and privacy teams are usually well-experienced in conducting risk assessments and privacy impact assessments. They are often well placed to identify and mitigate data-related harms. However, AI systems raise risks that extend beyond privacy, such as bias, transparency, and explainability, meaning privacy expertise must be supported by other capabilities.

Legal

Legal teams play a critical role in advising and ensuring compliance with existing obligations. The use of AI is subject to range of existing laws, including privacy, intellectual property, consumer protection, cyber security, anti-discrimination, negligence, and work health & safety. While lawyers bring expertise in law and regulation, they are often less familiar with AI technologies or governance frameworks and upskilling is usually required.

IT

IT teams often lead AI adoption, making them a natural early choice for governance.[17] However, traditional IT governance is not fit-for-purpose for AI. Unlike traditional software, AI can be dynamic, opaque, and non-deterministic. This creates distinct oversight challenges and potential conflicts of interest if the same team both implements and oversees AI.

Data

AI systems are heavily dependent on data so data governance is essential for effective AI governance. Yet while necessary, data governance alone is insufficient to address the full spectrum of AI risks and organisational impacts.[18]

However, even if one function has primary responsibility for AI governance, HTI's research strongly indicates that no single function can govern AI effectively in isolation: it requires an interdisciplinary approach.[19] Each domain should consider how its remit is affected by AI and respond accordingly, drawing on other teams as appropriate.

Choosing the form of the AI governance structure

Each form offers advantages and challenges for organisations, as set out below in Table 1. Organisations should select the form that delivers effective AI governance while best meeting their needs.

Table 1: Advantages and challenges of building dedicated AI governance bodies or relying on existing governance structures

Building dedicated AI governance bodies	Relying on existing governance structures
Tailored expert oversight: AI-specific governance bodies should have the skills and capability to address unique AI risks more effectively than general governance committees. Establishing a dedicated AI governance body means organisations build deep understanding and expertise in AI risk management and mitigation strategies. These governance bodies can provide early-stage guidance on AI use cases, helping business functions proactively identify and address risks before deployment.	**Leverages existing knowledge and processes:** Organisations can build upon existing governance mechanisms, reducing duplication and avoiding overlapping internal responsibilities and processes. This approach can leverage the expertise of existing teams. However, these teams may need to upskill to ensure they are across any new or amplified risks from AI systems.
Agility: Standalone AI bodies can be set up so that they are more responsive to AI's evolving risks and opportunities. This can help ensure timely risk mitigation and alignment with new regulatory requirements. However, to avoid delays and inefficiencies, these bodies must be able to effectively receive and provide information to and from existing structures.	**Not-fit-for-purpose:** Existing governance structures may be unsuited for AI governance as AI introduces amplified and emerging risks and considerations that may demand unique focus. Extending the remit of existing governance to address AI will require review and adaption of existing governance policies and processes, as well as upskilling and expanding the expertise of governance teams. There may also be governance gaps that emerge where issues fall outside existing functions.
Scalability: If a single AI governance body is responsible for reviewing and approving all AI use cases, this may create a bottleneck as AI adoption grows.	**Slow and complex:** Existing governance structures, particularly in larger organisations, may be slow, complex, and unsuited to address AI's rapid evolution and emergent risks, leading to slower decision-making, potential governance gaps, unidentified risks, and missed opportunities.

2. Where does authority reside for the AI governance structure?

In addition to deciding the *form* of AI governance structure, organisations should also consider *where authority sits*. This choice determines who gets to approve, veto, escalate and enforce policies related to AI system investment and use. It shapes risk control, speed, consistency, and alignment to strategy.

The key question here is whether authority over AI governance will be centralised (e.g. in a single body reporting to a senior executive) or decentralised across multiple business units or functions (e.g. each business unit decides for themselves or has their own committee).

Centralised authority

A centralised AI governance structure means there is a unified focal point for AI oversight, decision-making, and risk management.

