A Director's Guide to AI Governance

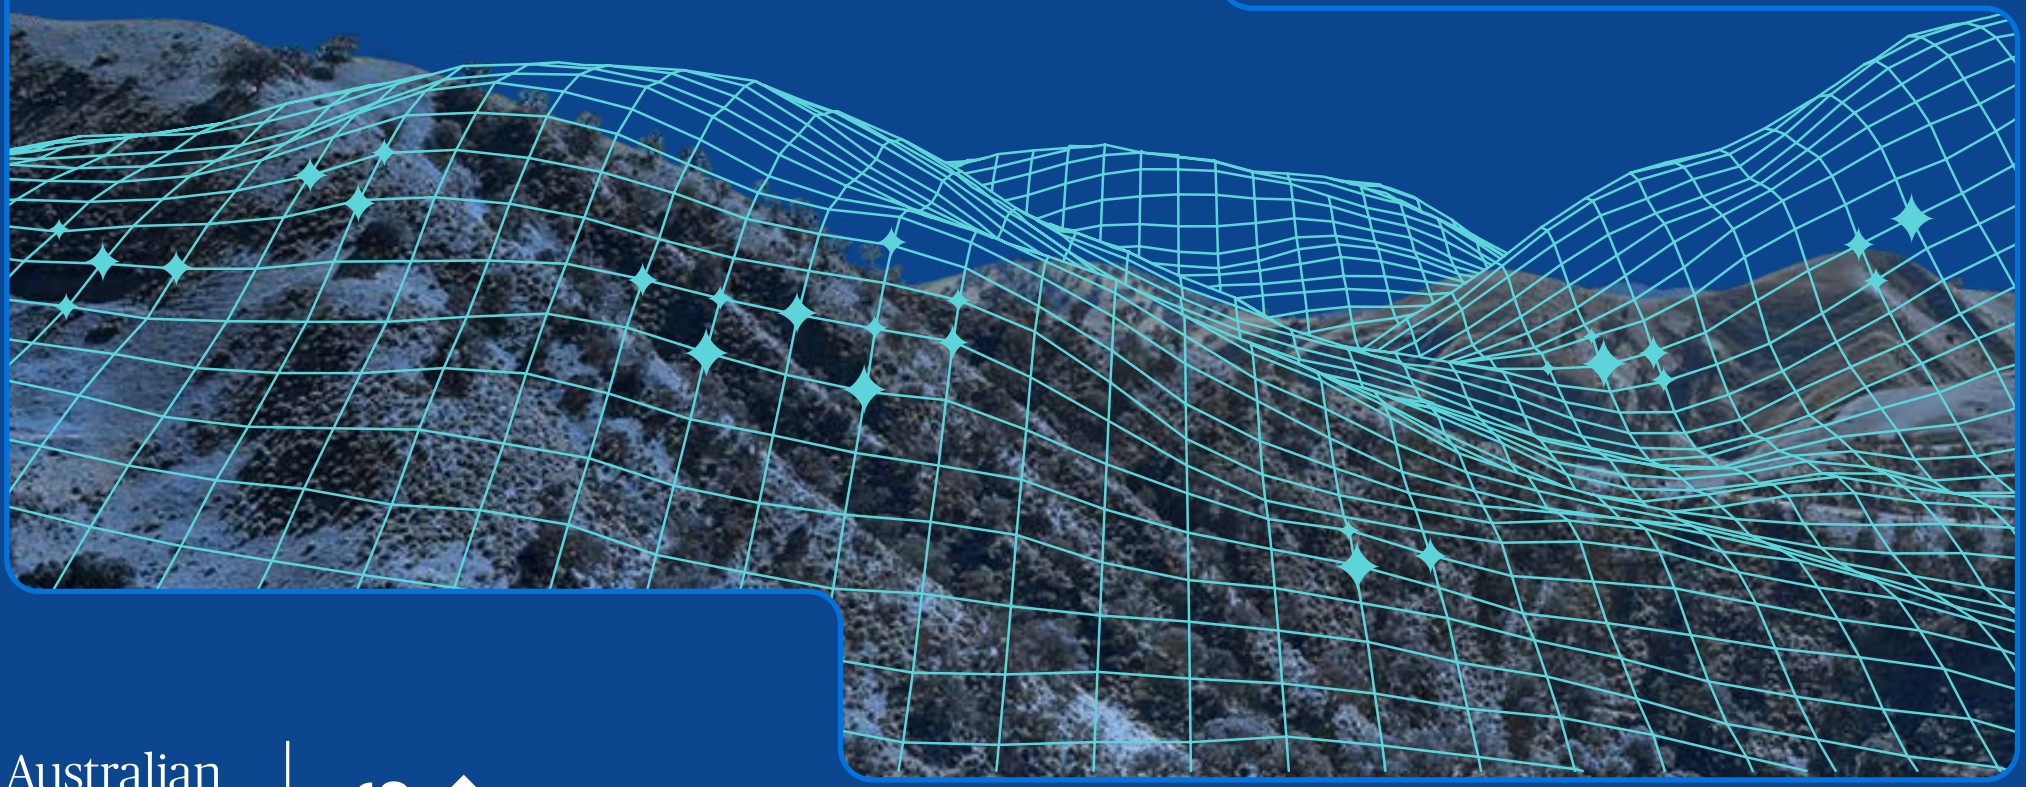

Australian Institute of Company Directors

UTS Human Technology Institute

Version 2 | June 2026

Contents

Case studies

Ministerial Foreword

Senator the Hon Tim Ayres

New technologies always catalyse change.

Steam-powered machines, internal combustion engines and industrial automation transformed the way we manufacture goods. Rail and steam transport revolutionised logistics. Electricity connection changed lives at home and at work. High-tech capabilities and the digital economy – unimaginable to thousands of previous generations – rewrote the nature of social, economic and industrial life.

All of those waves of technological progress carried risks. Not only society-level risks, but also harms to workers and communities, labour market disruption and environmental hazard. But the moment of their arrival posed another risk, silent yet equally dangerous: the risk of being left behind. It was the risk of missing out on investment and failing to capture the benefits of new technologies.

Firms that chose not to grapple with those parallel risks atrophied and failed. Countries that did the same – often constrained in their adaptive capacity by extractive colonial regimes and isolation from technological centres of gravity – fell further behind in economic, social justice, democratic and strategic terms.

Artificial intelligence (AI) – another general-purpose technology – is as significant as the steam engine, electricity or the digital economy.

Australians have always had a pragmatic approach to new technology. Business Council of Australia research demonstrates that innate Australian pragmatism is the lens through which most Australians see AI. Sceptical about the risks and especially apprehensive about what the technology means for work and jobs, Australians are nonetheless big adopters of AI at home and in the workplace.

The question is not whether Australians will use AI, but how, and in whose interest. All of us in positions of leadership – be they corporate or public sector roles – have a responsibility to deliver clarity and guidance. Both are pre-requisites of public trust in AI and the institutions responsible for stewarding its uptake. The Albanese Labor Government's National AI Plan – and Data Centre Expectations – is how we are clarifying the work we need to do together to get the most out of AI in Australia's interest.

The National AI Plan is all about capturing the opportunities that AI presents Australians and sharing those benefits broadly across the economy. AI should make life better for small and medium as well as large firms, low- and middle-income workers as well as professionals in office jobs, and regional communities as well as capital cities.

The Plan also articulates how we will work to keep Australians safe by updating legislative frameworks to keep regulation responsive and up to date. The Government's new AI Safety Institute is lifting the technical expertise and preparedness of agencies and regulators across government. Informed by experts at home and trusted AI safety scientists abroad, the Institute tests and monitors emerging AI capabilities to make sure that existing Australian laws – covering things like personal data, children's safety, online platform accountability, automated decision-making and workplace safety – remain fit for purpose.

Rapid technological change amplifies the need for thoughtful, adept, effective corporate leadership at the board level. Much has changed in policy terms, and in

technological terms, since the original *A Director's Guide to AI Governance* was released in 2024. This updated guide from the AICD and the Human Technology Institute (HTI) at the University of Technology Sydney grapples with organisational capability in the AI age.

The guidance emphasises the importance of human oversight. Boards are ultimately accountable for the adoption, implementation and outcomes of AI technologies in their organisations. Just as the Institute is increasing government capability and responsiveness, company boards will need to keep pace with change and routinely evaluate AI's use, its rewards and risks. That may involve regular AI stocktakes, workforce skills benchmarking, role mapping, technological monitoring and testing, workshops, and simulations to model future adaptation scenarios.

Boards that ask the right questions, set clear accountabilities and oversee effective risk controls will position their organisations to harness AI successfully. They will also earn the trust of their workers, customers and supply chains, which in turn drives better organisational performance. So many Australians are surveying the potential impacts of AI on the future of their occupations. Good corporate leadership is what must distinguish this wave of technological change from the unequal, sometimes inhumane, approach of industrialists in the age of steam and iron.

Boards will also shape AI adoption through the power of their example. Corporate adoption of analytical AI has been uneven and personal in practice, but boards will increasingly need to work with management to develop clear and transparent guidelines for the use of analytical, generative and agentic AI in their own deliberations and decision-making processes.

The AICD and HTI's updated, practical, accessible guidance for boards comes at a time when many are navigating AI adoption for the first time or establishing how to stay level with a technology that changes daily.

Much like the Government's National AI Centre's Guidance on AI Adoption, this document emphasises accountability, planning, risk measurement and management, transparency, testing and human control. These principles will stand Australian boardrooms in good stead to engage with AI in a way that is thoughtful, proportionate and forward looking.

Senator the Hon Tim Ayres
Minister for Industry and Innovation
Minister for Science

Partners' Foreword

AI is rapidly reshaping the way all Australian organisations operate, innovate and engage with employees, customers and clients.

Since the AICD and HTI first collaborated on AI governance resources in 2024, both the pace of technological development and the breadth of AI adoption in Australia have accelerated significantly. AI is no longer an experimental or uncertain technology. Underpinned by significant investments, it is increasingly embedded in core business processes, decision-making and customer engagement.

This updated resource reflects that changing reality and the implications for effective governance. It responds to strong demand from directors for practical, board-level guidance that keeps pace with both technological advancement and evolving regulatory and community expectations. The AICD and HTI's objective is to equip boards with the insight needed to navigate AI with confidence, balancing innovation with responsibility, including an awareness of the human impact.

AI presents profound opportunities for Australian organisations to enhance productivity, improve decision-making and deliver innovative products and services. At the same time, it introduces complex and often interconnected risks, including privacy, cyber security, bias and organisational culture. These risks are not static. They evolve alongside the technology, and often at a faster pace than traditional governance frameworks are designed to address.

For boards, this creates a clear imperative. Directors should have confidence that their organisations are not only adopting AI effectively but doing so in a way that aligns with their strategy, risk appetite and organisational values. This includes understanding its impact on employees, customers and broader stakeholders, and overseeing appropriate controls.

Importantly, AI governance intersects with existing board responsibilities covering cyber security, data governance, risk management and compliance with directors' duties. Boards do not need to become technical experts, but they do need to ask the right questions, seek appropriate assurance and foster a culture that supports responsible AI use.

This guide has been developed with those principles in mind. It is designed to be practical and accessible, providing a clear framework for directors to understand AI, identify key risks and opportunities, and embed effective governance practices within their organisations. It also reflects insights from extensive consultation with directors, regulators, academics and industry experts.

As AI continues to evolve, so too must governance practices. There is no one-size-fits-all approach, and boards will need to adapt their oversight to the specific context, complexity and maturity of their organisation's AI use. What remains constant, however, is the central role of the board in effective AI oversight.

We encourage directors to use this publication as a starting point to engage deeply with the issues, challenge assumptions and support their organisations to harness AI in a way that is both innovative and responsible.

Mark Rigotti FAICD
Managing Director & CEO
Australian Institute of Company Directors (AICD)

Professor Nicholas Davis
Industry Professor, Emerging Technology & Co-Director
University of Technology Sydney Human Technology Institute (HTI)

Snapshot

AI holds profound potential to reshape how all Australian organisations structure operations, innovate and deliver value to customers. It also brings unique and elevated risks. Central to AI adoption and impact is the human element, encompassing employees, customers and the broader community.

Boards play a key role in how Australian organisations balance the opportunity-risk dynamic of this transformative technology.

This publication presents governance guidance on AI in three parts:

PART 1 - AI Governance Foundations

Boards should be aware of the broader context in which AI investments are being considered, including rapid advances in technology, key relevant regulations and the opportunities and risks that are present in any organisational use of AI.

Sections	Key points for boards
AI is a transformative technology	• Developments in AI technology are advancing at an unprecedented rate with agentic AI at the forefront of its potential to transform organisations. • AI capabilities are increasingly being embedded within third-party products, while employees may also be using AI tools on a 'shadow' basis, without formal oversight. • AI governance is interdependent with board oversight of cyber security and data governance.
The role of the board and the regulatory landscape	• Boards have a key role in overseeing how an organisation uses AI, balancing the opportunities and risks of this transformative technology. • Board-level oversight of AI and its associated risks forms part of directors' existing duties. • Organisations are expected to govern AI within existing legal and regulatory frameworks, including obligations relating to privacy, consumer protection, discrimination, work health and safety, copyright and cyber security.
AI opportunities and risks	• AI systems present a range of opportunities for organisations, including improvements in productivity, product quality, customer service and employee experience. • At the same time, AI systems can introduce new commercial, reputational and regulatory risks, while also amplifying existing risks, including those related to cyber security and privacy. • AI has the potential to have significant human impacts, including on employment, and boards should maintain close oversight of these impacts.

PART 2 - AI Governance Operating Model

The AI Governance Operating Model developed by HTI provides a structured framework for board oversight across strategy, governance structures, governance practices, including risk management, and organisational enablers.

AI governance element	Key points for boards
Strategy	• The board should approach decision-making on AI and AI-related investments with discipline and have confidence that decisions align with the organisation's strategy, values and purpose. • Given the distinctive risks associated with AI, it is appropriate for the board to work with management to establish a risk appetite for AI that is ultimately reflected in internal processes and policies, including the risk management framework. • Effective AI implementation cannot occur in a vacuum. The board should have a clear view on the resourcing, data quality and prioritisation required to support AI adoption.
Structure	• The board has a key role in overseeing the use of AI in an organisation and should work with management to map AI-related responsibilities across the organisation. • Dedicated AI governance structures and bodies may not be necessary for every organisation. What is critical is that the board has visibility of how AI is managed, monitored and governed, underpinned by sufficient board AI literacy. • Roles and responsibilities for AI should be regularly reviewed and updated consistent with changing business operations, AI use cases and technological developments.
Practices	• AI can introduce unique risks, and elevate existing risks, for organisations. The board should oversee how these risks are appropriately identified and managed through existing risk management frameworks and controls. • The board should have visibility over how AI, cyber security and data governance risk controls intersect, and check that existing cyber controls (e.g. rapid patching) are also applied within AI environments. • External vendors of AI systems and supporting infrastructure are central to most AI systems and tools. The board should have visibility over the risks associated with these providers and their data protection settings.
Enablers	• A pre-condition of significant organisational investments in AI is enhancing the underlying digital, data and technology infrastructure. • An organisational culture that embraces AI knowledge, literacy and responsible use is critical to successful AI implementation and adoption. • Organisations need to be alive to the impact of AI on both internal and external stakeholders, with the board playing a particularly key role in overseeing the impact on employees and the associated support structures.

PART 3 - Measuring AI Returns

As organisations rush to adopt AI, often with significant investments, the board has a key role in understanding whether AI is driving genuine returns and organisational value.

Sections	Key points for boards
Measuring value and return from AI	• Greater value is realised where AI is used to transform processes, systems or business models, yet relatively few organisations have moved beyond the pilot phase in these strategies. • AI return on investment (ROI) assessments should focus on specific, measurable use cases and test if the investment has realised in terms of organisational efficiency and improved outputs.

Part 1 - AI Governance Foundations

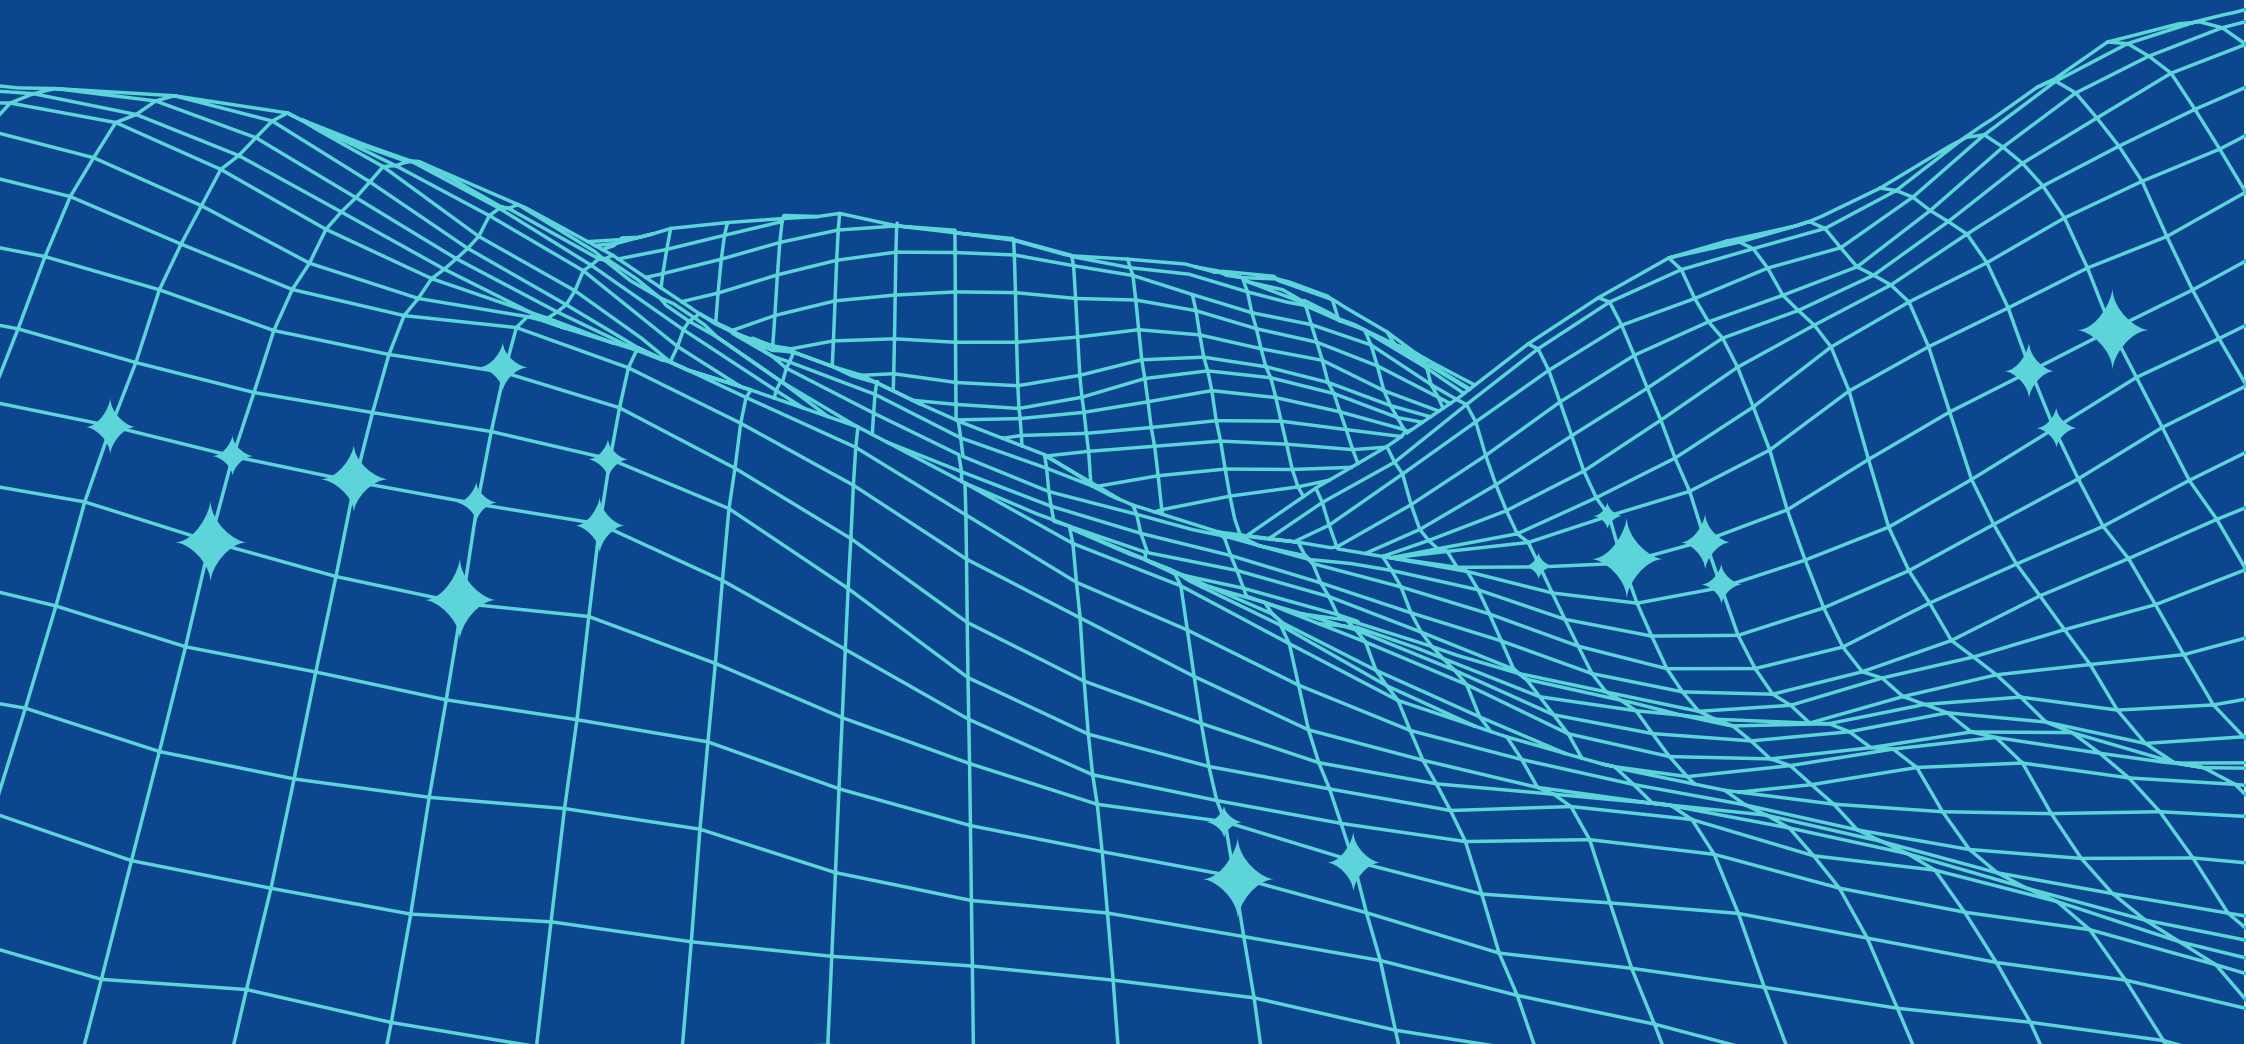

AI is a transformative technology

As a general-purpose technology, AI is increasingly reshaping how Australian organisations operate, compete and create value. It is both a significant transformational opportunity and a material source of risk.

Key points

- Developments in AI technology are advancing at an unprecedented rate with agentic AI at the forefront of its potential to transform organisations.
- AI capabilities are increasingly being embedded within third-party products, while employees may also be using AI tools on a 'shadow' basis, without formal oversight.
- AI governance is interdependent with board oversight of cyber security and data governance.

AI systems[1] and tools hold great potential to drive growth, productivity and innovation across Australian organisations, including through more tailored products and services. They are also a force that is resetting organisational practices, employee roles and customer behaviour. Organisations that move decisively on AI may gain a competitive advantage, while those that delay risk being left behind.

However, AI can also introduce new risks and amplify existing ones, including cyber security risks. Boards play a key role in balancing this dynamic between opportunity and risk, including navigating issues such as fairness, accountability, and the environmental and human impacts of AI.

Given the pace of technology change, it is critical that boards stay abreast of AI developments, continue to build their own AI literacy, and be conscious of how future advances may affect their organisations.

How is AI different from other technology?