Done well, it drives consistent standards and approvals, gives executives and the board line-of-sight across the AI portfolio, and makes it easier to coordinate scarce expertise and change efforts across the enterprise. These advantages are why centralisation is increasingly seen as better practice[20] in regulated contexts: for example, ASIC Report 798 found more mature financial licensees took 'strategic, centralised AI governance approaches' and generally 'demonstrated clear ownership and accountability for AI at an organisational level, including an AI-specific committee or council'.[21]

Centralisation also has its drawbacks. If all decisions are routed through a single body, it can quickly become a bottleneck, slowing approvals and frustrating innovation. For large, federated organisations, embedding a central authority can also be difficult where strong local risk and compliance processes already exist. These challenges can be managed through careful design: keeping central control over high-risk approvals, policy and standards, while delegating lower-risk decisions to business units within clear guardrails and agreed timeframes.

Where organisations choose to create a dedicated AI governance structure (as discussed in 5.1 above), HTI's view is that it should be generally supported by centralised decision-making authority. The ultimate owner should be a C-suite senior executive with lead responsibility for AI governance (what HTI describes as the Lead AI Executive[22]), such as the Chief Operations Officer, Chief Risk Officer, or Chief AI Officer.

Centralisation is typically most effective for organisations that:

- Operate in heavily regulated sectors, such as financial services
- Deploy AI across multiple business units
- Have significant data protection needs
- Use hierarchical structures and an existing centralised decision-making framework.

Decentralised authority

A decentralised or federated approach locates AI decision rights and governance responsibilities close to where AI systems are built, used and deliver value. In this model, business units (or teams) make their own investment and deployment decisions within reference to organisational policies, and are accountable for outcomes in their domains.

Done well, decentralisation of AI governance can better support different needs across an organisation, accelerate decisions, and surface practical risks early: the teams intending to deploy the AI systems can also be ones that understand their data, processes, and users best. This advantage can be enhanced through structured worker engagement to leverage the expertise and insights of employees closest to stakeholders.[23]

It is most effective where use cases are low-to-medium risk, where use cases are domain-specific (rather than enterprise-wide), and where senior leadership has confidence in the consistency of local risk management. It also suits federated operating models in which accountability for outcomes is firmly owned in the line of business.

However, without strong coordination, decentralisation can result in inconsistent standards, duplicated effort, and drift from established strategy and risk appetite. In Report 798, ASIC found that less mature licensees 'generally adopted decentralised approaches that leveraged existing frameworks'.[24] The effectiveness of their governance depended, in part, on the robustness of existing governance and risk management arrangements.[25] If the organisation cannot rely on consistent local discipline, or if the risks are significant, a decentralised model on its own is likely to be insufficient.

Without strong coordination, decentralisation can result in inconsistent standards, duplicated effort, and drift from established strategy and risk appetite.

Hybrid approaches

A third option is to deploy a hybrid approach that combines elements of both centralised and decentralised authority. A hybrid approach can balance strategic consistency with operational flexibility, making it particularly valuable in large organisations where AI is used in different ways across multiple business units.[26]

Under a hybrid approach, authority and decision-making is distributed based on risk, scale, and context. A central AI governance body tends to set enterprise-wide policies, strategies and minimum standards and retains the authority to approve high-risk use cases, while business units enjoy delegated authority to manage low-risk or operational AI deployments, provided they operate within approved guardrails.

A common way to operationalise this is a hub-and-spoke design. An AI Centre of Excellence (the hub) centralises expertise, develops templates and frameworks, supports teams with advice and reviews, and monitors the portfolio for trends, critical incidents and strategic drift.[27] Meanwhile, individual business units are responsible for domain-specific innovation (the spokes), completing standard review processes, escalate higher-risk uses or proposals to the centre for advice, challenge and approval.[28]

In practice, a hybrid approach often involves:

First, centralised policies and information flows to support coordination: a single intake and register for all AI initiatives to provide portfolio visibility; common risk-tiering and review criteria; and a shared set of minimum controls (privacy and security, safety/robustness, fairness, transparency, legal, and workforce/community impact). These help local teams to make consistent, evidence-based decisions that draw on the thinking and effort of other business units.