AI is a form of software that is particularly good at pattern recognition, classifying information, making predictions based on those patterns, generating outputs and, increasingly, performing tasks with a degree of autonomy.

AI systems can provide significant productivity, efficiency and customer experience benefits, and often outperform other types of technology. AI systems are also highly versatile and scalable, as they can be adapted and deployed to new contexts at a relatively low cost. AI functionality is increasingly embedded in standard software packages, such as software as a service (SaaS) offerings.

Box 1: What is an AI System?

The OECD defines an AI system as follows:

An AI system is a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment.

[1] See Box 1: What is an AI System?

Governance implications

Pre-AI or legacy software systems are built with explicit rules coded by developers, making their behaviour generally more predictable and understandable, even where the software itself is extremely complicated.

By contrast, AI systems are generally not programmed with fixed rules that determine every output. Instead, developers train models on data so they can identify patterns and relationships and generate outputs, predictions, or decisions. Some models rely on billions of internal parameters to infer these patterns. This can make AI systems less predictable and more difficult to explain than traditional software.

These challenges are becoming more pronounced with agentic AI systems, which can create content, interact with users and perform tasks with increasing levels of autonomy.

Different categories of AI

AI is not a single technology. Organisations are using a variety of AI systems with different capabilities, risk profiles and governance implications. Those that are currently being used by organisations may contain any, or all, of the following categories outlined in Figure 1.

Figure 1: Key categories of AI being used in Australian organisations

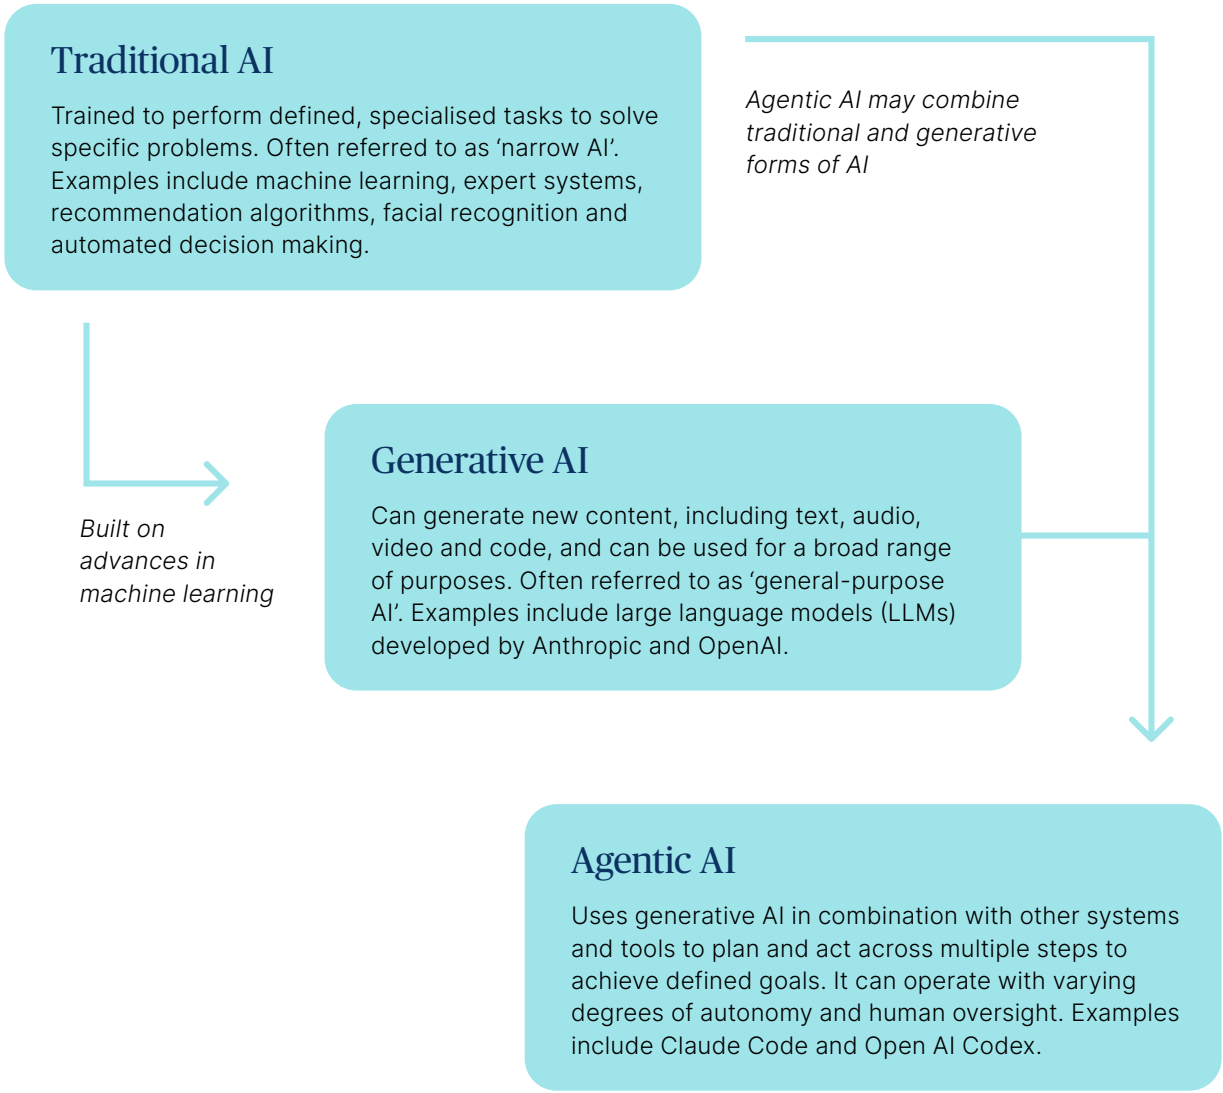

AI adoption in Australia

AI adoption in Australia has accelerated sharply. HTI research conducted with 419 directors, senior executives and decisions-makers in late 2025 found that 90% of surveyed organisations are using or planning to use AI, up from 64% in 2022. The results from the forthcoming 2025 survey are referred to as the HTI Corporate Leaders Survey in this guide.

Adoption varies significantly by industry and organisation size. The industries with highest AI adoption are retail trade and health and education, with services and hospitality close behind. According to the Australian Government's National AI Centre (NAIC), in 2025, 82% of those with 200–500 employees have adopted AI, compared with 33% of micro-businesses.

A number of dynamics help explain this growth. First, AI has moved from pilot stage to deployment. Since the arrival of ChatGPT in late 2022, generative AI has become increasingly embedded within third-party products. Second, cloud-based platforms have 'democratised' access to AI, enabling non-technical teams to adopt AI tools, without specialist teams needing to build or manage them. Lastly, the potential of the technology has increased exponentially supported by vast investments in AI infrastructure and computing power.

The use of agentic AI is also increasing. In the HTI Corporate Leaders Survey, nearly a third of organisations (32.5%) reported at least one use of agentic AI, with a further 12.9% planning adoption. Agentic AI is being deployed in IT and security (35.6%), marketing and sales (34.5%), and back-office activities (31.6%). A 2026 Deloitte report similarly found that 23% of businesses are already using agentic AI to at least a moderate extent, with this figure expected to increase to 74% within the next two years.

AI uptake by SMEs

AI adoption among small- and medium-sized enterprises (SMEs) is increasing, with evidence that 42% of SMEs are now using AI. Data-heavy industries such as property services, finance and insurance, and business services are leading adoption.

Source: NAB (2026) Embracing AI: Adoption & Key Opportunities Identified by SMEs

AI tokens and the cost of AI models

AI tokens are the units that large language models (LLMs) use to process information (see Box 2). Rather than reading text as full words or sentences, AI systems break inputs into tokens which may represent whole words, sub-words, characters, or punctuation. This process, known as tokenisation, converts human language into numerical representations that an AI system then analyses mathematically.

Organisations and their boards should be aware of the concept of tokenisation as the cost of using an AI system in an organisation (e.g. ChatGPT) is now generally priced on a per token basis. Organisations are typically charged separately for input tokens (what is sent to the system) and output tokens (what it generates), with output tokens often costing more due to the computational demand.

How employees are educated in prompt design and efficient use of AI systems is an increasing focus for Australian organisations.

While token usage influences AI costs, boards should also consider broader financial, operational and governance implications. See AI risks and harms for further discussion

Box 2: 'Tokenomics'

At the time of publication in June 2026 there was significant public commentary from Australian business leaders on the rapidly increasing cost of AI models due to significant growth in use and the greater computational complexity of agentic AI models. This increasing cost is in a context of continuing uncertainty on the return on AI investments.

Measuring the return of AI investments is discussed further in Part 3 of this publication.

Source: AFR, AI's big cost problem is going to get only worse, 2 June 2026

AI use is often embedded with limited visibility

AI use is not always visible to boards. AI encompasses a wide range of technologies and capabilities, many of which are being embedded within third-party products, sometimes without clear disclosure. This can make AI use more difficult to identify and govern.

Directors should be alert to terminology that may signal AI use within their organisation. As outlined in Box 3, terms such as 'model', 'algorithm', 'predictive analytics' or 'agent' may indicate the presence of AI systems. As AI advances rapidly, boards should take a broad view of what constitutes an AI system within their organisation.

Box 3: Terms that can signal AI use

- **Model or algorithm**: Software designed to provide recommendations, optimise systems, or prioritise actions.
- **Expert systems**: Systems that use a knowledge base, inference engine and logical rules to mimic human decision-making.
- **Training data**: Data used to train or fine-tune an AI algorithm.
- **Natural language systems**: Systems that can understand and use natural language.
- **Process automation**: Technologies such as robotic process automation that perform repetitive tasks.
- **Automated decision-making**: The use of a rules-based or a self-learning algorithm to make a decision.
- **Agents**: AI systems that perform tasks autonomously, such as developing software code.

How AI can be used by organisations

HTI has identified a range of common AI use cases that can help boards consider how and why AI is being deployed within their organisations.

- **Tactical tools**: Individual or ad hoc use of AI by staff to deliver localised productivity gains. For example, a secure AI assistant to draft correspondence, summarise long documents, or accelerate analysis.
- **Process transformation**: AI deployed to redesign specific end-to-end processes, with measurable efficiency or service improvements as the goal. For example, a customer-facing AI agent handles routine queries, allowing call centre staff to focus on more complex cases.
- **Systematic uplift**: AI applied to core enabling infrastructure or standard ways of working across the organisation, intended to improve flexibility and efficiency at scale. For example, a generative AI system that modernises legacy code or helps identify and patch cyber vulnerabilities.
- **Organisational transformation**: AI used to reshape the operating model. For example, redesigning HR and technology functions to operate as a more integrated function to accommodate how AI agents work alongside employees.

Boards should consider which category their organisation's AI investments fall within and whether the governance frameworks in place are adequate for that level of ambition.

Intersection with cyber security and data governance

AI systems are fundamentally dependent on data and digital infrastructure. As a result, AI governance cannot be considered in isolation from broader data governance and cybersecurity frameworks.

Comprehensive data governance provides the foundation for effective AI use by helping ensure data is accurate, lawful, relevant and appropriately managed across its lifecycle. Strong cyber security controls help safeguard the confidentiality, integrity and availability of that data and the systems that process it.

Weaknesses in either domain can directly undermine AI performance, introduce bias, and expose organisations to significant operational, regulatory and reputational risk.

This publication references existing AICD resources on cyber security and data governance where relevant (Figure 2). The AICD and HTI encourage directors and others involved in organisational governance to refer to these resources for additional guidance in these important areas.