Second, clear delegation and escalation rules: generally, decentralised approaches empower business units to approve low and medium-risk use cases within defined parameters, while high-risk proposals are escalated to a central body or executive sponsor for separate challenge and approval. The triggers for this kind of escalation must be explicit, well understood and verifiable (e.g. potential for serious customer harm, use of sensitive data in an AI system, or automation of significant decisions).

Third, independence and assurance: particularly in regulated entities or where risk management is a priority, delivery teams should not be the sole approvers of their own systems. Independent challenge and review (from risk, legal, internal audit, or a qualified external advisor) helps counter optimism bias and “marking our own homework”.

Fourth, structured stakeholder engagement processes: a strength of decentralised approaches is that worker, customer and community perspectives can be included in governance decisions by design. Structured input channels, such as workforce impact assessments or targeted user consultations, can help surface risks while also ensuring systems deliver their intended value in a trustworthy way.

Choosing where authority resides for the AI governance structure

Each approach has advantages and challenges, as outlined below. Ultimately, in HTI's view, whilst there is value in allowing business units to develop and experiment with AI use cases, it is important to have consistency in approval of high-risk AI use cases, which is best achieved through a centralised process.

Table 2: Advantages and challenges of centralised, decentralised, and hybrid AI governance structures

Centralised	Decentralised	Hybrid
Standardisation and strategic alignment: Ensures consistent application of AI policies and controls across departments. Consolidates expertise and provides enterprise-level oversight to align AI activities with strategic priorities. **Efficient but potentially slow:** Centralised review can avoid duplications and leverages cross-functional expertise, but can create bottlenecks as AI use scales. **One-size-fits-all:** A single body may lack deep contextual knowledge of each business unit, leading to decisions that overlook operational needs and domain-specific risk.	**Tailored and responsive:** Business units can develop AI solutions that directly address their needs and adapt quickly to local operational contexts. Business units often have deeper knowledge of domain-specific risks, data, and constraints. Hence, local teams may be best placed to identify risks early and take steps to mitigate them accordingly. **Risk of inconsistency:** A lack of strong coordination may result in governance standards and policies which vary across units or jurisdictions. **Lack of enterprise accountability**: Without senior-level ownership, decentralised approaches may lack executive visibility and oversight.	**Balanced governance:** Distributes authority based on risk or context, combining central oversight for higher-risk use cases with flexibility for local experimentation. **Reduces bottlenecks:** Operational units can act within clear guardrails, with approval by central bodies only required for higher-risk use cases. **Requires effective coordination:** To avoid fragmentation, hybrid models must ensure shared standards, clear roles, and good communication between central bodies and operational units.

Informal governance structures

This snapshot has focused on the development of formal AI governance structures. However, informal structures also play a vital role in effective AI governance. While they lack formal authority or decision-making powers, they can provide advice support innovation, and foster cultural change. Examples include working groups, internal or external advisory networks, discussion boards, and communities of practice.

Compared to formal processes, these mechanisms can offer greater flexibility and responsiveness, lower barriers to engagement, and richer operational insights. Organisations often benefit from embedding both formal oversight for high-risk or strategic decisions, alongside informal mechanisms to build capability, provide advice, and support implementation.

For example, AstraZeneca has implemented two complementary structures: an AI Resolution Board and Responsible AI Consultancy Service. These structures were created to support the sharing of best practices and to educate staff about the risks of AI systems.

The AI resolution board provides formal governance oversight by reviewing and approving high-risk AI use cases. The Responsible AI Consultancy Service conversely acts as an informal advisory body with three key objectives:

- provide ethical guidance
- support the practical embedding of AI ethics principles
- overseeing AI project governance.[29]

This dual structure illustrates how formal and informal mechanisms work best in tandem. The Consultancy Service helps teams seeking guidance and fosters a culture of responsible AI development without immediately imposing formal review processes, while the AI Resolution Board ensures appropriate, formal scrutiny for higher-risk applications. Together, they support a culture of responsible AI development and enable risk-proportionate governance across the organisation.