Figure 2: AICD resources and education on digital governance oversight

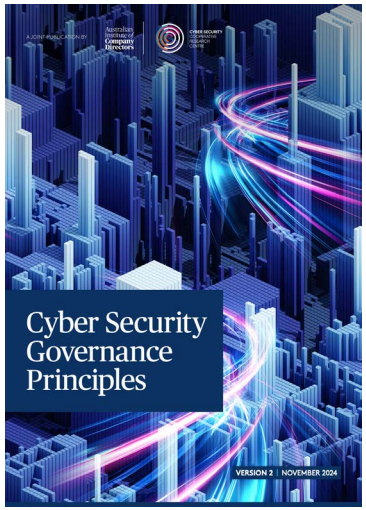

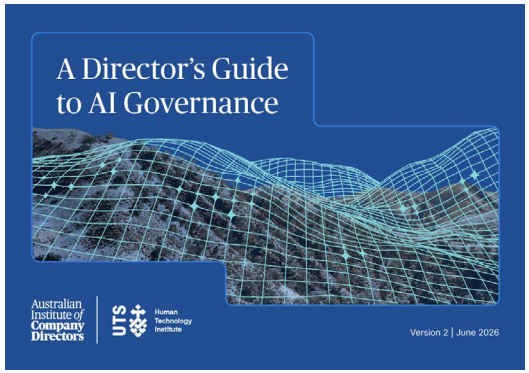

Director education:

- AI Fluency for Directors Sprint (in partnership with University of Sydney and Deloitte)
- AI Governance for Directors Webinar Series

The role of the board and the regulatory landscape

AI can be a valuable tool that can assist boards and individual directors in carrying out their roles and meet their obligations. However, it does not replace the essential human oversight role played by the board and individual directors.

Key points

- Boards have a key role in overseeing how an organisation uses AI, balancing the opportunities and risks of this transformative technology.
- Board-level oversight of AI and its associated risks forms part of directors' existing duties.
- Organisations are expected to govern AI within existing legal and regulatory frameworks, including obligations relating to privacy, consumer protection, discrimination, work health and safety, copyright and cyber security.

Role of the board

The board plays a critical role in overseeing that the organisation's use of AI aligns with its strategy, risk appetite and values.

The opportunities of AI are unlikely to be realised, and the risks are unlikely to be effectively managed, without active board oversight.

Active board oversight of AI involves setting expectations for responsible AI use, overseeing material AI-related opportunities and risks, and ensuring appropriate governance, reporting and assurance arrangements are in place.

In practice, AI investments and deployment are often governed through existing IT, digital or transformation governance structures. However, the transformative nature of AI warrants heightened board attention, given its potential impacts across the organisation.

The board should satisfy itself that appropriate controls are in place to manage the legal, ethical, reputational and operational risks associated with AI. This is explored further in Part 2 – AI Governance Operating Model.

Directors do not need to be AI experts. But they do need a minimum viable understanding of AI to ask the right questions, make informed strategic decisions, and ensure appropriate governance structures and processes are in place. Box 4 outlines a set of practical suggestions, developed by Shirley Chowdhary, to help boards start to embrace AI.

Box 4: Ways boards can harness AI

1. **Oversight**: Building awareness & assurance
2. **Boardroom wisdom**: Generating sharper director thinking
3. **Strategy**: Shaping long-term value
4. **ESG**: Reframing culture
5. **Resilience**: Managing dependency & risk

Further detail on potential applications of AI in the boardroom is available on the AICD website.

AI and directors' duties

Board-level oversight of AI and AI-related risk management forms part of directors' existing duties under the *Corporations Act 2001* (Cth) (Corporations Act).

When making decisions and providing oversight of their company's development and use of AI systems, directors are required to act with due care and diligence, in good faith, and for a proper purpose.

The Australian Securities and Investments Commission (ASIC) has emphasised the importance of board oversight of AI and related risk settings: 'Company directors and officers must discharge their duties with a reasonable degree of care and diligence. These duties extend to the adoption, deployment and use of AI. Directors and officers should be aware of the use of AI within their companies, the extent to which they rely on AI-generated information to discharge their duties and the reasonably foreseeable associated risks.'[2]

The AICD practice statement, Directors' oversight of company compliance obligations and the supporting legal opinion (October 2024) outlines the duty of care and diligence that is owed by company directors. This duty is central to how boards approach non-financial risk, including AI governance.

Ethical considerations

As outlined in this publication, AI adoption by organisations can give rise to particular risks because AI systems may behave in ways that are less predictable, less transparent and harder to test or fully explain in advance.

As AI is increasingly integrated into decision-making, customer interactions and core operations, boards may need to grapple with novel questions about accountability, human oversight and ethical judgement. As the legal and regulatory boundaries are still blurry, boards will inevitably be required to exercise judgements about the appropriate adoption and use of AI systems.

Boards may therefore need to grapple with ethical questions that extend beyond legal compliance, including the fundamental question: 'We can, but should we?'

The AICD and The Ethics Centre joint publication, Ethics in the Boardroom (2nd edition), provides a framework for boards to approach ethical decision making.

AI as a governance tool

Generative AI has the potential to be a valuable tool that can assist boards and individual directors to carry out their roles and meet their obligations. However, boards and directors should approach the use of AI in governance functions with caution, recognising that AI tools do not replace the essential human oversight role played by the board and individual directors, nor do they mitigate blind spots or substitute for critical thinking.

AICD research in AI Use by Directors and Boards (2025) found that AI use in Australian boardrooms remains cautious and limited, with adoption at the time driven by informal individual experimentation rather than collective board decisions.

> "The article [AICD publication AI use by directors and boards: Early insights] repays close reading and explains that Australian boards are cautiously but increasingly experimenting with AI as a governance support tool, shifting the focus from how companies generally oversee AI in their organisations to how AI might assist the directors themselves in the discharge of their duties."
>
> **Justice Michael Lee**

[2] ASIC, Beware the gap: Governance arrangeents in the face of AI innovation (Report No 798, October, 2024).
[3] Paul Smith, 'Telstra's board was drowning in data. An AI bot has changed everything.' Australian Financial Review (23 April 2026).

Since that publication, there have been advances in how boards are assessing the use of AI to support their responsibilities. For example, Telstra has highlighted the development of an in-house agent to assist directors in analysing board packs.[3] Similarly, in *ASIC v Bekier* (Liability Judgment) [2026] FCA 196, Justice Michael Lee commented on the use of AI in analysing board papers, noting the need for caution.

> "There is nothing inherently objectionable in obtaining such assistance [from AI], but what ought not occur is that this development becomes an excuse for a failure to instil discipline in the provision of information to directors or leads to a quiet normalisation of private reliance by them upon computer-generated distillations, unregulated by any agreed policy."
>
> **Justice Michael Lee**

The following principles should be prominent in deliberations on AI use in the boardroom:

- Directors should not rely on AI-generated summaries or analysis as a substitute for their own review and interrogation of board papers.
- Boards should ensure the organisation's AI register, inventory and/or risk frameworks are updated to reflect management, committee and board use of AI.
- Directors should have a level of AI literacy that enables them to understand the strengths and deficiencies of AI tools that could be used to support governance tasks, including risks associated with bias, data quality, opacity and security.
- Directors should take care to ensure the board or individual director use of AI does not undermine confidence in management or blur lines of accountability.[4]

Where organisations are considering using AI to prepare board minutes, the AICD and Governance Institute of Australia have provided practical guidance through the joint statement Effective Board Minutes & the use of AI.

Regulatory obligations

In December 2025, the Australian Government's National AI Plan confirmed that Australia will not introduce standalone AI legislation and will instead rely on existing laws to regulate AI, while leaving open the possibility of targeted reforms where regulatory gaps emerge.

This places an onus on organisations to ensure AI systems are governed appropriately within existing legal and regulatory frameworks.

A range of existing legal obligations may be relevant where an organisations develops, procures or uses AI systems, including obligations relating to privacy, industrial relations, consumer protection, market disclosure, anti-discrimination, copyright and work health and safety. Table 1 provides an overview of key legal and regulatory frameworks relevant to AI.

Important notice: References to legislation and key resources in this publication are current as at June 2026. However, given the pace of change in AI, privacy and cyber security obligations, readers are encouraged to stay informed of developments.

This publication is intended as general guidance and does not constitute legal advice. The partners recommend seeking independent advice on legal, regulatory and technical matters.

[4] See supporting AICD resources to support boards in Appendix D.

Table 1: Overview of key regulatory frameworks

Law	Summary
Privacy Act 1988	The Privacy Act applies to the collection, use, storage, disclosure and destruction of personal information. The Act applies where personal information is used to train, test or use an AI system, unless the organisation is exempted from its operation. From 10 December 2026, organisations must disclose in their privacy policies where they use personal information in certain automated decision making.
Australian Consumer Law (ACL)	ACL prohibits misleading or deceptive conduct, unconscionable conduct, and false or misleading representations in trade or commerce. It also contains consumer guarantees and a product liability regime for defective goods. ACL may apply to cases of 'AI-washing', or other misleading representations about the capabilities, performance or use of AI systems.
Anti-discrimination laws	Commonwealth and State anti-discrimination laws prohibit direct and indirect discrimination based on protected attributes such as race, sex, age and disability. Organisations should ensure they do not engage in unlawful discrimination when using AI systems, including where a system has been trained on data reflecting historical biases.
Work health and safety laws	Australia has harmonised work health and safety laws that impose a primary duty of care to ensure, so far as is reasonably practicable, the health and safety of workers and other persons. These obligations may be relevant where organisations develop, procure or deploy AI systems. The NSW Parliament has also enacted a duty to ensure, so far as is reasonably practicable, that workers are not exposed to health and safety risks arising from the allocation of work by a digital work system.
Law of negligence	Organisations that use AI systems owe a duty of care to people affected by that use, so as to avoid causing reasonably foreseeable harm. An organisation may breach this duty of care where harm results from foreseeable risks associated with the AI system and the organisation failed to take reasonable steps to prevent or mitigate those risks.
Copyright Act 1968	The Copyright Act protects original works created by human authors. The use of copyright-protected material in connection with the training or use of AI systems may involve acts of reproduction that infringe copyright – unless the organisation has the consent of the copyright holder, relies on a licensing arrangement or the use falls within a fair dealing exception.
Fair Work Act 2009	Under the Fair Work Act, and most modern awards and enterprise agreements, the introduction of AI will often constitute a 'major workplace change' where it is likely to have significant effects on employees (e.g. changes to roles, hours, skills requirements or job loss). In those circumstances the employer undertakes certain steps, including consultation and consideration of measures to mitigate adverse impacts.

Cyber security risk and regulatory obligations

As discussed further in AI opportunities and risks, there is a significant interaction and interdependency between cyber security and data governance settings, and the adoption of AI by organisations.

Both ASIC and the Australian Prudential Regulation Authority (APRA) have called for greater board vigilance on how organisations are adjusting cyber and data risk settings to account for rapid advancements in AI. Guidance from APRA, ASIC and Australian Signals Directorate (ASD) is highlighted in Box 5.

The *Cyber Security Act 2024* (Cth) and separately the *Security of Critical Infrastructure Act 2018* (Cth) impose a range of cyber security and related critical asset risk management and reporting requirements that may be relevant to AI adoption.

Box 5: ASD, APRA and ASIC guidance

In May 2026, the ASD, APRA and ASIC each issued separate guidance in a short period on the elevated cyber risks associated with AI adoption.

The guidance collectively had key messages for boards on overseeing proactive steps to enhance cyber risk controls in an AI age. This included rapid patching of vulnerabilities and strong privileged access management.

An AICD summary of this guidance is available on the AICD website.

Australian Government guidance

The Australian Government has provided guidance on the ethical principles and AI governance practices that support responsible AI use.

AI Ethics Principles

Australia's AI Ethics Principles provides a voluntary set of principles to guide the responsible design, development and use of AI (Box 6). They align closely with the OECD AI Principles and reflect a consistent international approach to safe and responsible AI.

Guidance for AI Adoption

In October 2025, the NAIC published its Guidance for AI Adoption, which sets out six essential practices for managing AI systems across their lifecycle (known as the AI6, Box 7).

The AI6 translates the Australian's AI Ethics Principles into practical governance actions. It replaces the earlier Australia's Voluntary AI Safety Standard, consolidating the former 10 voluntary guardrails into six practices.

The AI6 complements the AI Governance Operating Model detailed in this publication.

Two versions of the NAIC guidance are available: Foundations, for organisations getting started in adopting AI; and Implementation guidance, for governance professionals and technical experts.

Box 6: Australia's AI Ethics Principles

- Human, societal and environmental wellbeing
- Human-centred values
- Fairness
- Privacy protection and security
- Reliability and safety
- Transparency and explainability
- Contestability
- Accountability

Box 7: The AI6 – Six Essential Practices

The AI6 identifies six essential practices:

1. **Decide who is accountable**: Assign clear accountability for AI systems across their lifecycle, including roles for approval, operation, monitoring and oversight.
2. **Understand impacts and plan accordingly**: Identify and assess potential impacts on people, communities and the organisation, and plan appropriate safeguards and response mechanisms.
3. **Measure and manage risks**: Establish processes to identify, assess and manage AI-related risks on an ongoing basis, integrated with existing risk frameworks.
4. **Share essential information**: Provide appropriate transparency to users and stakeholders about how AI systems are used and how decisions are made.
5. **Test and monitor**: Test AI systems before deployment and continuously monitor performance to identify issues, drift or unintended outcomes.
6. **Maintain human control**: Ensure meaningful human oversight, including the ability to intervene where necessary.

International regulatory landscape

Globally, AI regulation is evolving rapidly but jurisdictions are adopting markedly different approaches to it. An overview is provided in Table 2.

Table 2: Selected international approaches to AI regulation

Country / Region	AI regulatory approach
European Union	The European Union has adopted the most comprehensive AI-specific regulation through its Artificial Intelligence Act (EU AI Act). The Act establishes a risk-based regulatory framework, with graduated obligations depending on the level of risk posed by an AI system. Implementation is phased, and the Act has extraterritorial application, meaning it may apply to organisations outside the EU in certain circumstances.
United Kingdom	The United Kingdom has adopted a pro-innovation, principles-based framework that relies on regulators to apply existing laws in line with five key principles: safety, transparency, fairness, accountability and contestability.
United States	The United States does not have a comprehensive federal AI statute. A growing number of US states have enacted AI-related laws that address specific risks and use cases.
Canada	Canada has not enacted comprehensive AI-specific legislation. Although the proposed Artificial Intelligence and Data Act was introduced into Parliament, it did not progress. AI is currently regulated through existing federal and provincial laws.
China	China has a regulatory framework that integrates AI governance into existing laws, supplemented by AI-specific administrative measures.

AI opportunities and risks

The potential benefits of AI heighten community expectations that AI systems are governed effectively and used lawfully, ethically and transparently, particularly in relation to privacy, security, bias, accountability and human impacts.

Key points

- AI systems present a range of opportunities for organisations, including improvements in productivity, product quality, customer service and employee experience.
- At the same time, AI systems can introduce new commercial, reputational and regulatory risks, while also amplifying existing risks, including those related to cyber security and privacy.
- AI has the potential to have significant human impacts, including on employment, and boards should maintain close oversight of these impacts.

AI benefits and opportunities

AI has the potential to deliver significant benefits for Australian organisations, including productivity gains, improved decision-making, enhanced customer and stakeholder engagement, and new sources of strategic insight. For boards, the challenge is not whether to engage with AI, but how to realise these benefits in a way that supports long-term value creation while remaining aligned with sound governance, ethics and risk management.

At the same time, increasing competitive pressure is accelerating AI adoption, with many organisations feeling compelled to move quickly to avoid strategic disadvantage as peers, competitors and global counterparts embed AI into core operating models.

Potential benefits include:

- **Improved productivity:** Some AI systems can reduce the burden of administrative tasks through new forms of automation. Others can increase employee productivity and support higher-value work by helping teams analyse trends, summarise information and generate new content. For many organisations, these are among the primary benefits of adopting AI.[5]
- **Reduced errors and improved quality**: While AI systems can be prone to errors when inputs fall outside their training data, they can perform significantly better than other approaches, including many human experts, for well-defined mechanical or repetitive tasks, particularly those involving pattern recognition.
- **New products and services:** AI systems can support innovation by assisting organisations to design, prototype and test new products and services. This includes developing new digital or software products and supporting the maintenance of legacy products where there may be scarce expertise.
- **Improved customer experience:** Through their ability to engage in natural language and scale digitally, AI systems can reduce customer wait times, synthesise existing information for accessibility, and deliver more personalised customer experiences.
- **Improved employee experience:** AI systems can help to reduce the time spent on low-value administrative tasks so that staff can focus on high-value work and innovation. Some AI systems (such as generative AI) can also guide workers through more complex tasks and assist in problem-solving.

[5] Australian Industry Group, Technology Adoption in Australian Industry, Commercial, Workforce and Regulatory Drivers, (Report, October 2024).

These benefits can be significant, but they are not automatic. Realising them requires strategic oversight, effective deployment and appropriate governance.

At the same time, under-investment in AI capabilities may leave organisations vulnerable to falling behind their competitors in terms of costs, product and service innovation, customer service and talent acquisition and retention. Organisations that have invested strategically in AI capabilities are reportedly outperforming their peers, with stronger returns from AI systems, compared to slower adopters.[6]

Box 8: AI and medical imaging

Research indicates that AI-enabled medical screening and imaging can improve diagnostic accuracy and sensitivity across a range of conditions when used alongside clinicians rather than as a replacement. AI assistance can increase detection rates and reduce clinician workload, supporting earlier intervention and more consistent clinical decision-making.

Source: National Institute of Health (2025) Artificial intelligence in healthcare and medicine: clinical applications, therapeutic advances and future perspectives

AI challenges and opportunities for small businesses and NFPs

Small businesses, NFPs and charities in Australia face unique AI governance challenges, particularly where resources and specialist expertise are limited. However, there are still opportunities to unlock gains that can improve business operations and deliver improvements in products and services for customers, clients and beneficiaries.

Challenges

Small organisations often lack dedicated IT staff or data specialists, making it more difficult to assess the AI-related risks and opportunities, and to implement and govern AI systems effectively.

The upfront and ongoing costs of AI systems, including via software subscription models, can be burdensome for organisations operating on tight budgets. Additionally, staff training on AI best practices can often take a backseat to more immediate operational concerns.

Many charities and NFPs may also face heightened data governance and privacy risks due to the nature of the information they collect, including information relating to vulnerable members of the community.

This dynamic may make the organisation inherently cautious about implementing AI given the potential risks and tension with the purpose of the organisation. Further, resource constraints may limit the ability of a small business or NFP to respond to elevated cyber risks from AI, for example quickly patching vulnerabilities in legacy systems.

Opportunities

Despite these challenges, smaller organisations can increasingly deploy AI to drive organisational improvements.

Limited size and resources need not be an insurmountable barrier to using accessible and low-cost AI tools. For instance, SMEs may benefit from using AI features within accounting products to reduce administrative workload. Similarly, AI features within social media platforms can assist SMEs with more effective targeted advertising and marketing initiatives. Charitable organisations may also benefit from AI-generated insights into fundraising patterns.

As AI capabilities become increasingly embedded within existing software-as-a-service (SaaS) products, the opportunities for SMEs and NFPs to harness AI's potential will only increase.

[6] For example, see Boston Consulting Group (2024) Where's the Value in AI?.

AI risks and harms

The characteristics of AI systems that distinguish them from traditional software can also amplify existing risks and harms, and create new ones.

Directors should understand the key risks for organisations associated with deploying AI systems and the types of harm they may present.

Where do AI risks and harm come from?

The use of AI systems can create or amplify harms in three key ways, as outlined in Table 3.

Table 3: Overview of three key sources of risk

Key source of risk	Summary and examples
AI systems failure	Where AI systems create harm because they fail to perform as intended. • Poor system performance • Biased system performance • System fragility or unreliability • Security failures or vulnerabilities
Malicious or misleading use	Where AI systems are deliberately used in ways that create or amplify risk of harm. • Use in scams and other fraudulent activity • Misinformation at scale • AI-powered cyber attacks • Price or service discrimination
Inappropriate or reckless use	Where AI systems are used in a way that does not sufficiently consider the risk of harm. • Use of facial recognition technology without authority • Mass surveillance of the community • Excessive deployment across a workforce leading to loss of critical skills • Overuse of generative AI resulting in significant energy consumption and emissions

Key areas of AI related risk

Table 4 summarises key areas of AI-related risk that are front of mind for Australian boards, according to the HTI Corporate Leaders Survey.

How boards can oversee effective controls to mitigate these risks is discussed further in Part 2 – AI Governance Operating Model.

Table 4: Overview of key AI-related risks

Risk	Description
Cyber security	AI systems introduce new cyber threat vectors and amplify existing vulnerabilities. Malicious actors can use generative AI to produce convincing phishing content or deepfakes. An organisation's own AI systems can also be targeted, for example through prompt manipulation to extract sensitive information or training data poisoning that distorts model behaviour. These risks are often compounded by gaps in basic controls such as access management.
Privacy and data governance	AI systems can create privacy and governance risks where they collect, use or generate personal information. Privacy obligations apply not only to data inputted into AI systems, but also to outputs that contain personal information, including inferred or incorrect information. Risks include the use of personal information beyond its original purpose and the inadvertent exposure of sensitive data.
Output issues: quality and explainability	AI systems can produce outputs that are inaccurate or misleading. Generative AI systems are particularly prone to 'hallucinations', which are coherent but incorrect outputs. This is because they generate responses probabilistically rather than operating within the defined parameters of traditional AI. Explainability is a related challenge. It is often not possible to explain how a model arrived at a particular output, or to fully test and validate its reasoning. Where AI is used in decisions with legal or similarly significant effects, this can make it difficult to justify outcomes or provide meaningful explanations.
Employee and workforce impacts	The adoption of AI systems is expected to significantly impact workforces. While some roles and tasks may be automated, AI is also likely to augment workers by supporting productivity and enabling new ways of working. This issue is discussed further in the the human impact section of this publication.
Fairness and discrimination	Bias is one of the most well-documented concerns related to AI systems, reflecting the potential for AI to inherit and amplify biases present in real-world data. It can result in AI systems producing outcomes that are unlawful and discriminatory, disadvantaging individuals or groups based on protected characteristics. Bias may emerge from pre-existing biases in the real world, non-representative datasets, or the selection of algorithmic approaches that embed the bias.
Shadow AI	When staff use publicly available AI systems (such as ChatGPT) at work, they may not be aware of the risk that information uploaded into them will lose confidentiality, or the potential for the system to generate information that seems credible but is not correct (i.e. hallucinations).

Agentic AI – Potential for elevated risks

Agentic AI systems rely on generative AI, and therefore inherit many of the same risks, including the risk of hallucinations. In addition, agentic AI may also create new risks as a result of its defining features: its ability to operate autonomously, in complex environments and at scale.

A key risk is that an agent does something that was not intended or authorised by the organisation deploying it. This may occur because the instructions were unclear, the agent misinterprets its goal, or pursues it in unintended ways, or makes errors while executing tasks.

Two additional characteristics can amplify these risks.

- Risks may increase where agentic systems are required to orchestrate across multiple agents or other organisational systems. At the current stage of the technology, AI agents can struggle with ambiguity or complex contexts. Coordination across multiple agents or systems may not always operate smoothly. In addition, orchestration across multiple systems can expand the attack surface available to malicious actors. They may be able to target not only the agent itself, but also the other systems with which it interacts.[7]
- Risks can be amplified by the scale and continuity at which agentic systems operate. Where systems operate with high levels of autonomy, errors may go undetected for longer periods. If those systems operate continuously, such as around the clock or across multiple organisational systems, errors may compound significantly before they are addressed.

Box 9: Liability for the actions of AI agents?

AI agents are not employees or contractors acting on behalf of the organisation. They are IT systems.

Although Australian courts are yet to provide specific guidance on liability for AI agents, the issue has arisen in Canada. In that case, the Civil Resolution Tribunal of British Columbia swiftly dismissed Air Canada's attempts to argue that it was not responsible for misleading statements made by a chatbot on its website.

The article Agentic AI: rogue agents, real liability (2026), from the law firm Mallesons provides an overview of this complex emerging legal area.

The human impact

The decisions that organisations make about how to deploy AI are not just commercial. They can affect the lives of workers, customers, communities and the organisation's relationship with each. The speed of AI deployment by organisations is occurring in an environment where there is broad and deep mistrust of the technology in the Australian community.[8]

Boards play a key role in identifying these human impacts and overseeing steps to mitigate potential harms.

Some of these decisions are immediate. AI promises productivity gains, but organisations face choices about how those gains are realised, including whether AI is used to reduce headcount, or augment the workforce. These decisions involve trade-offs between efficiency and the value of human contribution, judgement and relationships. Where roles are displaced, an organisation's approach to transition, reskilling and redeployment can shape employee trust, retention, and its ability to attract talent over time.

Other decisions determine the organisation's wider footprint. AI systems can disproportionately affect groups who are already vulnerable, particularly where systems are used in decisions about access to products and services. Even relatively small errors or biases in a system can result in significant harms when applied at scale. Decisions about how to design, deploy and oversee AI can compound or mitigate these effects.