Conclusion

AI governance structures are not just administrative formalities. They are a critical component of an organsations AI governance operating model and the levers through which boards and executives exercise real oversight and enable safe innovation. They provide the accountability and decision-making discipline that ensures AI adoption aligns with strategy, risk appetite, and values.

As AI adoption accelerates and regulatory scrutiny increases, boards cannot rely purely on informal or inconsistent arrangements. Boards must satisfy themselves that clear and effective structures exist that include defined roles, cross-functional expertise, senior sponsorship, and the authority to make binding decisions. And those structures must be active: meeting regularly, monitoring systems in production, escalating risks, and learning as the organisation matures.

Organisations that take a deliberate, forward-looking approach to AI governance will not just avoid missteps, but also unlock AI's value with confidence and trust.[30] Ultimately, responsibility for AI governance sits with the board. By establishing and resourcing robust governance structures, boards can ensure their organisations are not simply compliant, but strategically positioned to innovate safely, maintain stakeholder trust, and compete in an AI-enabled future.

Organisations that take a deliberate, forward-looking approach to AI governance will not just avoid missteps, but also unlock AI's value, with confidence and trust.[32]

Learn more about AI governance with HTI

The AI Corporate Governance Program is an initiative of the UTS Human Technology Institute (HTI). Its aim is to broaden understanding of corporate accountability and governance in the use of AI. HTI's AI Corporate Governance Program analyses current and emerging AI governance practices and provides organisations with the resources and strategic insight to navigate this developing terrain.

HTI's AI Corporate Governance Program is supported by philanthropic partner Minderoo Foundation, and project advisory partners KPMG, Gilbert + Tobin, and Atlassian.

For more information, or to join our AI Governance mailing list, please contact our AI Governance Team:

Professor Nicholas Davis
Industry Professor, Emerging Technology and HTI Co-Director
nicholas.davis@uts.edu.au

Gaby Carney
Senior Fellow, Strategic AI
gaby.carney@uts.edu.au

Llewellyn Spink
AI Corporate Governance Lead
llewellyn.spink@uts.edu.au

Myfanwy Wallwork
AI Governance Lead
myfanwy.wallwork@uts.edu.au

Jack Goldsmith
AI Corporate Governance Specialist
jack.goldsmith@uts.edu.au

Endnotes

1. Australian Securities and Investments Commission, *Report 798 Beware the Gap: Governance Arrangements in the Face of AI Innovation* (Report, 29 October 2024) 3 <https://download.asic.gov.au/media/mtllqjo0/rep-798-published-29-october-2024.pdf>.

2. Fifth Quadrant , *Australian Responsible AI Index 2025* (Report, 26 August 2025) <https://www.fifthquadrant.com.au/content/uploads/Australian-Responsible-AI-Index-2025_Full-Report.pdf>

3. Allens, AI Governance Toolkit: A Guide for Boards and General Counsel (Report, 2024) <https://www.allens.com.au/globalassets/pdfs/insights/data-privacy/ai-governance-toolkit.pdf>; Australian Institute of Company Directors & Human Technology Institute, *A Director's Guide to AI Governance* (Report, June 2024) <https://www.aicd.com.au/content/dam/aicd/pdf/tools-resources/director-resources/a-directors-guide-to-ai-governance-web.pdf>; Human Technology Institute, *AI Governance Snapshot: Essential Components of AI Governance* (Report, January 2024) <https://www.uts.edu.au/globalassets/sites/default/files/2024-01/ai-governance-snapshot---essential-components-of-ai-governance.pdf>