The human impacts of AI systems warrant close board attention. Decisions about AI adoption and use will influence how the organisation is perceived by its workforce, customers, regulators, and the broader community. As such, these considerations should form part of the organisation's decision-making framework for AI.

Sustainability considerations

AI systems have significant energy and resource needs across their lifecycle, including the manufacture of chips used in AI models and servers.

In practice, the environmental impact of using an AI system will vary depending on factors such as the nature of the query, the amount of processing required, the size of the model, and the energy sources used to power the supporting infrastructure. Data centres also require substantial amounts of water, primarily for server cooling.

Boards should be cognisant of the potential environmental impact associated with large-scale AI adoption. For organisations covered by Australia's mandatory climate reporting regime, emissions generated through the use of AI services downstream of the organisation (for example, at the data centre) may be reportable as scope 3 emissions.

[7] For more information, see Gradient Institute's 2025 report, Risk Analysis Techniques for Governed LLM-based Multi-Agent Systems.
[8] Tech Policy Design Institute, Earning Trust: Unlocking AI Adoption for Australians, May 2026.

Part 2 - AI Governance Operating Model

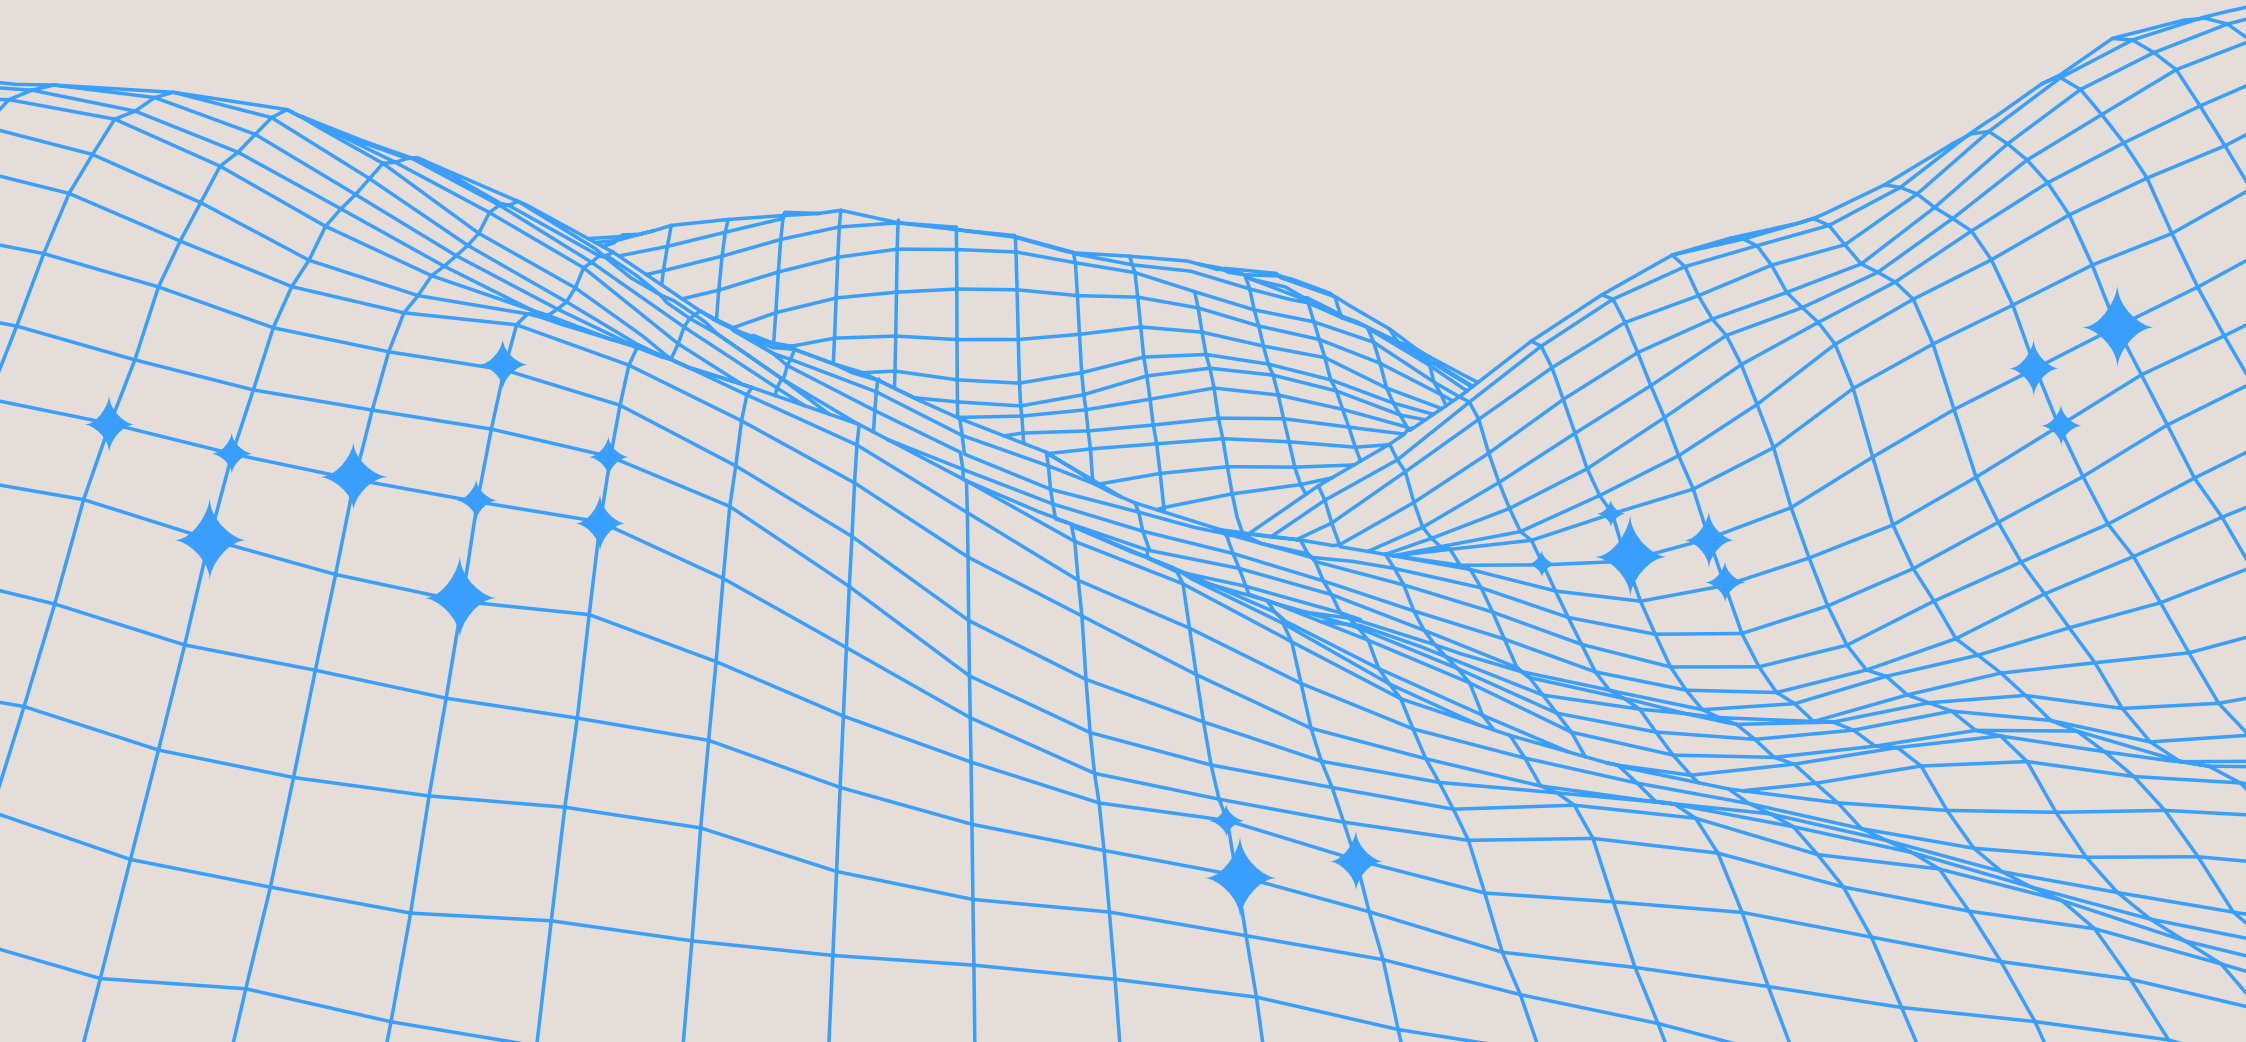

AI Governance Operating Model

Part 2 of this guidance presents a structured approach for board oversight of AI deployment and use in an organisation. The approach is based on the AI Governance Operating Model developed by HTI through its work advising organisations on AI governance. Each component of the model, shown in Figure 3, is explored in the sections that follow.