4. For example, see: Australian Securities and Investments Commission, *REP 798 Beware the Gap: Governance Arrangements in the Face of AI Innovation* (Report, 29 October 2024) <https://download.asic.gov.au/media/mtllqjo0/rep-798-published-29-october-2024.pdf>; Info-communications Media Development Authority and Personal Data Protection Commission, *Model AI Governance Framework Second Edition* (Report, 21 January 2020) 22 <https://www.pdpc.gov.sg/-/media/files/pdpc/pdf-files/resource-for-organisation/ai/sgmodelaigovframework2.pdf>; Centre for Information Policy Leadership, *Building Accountable AI Programs: Mapping Emerging Best Practices to the CIPL Accountability Framework* (Report, February 2024) <https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_building_accountable_ai_programs_feb24.pdf>.

5. Centre for Information Policy Leadership, *Building Accountable AI Programs: Mapping Emerging Best Practices to the CIPL Accountability Framework* (Report, February 2024) 14 <https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_building_accountable_ai_programs_feb24.pdf>.

6. Human Technology Institute, *AI Governance Snapshot: Essential Components of AI Governance* (Report, January 2024) <https://www.uts.edu.au/sites/default/files/2024-01/AI%20Governance%20Snapshot%20-%20Essential%20Components%20of%20AI%20Governance.pdf>.

7. IAPP, *AI Governance Profession Report 2025* (Report, 2025) 17 <https://iapp.org/media/pdf/resource_center/ai_governance_profession_report_2025.pdf>.

8. Australian National Audit Office, *Governance of Artificial Intelligence at the Australian Taxation Office* (Auditor-General's Report No 26, 24 February 2025) <https://www.anao.gov.au/work/performance-audit/governance-of-artificial-intelligence-the-australian-taxation-office>.

9. Gartner, *Gartner Poll Finds 55% of Organisations Have an AI Board* (Press Release, 26 June 2024) <https://www.gartner.com/en/newsroom/press-releases/2024-06-26-gartner-poll-finds-55-percent-of-organizations-have-an-ai-board>.

10. IAPP, *AI Governance Profession Report 2025* (Report, 2025) <https://iapp.org/media/pdf/resource_center/ai_governance_profession_report_2025.pdf>.

11. Australian Small Business and Family Enterprise Ombudsman, *Number of Small Businesses in Australia* (Small Business Data Portal, August 2024) <https://www.asbfeo.gov.au/small-business-data-portal/number-small-businesses-australia>.

12. Human Technology Institute, *Shaping Our Future Symposium: Breakout Activities* (News Item, 28 February 2024) <https://www.uts.edu.au/human-technology-institute/news/shaping-our-future-symposium-breakout-activities>.

13. Centre for Information Policy Leadership, *Building Accountable AI Programs: Mapping Emerging Best Practices to the CIPL Accountability Framework* (Report, February 2024) <https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_building_accountable_ai_programs_feb24.pdf>.

14. IAPP, *AI Governance Profession Report 2025* (Report, 2025) <https://iapp.org/media/pdf/resource_center/ai_governance_profession_report_2025.pdf>.

15. Artificial Intelligence/Machine Learning Risk & Security Working Group, *Artificial Intelligence Risk & Governance* (Report, 11 January 2023) 12 <https://ai.wharton.upenn.edu/white-paper/artificial-intelligence-risk-governance/>.

16. Centre for Information Policy Leadership, *Building Accountable AI Programs: Mapping Emerging Best Practices to the CIPL Accountability Framework* (Report, February 2024) <https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_building_accountable_ai_programs_feb24.pdf>.

17. Centre for Information Policy Leadership, *Building Accountable AI Programs: Mapping Emerging Best Practices to the CIPL Accountability Framework* (Report, February 2024) 9; Matti Mäntymäki et al, 'Defining organizational AI Governance' (2022) 2 *AI and Ethics* 603, 606 <https://link.springer.com/article/10.1007/s43681-022-00143-x>.

18. Matti Mäntymäki et al, 'Defining organizational AI Governance' (2022) 2 *AI and Ethics* 603, 607 <https://link.springer.com/article/10.1007/s43681-022-00143-x>.

19. For further information on the specific people and skills needed for effective AI Governance, see: Human Technology Institute, *People, Skills and Culture for Effective AI Governance* (Report, December 2024) <https://www.uts.edu.au/globalassets/sites/default/files/2024-12/HTI-Snapshot-4-People-Skills-Culture-Dec.pdf>.

20. Centre for Information Policy Leadership, *Building Accountable AI Programs: Mapping Emerging Best Practices to the CIPL Accountability Framework* (Report, February 2024) 38 <https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_building_accountable_ai_programs_feb24.pdf>.

21. Australian Securities and Investments Commission, *REP 798 Beware the Gap: Governance Arrangements in the Face of AI Innovation* (Report, 29 October 2024) 23 <https://download.asic.gov.au/media/mtllqjo0/rep-798-published-29-october-2024.pdf>.

22. For further information on the role and responsibilities of the Lead AI executive, see: Human Technology Institute, *People, Skills and Culture for Effective AI Governance* (Report, December 2024) <https://www.uts.edu.au/globalassets/sites/default/files/2024-12/HTI-Snapshot-4-People-Skills-Culture-Dec.pdf>.

23. HTI's research into the impact of AI and automation on workers highlights that many workers feel like 'invisible bystanders' in AI adoption, with solutions imposed upon them without meaningful consultation. However, when properly engaged, workers provide valuable insights into ethical, legal, and operational risks of AI systems (Human Technology Institute & Essential Research, *Invisible Bystanders: How Australian Workers Experience the Uptake of AI and Automation* (Report, University of Technology Sydney, May 2024) <https://www.uts.edu.au/globalassets/shared-media/documents/CSJI/essentialresearchuts_invisible_bystanders_0524.pdf>).

24. Australian Securities and Investments Commission, *REP 798 Beware the Gap: Governance Arrangements in the Face of AI Innovation* (Report, 29 October 2024) 25 <https://download.asic.gov.au/media/mtllqjo0/rep-798-published-29-october-2024.pdf>.

25. Australian Securities and Investments Commission, *REP 798 Beware the Gap: Governance Arrangements in the Face of AI Innovation* (Report, 29 October 2024) 25 <https://download.asic.gov.au/media/mtllqjo0/rep-798-published-29-october-2024.pdf>.

26. Centre for Information Policy Leadership, *Building Accountable AI Programs: Mapping Emerging Best Practices to the CIPL Accountability Framework* (Report, February 2024) <https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_building_accountable_ai_programs_feb24.pdf>.

27. Kelly Gonzalez and Jason Milnes, 'Building an AI Centre of Excellence: what you need to know', *Bell Business* (Blog Post, 28 October 2024) <https://explore.business.bell.ca/blog/building-an-ai-coe>.

28. Masood, Adnan, 'The AI CoE Imperative — How to Build a Center of Excellence that Turns Experimental Bets into Scalable, Strategic Assets', *Medium* (online, 10 July 2025) <https://medium.com/@adnanmasood/the-ai-coe-imperative-how-to-build-an-ai-center-of-excellence-that-turns-experimental-bets-into-a4e0018aa84f>.

29. Pekka Pohja et al, 'Explainable Artificial Intelligence: A Systematic Review of Terminology, Methods and Metrics' (2022) 4 *Frontiers in Computer Science* <https://www.frontiersin.org/journals/computer-science/articles/10.3389/fcomp.2022.1068361/full>.

30. Allens, AI Governance Toolkit: A Guide for Boards and General Counsel (Report, 2024) < https://www.allens.com.au/globalassets/pdfs/insights/data-privacy/ai-governance-toolkit.pdf>.

Human Technology
Institute

For more information

Human Technology Institute
hti@uts.edu.au

University of Technology Sydney
PO Box 123
Broadway NSW 2007

uts.edu.au

UTS CRICOS 00099F