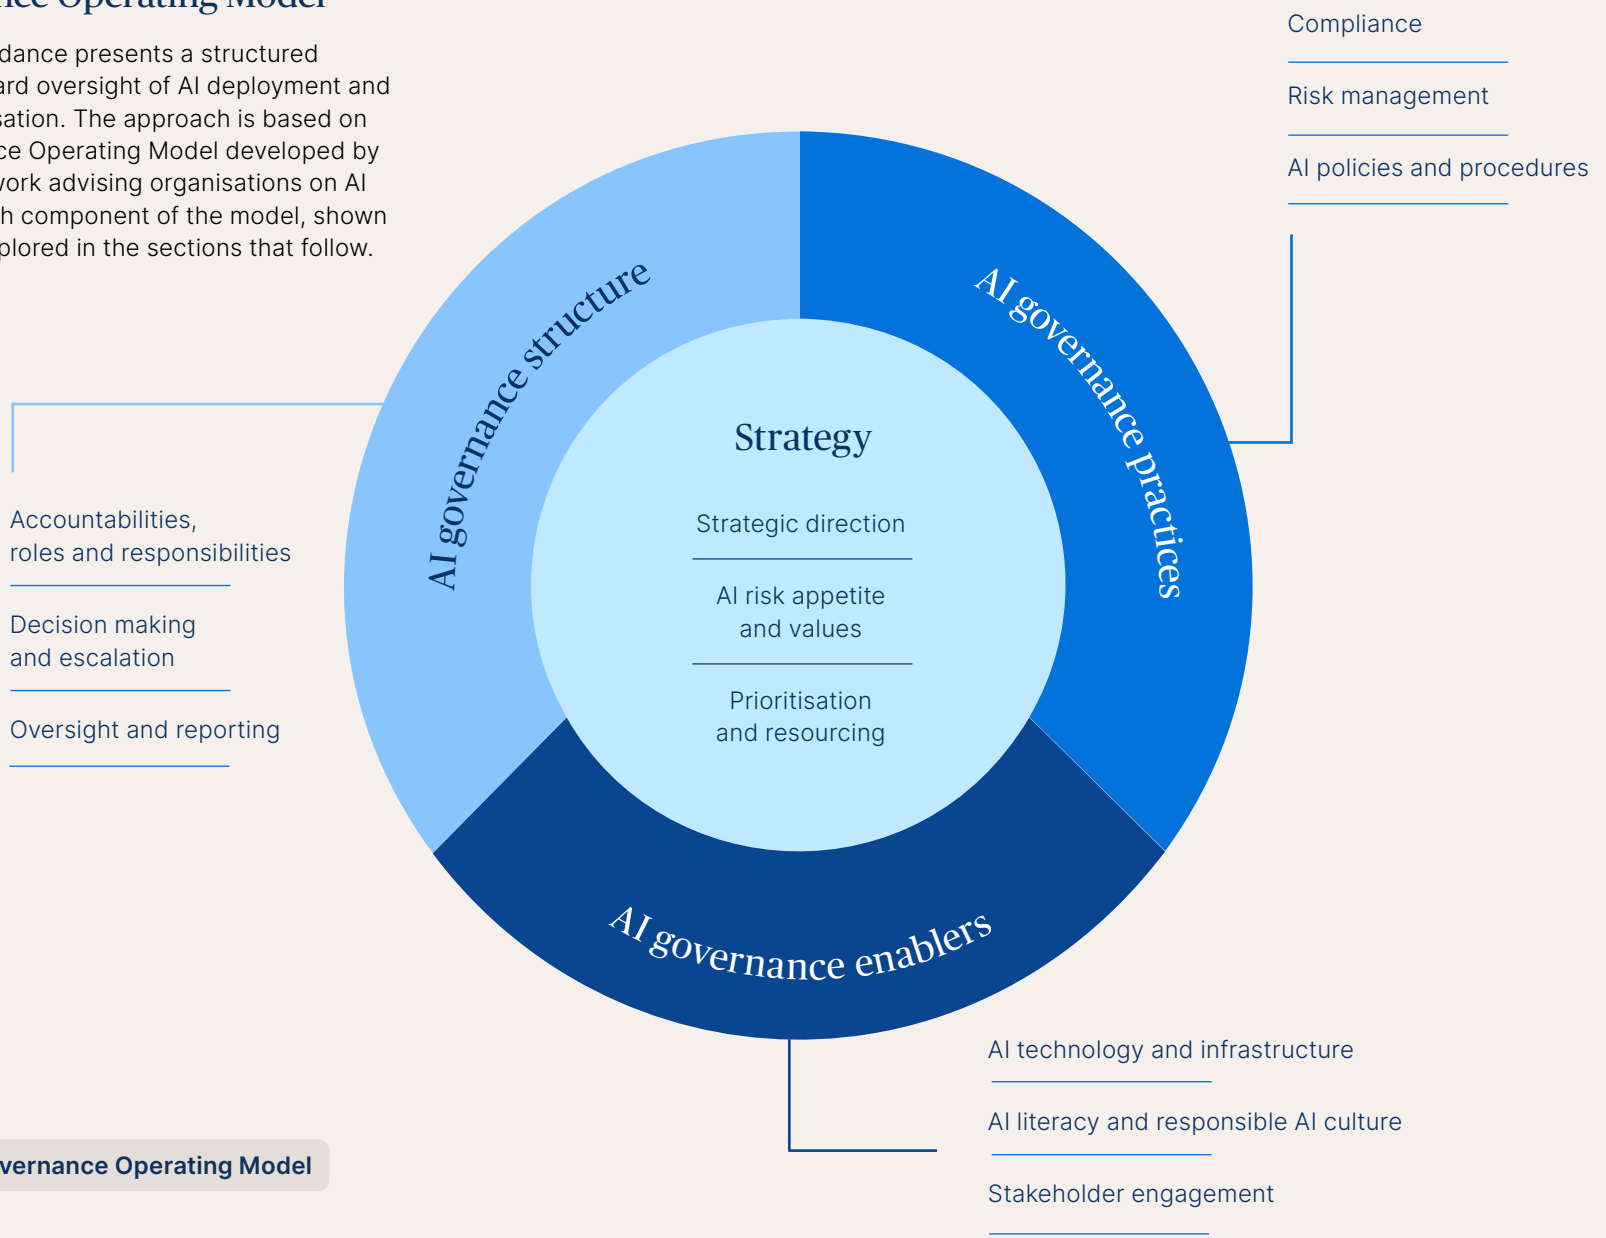

Figure 3: AI Governance Operating Model

Strategy

Decisions about AI adoption and use should be grounded in organisational strategy, supported by appropriate governance arrangements, and informed by a clear understanding of risks, opportunities and implementation requirements.

Key points

- The board should approach decision-making on AI and AI-related investments with discipline and have confidence that decisions align with the organisation's strategy, values and purpose.
- Given the distinctive risks associated with AI, it is appropriate for the board to work with management to establish a risk appetite for AI that is ultimately reflected in internal processes and policies, including the risk management framework.
- Effective AI implementation cannot occur in a vacuum. The board should have a clear view on the resourcing, data quality and prioritisation required to support AI adoption.

The board should discuss the strategic opportunities that AI presents for the organisation and, where appropriate, develop an AI strategy in conjunction with management.

Strategic direction

AI adoption and investment decisions should be aligned with the organisation's broader strategy, particularly its approach to digital and technology investments.

As with all significant capital and operational investments, the board should adopt a disciplined approach to considering the business case for AI investments. Pursuing 'AI for AI's sake' can result in investments that are not aligned with the organisation's strategy and may ultimately risk destroying value. It is often appropriate for organisations to trial or run pilot AI use cases to assist in informed organisational decision-making.

Based on these discussions, the board should set the strategic direction for management to implement. This direction will guide the organisation when making investment decisions about proposed AI systems and use cases. In practice, use cases that are not aligned with strategy can be filtered out early, before significant investment has been made.

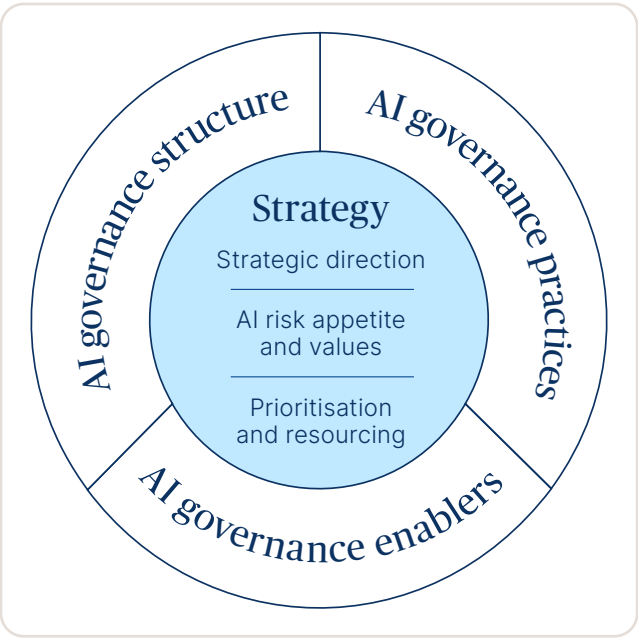

This enables management to focus on AI tools and use cases with genuine strategic value, rather than simply responding to the latest developments in the market.

How to measure the value from AI investments is discussed further in Part 3 – Measuring AI Returns.

AI risk appetite and values

The board should discuss its risk appetite for AI systems and communicate this to management. The risk appetite guides the organisation's deployment of AI tools, the controls adopted to manage any risks, and future investments in AI systems. In practice, the board may set different risk appetites across different AI technologies, tools and stakeholder groups (for example, generative versus agentic AI systems, or staff-facing versus customer-facing applications).

The board should also articulate the organisation's AI values, as these will guide its adoption and use of AI. These values will often reflect the ethical principles underlying responsible AI, including commitments to fairness, transparency and accountability. Once set, AI values can be translated into practice through other AI governance initiatives.

Prioritisation and resourcing

The board should ensure that management identifies and prioritises the AI governance initiatives needed to support the lawful and responsible use of AI. While some AI governance initiatives will be foundational, others can be implemented over time (for example, as the organisation scales up its AI use).

Key to prioritisation is establishing a comprehensive view of the current state of how AI and data are used across the organisation, and the internal and external resources that support them.

Baseline

Before developing and implementing an AI strategy, it is good practice for an organisation to undertake a detailed data inventory, stocktake or mapping exercise of how AI, machine learning and other data tools are being used, and importantly, which datasets and resources this use relies on. The importance of an AI inventory or register is discussed further in AI policies and practices.

This baseline provides the board and management with visibility over existing capabilities, data quality, risks and dependencies, ensuring future AI investments are grounded in operational reality rather than unrealistic confidence or ambition. This is particularly important for smaller organisations, where limited financial and human resources may mean AI investments need to be highly targeted and consistent with existing technological capabilities and skills.

Data quality

Organisational data, and the quality of this data, is the foundation and building blocks of effective AI. High-quality data is paramount because internal AI tools are trained on, and will continue to draw on, organisational datasets to produce results. Poor data governance, including underlying data quality issues, can result in misleading insights, flawed decision-making, and unintended biases, such as discriminatory customer profiling or hiring practices.

Internal and external resources

Management should provide the board with a clear understanding of the resourcing changes needed to support the organisation's AI strategy.

For larger organisations, AI deployment may require specialist internal teams and experts with knowledge of how to train and deploy AI systems and tools. This may include data scientists and engineers, AI architects, and specialists in underlying cloud computing and infrastructure (for example, Azure engineers). For smaller organisations this level of internal expertise may not be feasible given the cost and limited pool of people with these skills. These organisations are more likely to rely on external resources and assistance.

In practice, all organisations will rely on external providers to some extent, given that many AI tools are built on foundation models developed by organisations such as Anthropic and Google. Further, AI features are increasingly being embedded within existing SaaS products or offered through additional licences. Management should provide the board with a clear picture, supported by business cases, of the external resources and additional spending required to pursue the AI strategy. This should include an assessment of the training and support employees will need to effectively harness the benefits of AI investments.

More information – Data Quality

Further guidance on board oversight of data governance, including data quality is available in the AICD, MBS and Allens 2025 publication Data Governance Foundations for Boards.

Boards of SMEs and NFPs – AI Strategy

- Understand where AI is already being used in the organisation, including AI features in third-party products and services.
- Discuss the organisation's risk appetite for AI use with management.
- Understand whether new resources and significant investment is required to support AI deployment.
- Assess where AI could add value and support strategic objectives, including opportunities to enhance existing processes and resources (such as current SaaS subscriptions) in a cost-effective way.

Questions for directors to ask

1. Do we as a board have a clear understanding of the organisation's strategy with AI?
2. Have we communicated to management what we see as the key AI opportunities, and risks to be managed?
3. What is our risk appetite for AI use, and will this differ across different use cases or stakeholder groups?
4. Do we have agreed measures and indicators of AI progress and ultimate success, including return on investment?

Governance red flags

1. The organisation is adopting a large number of AI tools without reviewing alignment with strategy and/or potential return on investment.
2. AI initiatives are disconnected or misaligned from the organisation's core strategy.
3. The board hasn't considered the risk appetite for AI adoption consistent with the overall organisation's risk appetite, risk register and risk management framework.
4. We do not have a complete understanding of current resources and internal expertise and whether they will support our planned AI investments.

AI governance structure

An effective AI governance structure establishes clear roles and lines of accountability, defines decision-making processes, and creates reporting and escalation pathways through management to the board.

Key points

- The board has a key role in overseeing the use of AI in an organisation and should work with management to map AI-related responsibilities across the organisation.
- Dedicated AI governance structures and bodies may not be necessary for every organisation. What is critical is that the board has visibility of how AI is managed, monitored and governed, underpinned by sufficient board AI literacy.
- Roles and responsibilities for AI should be regularly reviewed and updated consistent with changing business operations, AI use cases and technological developments.

Accountabilities, roles and responsibilities

The board

While directors are not expected to become AI experts, it is important that directors have a sound understanding of AI and its associated risks and opportunities. If the board considers that it lacks these skills, it has several options:

- undertake board training to build the board's AI literacy;
- recruit additional members with the necessary skills and experience; and/or
- establish an advisory committee with AI experts that can be called on for advice.

See the section Role of the board and regulatory obligations for a more detailed discussion, and Box 10 for survey findings on where AI responsibility sits within Australian organisations.

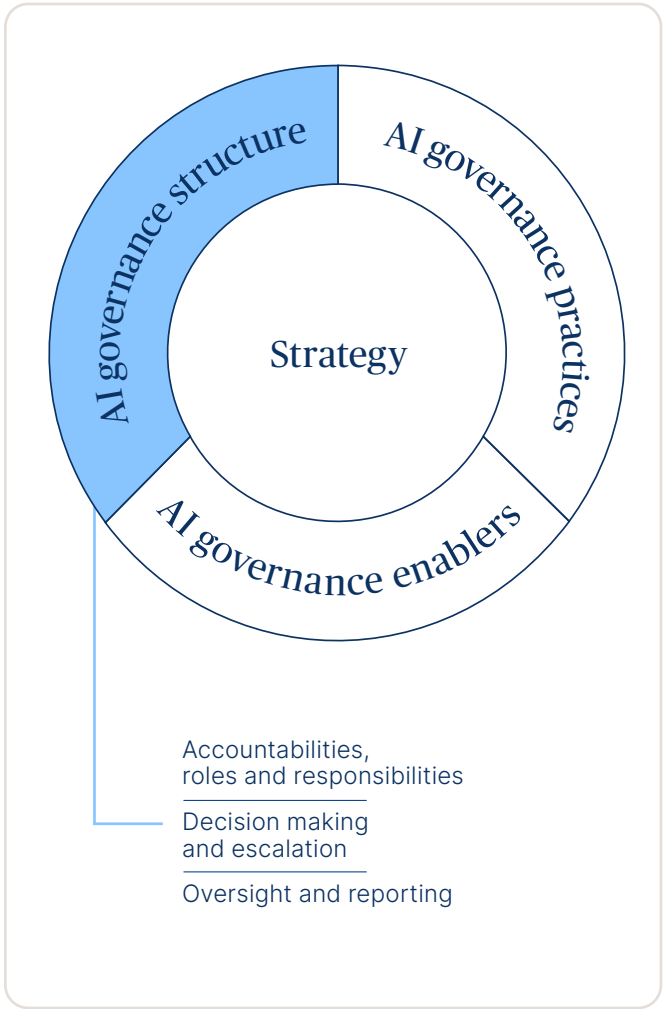

AI governance accountability

APRA's 2026 guidance set a clear expectation for its regulated population that governance structures should encompass 'ownership and accountability across the AI lifecycle, from design and development through to deployment, monitoring and decommissioning'. This principle is a valurable governance yardstick for all organisations.

At a minium, most large organisations should have a senior executive, or executives, who are accountable for AI governance. Responsibilities may include establishing the AI governance framework and overseeing its operation. In smaller organisations, the CEO may take on this role. In larger organisations, it may be a standalone role (such as a Chief AI Officer or Chief Data Officer), or responsibility may be shared across roles such as the Chief Technology Officer, Chief Information Officer and other senior roles.

Management responsibilities

In larger organisations, responsibility for AI and broader data governance is cascaded through management and embedded in specific roles. Individual AI and data responsibilities should be documented in position descriptions or role statements.

A frequent challenge is that AI responsibilities may be shared across different roles or blurred with related functions such as cyber security, data governance, and digital and IT operations and projects. Role maps, scenario testing or workshops can support staff to better understand where responsibility for AI and data sits, and how it overlaps with cyber security, data governance and broader digital functions.

Role of external providers

An organisation's documented roles and responsibilities for AI should reflect the critical role played by external providers. As discussed in the Strategy section, all organisations rely to some extent on external providers for AI tools and systems. The board should have visibility of the role these providers play in AI systems and how risks associated with these relationships are managed.

Information flow diagrams, which map how key data moves between the organisation and AI systems, including points of internal and external access and associated controls, can be a useful tool to support board oversight. In addition, these tools can also help with controls for supplier and vendor risks, discussed further in the AI governance practices section.

Box 10. Who has AI responsibility in Australian organisations?

The HTI Corporate Leaders Survey asked respondents which role(s) are primarily responsible for ensuring that AI systems operate safely, lawfully and responsibly.

The roles identified were:

- CEO or Managing Director (52.9%)
- Chief Technology Officer, Chief Information Officer or Chief Data Officer (46.1%)
- A legal or compliance executive (such as the Chief Legal Officer, General Counsel, etc) (31.8%)
- A risk focused executive (e.g. a Chief Risk Officer) (23.4%), and
- Board (21.8%)

The CEO was more likely to be identified as responsible in small and medium sized organisations; and a technology-focused executive was most likely in large organisations with more than 200 staff.

Decision-making and escalation

The board and management should consider how AI governance will be embedded across the organisation, including whether it is documented in a formal AI governance framework. Such a framework can encompass the approval of high-risk AI systems and use cases, oversight of AI risk management, approval of AI policies and processes, and reporting to senior management and the board.

Organisations may choose to use an existing governance process or to establish a dedicated AI committee, which is often management led. The benefit of a dedicated committee is that it can bring cross-functional expertise from across the business, support well-rounded decision-making and build internal capability in AI governance. It can also speed up decision-making by centralising decisions about new AI systems and use cases.

In smaller organisations, it may not be practical to establish an AI governance committee and management may rely on existing governance processes. Further, formalised documentation of responsibilities for AI may not be necessary. Nonetheless, individuals within the organisation should have a clear understanding of their responsibilities and contribution to the responsible use of AI.

The board should have confidence that AI governance structures and the relevant responsible people have the necessary authority and resources.

Oversight and reporting

Both boards and senior management rely on reporting about key risk and other performance indicators for effective AI governance. The board should work with management to determine the KPIs and cadence of reporting that it expects in relation to AI use and governance.

Reporting should provide a clear view of where AI is being used across the organisation, how systems are performing, and where risks are emerging or increasing. It should also help the board to assess whether AI use remains aligned with the organisation's strategy and risk appetite.

Directors should expect that board reporting is presented without unnecessary technical language that may act as a barrier to assessing AI performance and risks and engaging with management.

External assistance

Given the board imperative to monitor and stay across evolving AI risks and key capabilities, independent external experts can play an important role in providing an outside perspective.

For larger organisations, the board may wish to engage an external expert to periodically advise the board on AI governance. This can provide a more transparent view of the organisation's AI capability, including risks and comparisons with peer organisations.

That said, organisations should be cautious about becoming overly reliant on external experts given the increasing importance of AI and broader digital capability to all organisations.

Review

Boards should oversee the periodic review and updating of AI governance structures and accountabilities to ensure they remain fit for purpose as technologies, risks, key vendors and organisational use cases evolve. This includes testing whether roles, decision-making processes and reporting arrangements continue to support effective oversight, and requiring management to adapt governance frameworks where gaps or emerging risks are identified.

Boards of SMEs and NFPs – AI Governance Structure

- Assign accountability for AI oversight to an appropriate senior leader.
- Establish a proportionate process for approving the use of AI systems in line with their potential impact on the organisation and stakeholders.
- Establish when and how the board will be updated or consulted on AI systems, including working with management to develop targeted metrics on AI use, effectiveness, and risks.
- Consider how external assistance may help AI deployment in a cost-effective way.

Questions for directors to ask

1. Do we have a sufficient understanding of how AI technologies are being adopted and used across the organisation, and the associated opportunities and risks with this use?
2. Has the organisation identified a role with key accountability for AI governance and use?
3. Has management implemented an effective governance structure for oversight and decision making on AI use, including high-risk uses? Should we establish a management-led AI committee?
4. Do we received timely and comprehensive reporting on AI use throughout the organisation and is this reporting periodically updated?

Governance red flags

1. AI is adopted in an ad hoc and unplanned way without consideration of whether specific governance process and risk controls are needed.
2. It is not clear who has responsibility for AI governance in the organisation.
3. Agentic AI risks are not highlighted in the risk management framework and corresponding risk controls.
4. AI risks are not reported to the board or are not escalated in a timely way.

Case study: Commonwealth Bank of Australia (CBA)

At the CBA, the board holds the CEO and the executive leadership team accountable for the management of AI-related risks and opportunities. AI is treated as a material risk under the risk management framework (RMF) with the risk overseen at board and management levels through existing risk committee structures.

Management-led business unit risk committees assist in managing risk in line with the RMF and in respect of AI may review and evaluate AI models. Separately, CBA has a Model Risk Governance Committee which has the purpose to oversee the design and operation of the Model Risk Framework and support approval of AI models.

Lastly, CBA has an AI Risk Committee that oversees the AI risk framework and provides challenge and advice for higher risk AI use cases.

CBA, Our Approach to Adopting AI (December 2025)

AI governance practices

The board should understand how AI governance and risk practices reflect the organisation's AI strategy, risk appetite and principles for the use of AI.

Key points

- AI can introduce unique risks, as well as elevate existing risks, for organisations. The board should oversee how these risks are appropriately identified and managed through existing risk management frameworks and controls.
- The board should have visibility over how AI, cyber security and data governance risk controls intersect, and check that existing cyber controls (e.g. rapid patching) are also applied within AI environments.
- External vendors of AI systems and supporting infrastructure are central to most AI systems and tools. The board should have visibility over these the risks associated with these providers and their data protection settings.

Compliance

As detailed in The role of the board and the regulatory landscape section of this publication, a range of existing laws apply to organisations using AI systems. Effective reporting and testing processes should be in place to allow the board to assess whether the organisation is meeting its regulatory obligations as they relate to AI.

More broadly, organisations may have contractual obligations with key suppliers and customers that limits or prevents their information from being used by AI systems or tools, including for training purposes. Management should be able to assure the board that it has full visibility of the organisation's regulatory and commercial obligations as they relate to AI deployment.

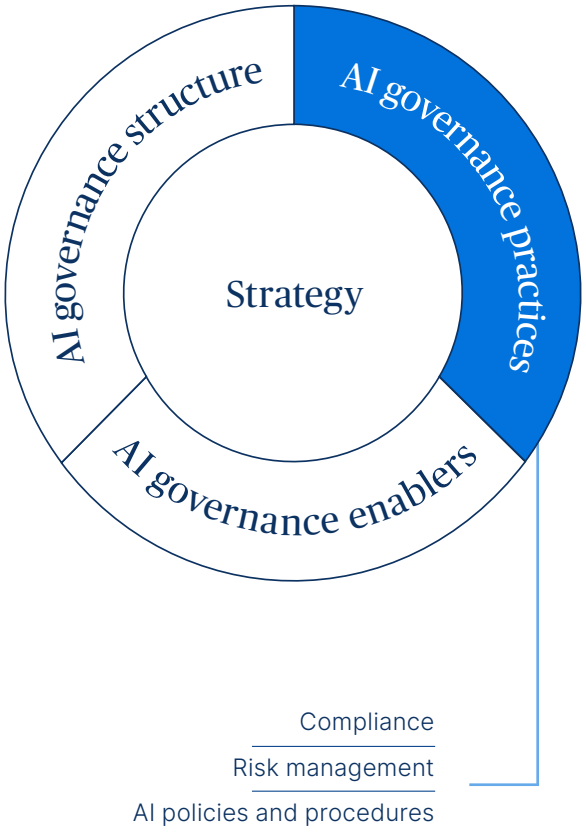

Risk management

The board should oversee a comprehensive organisational process to identify and manage AI-related risks. This should be aligned with the board's AI risk appetite. For many organisations, the identification and management of AI risks will be incorporated into existing risk management frameworks and control processes.

Once AI risks have been identified, organisations should implement effective controls to manage them. The controls required will depend on a range of factors, including the type of AI being used, the use case, the nature and severity of the risk, the organisation's risk appetite, and the context in which the AI system is operating.

Both risks and controls should be reviewed regularly as AI technologies develop and the threat environment evolves. Building on the risks outlined in Table 4, Table 5 outlines common controls for these risks that corporate leaders should be aware of.

Table 5: AI risks and common controls

Risks or risk drivers	Common controls
Cyber security	• Rapid patching • Existing controls (e.g. encryption, access controls, monitoring, incident response plans) updated to address AI-specific risks • AI threat modelling, red-teaming, and penetration testing
Privacy	• Privacy impact assessments for AI systems • Privacy policies and notifications addressing AI-specific data practices • Secure data throughout the AI system lifecycle
Output quality	• Model testing, validation and fine-tuning • Staff training on AI limitations • Human review of AI outputs, particularly in high-stakes contexts
Output explainability	• Documentation of model design, inputs, and intended use • Processes for explaining AI-driven decisions to affected individuals
Fairness and discrimination	• Representativeness and quality assessment of training data • Fairness and bias testing across protected attributes • AI impact assessments
Shadow AI	• Acceptable use policies or guidelines for employee AI use • Staff training on AI risks and obligations • Monitoring of AI use across networks
Agentic AI	• Documented risk appetite on what tasks an agent, or agents, can perform • Access controls limiting what agents can read, modify, or action • Logging and auditability of agent actions

Cyber security and data governance

The HTI Corporate Leaders Survey found that cyber security was the most concerning risk for board oversight in the context of AI. AI risks are inherently interconnected with both cyber security resilience and data governance. Often, boards will consider these risk areas in tandem.

While data governance traditionally focuses on collection, quality, accessibility and compliance, it also provides the foundation for effective AI system implementation and use. Further, datasets should be comprehensively protected through strong cyber controls, which are foundational to ensuring AI systems are both effective and trusted by internal and external users. The following examples illustrate the synergies between AI, data governance and cyber security:

- Clear data classification and lifecycle management underpin both cyber security (for example, least-privilege access) and the deployment of AI systems, helping to protect against data being misused or leaked through AI models or cyber incidents.
- Strong cyber controls, including secure configurations, vulnerability testing third-party risk management, apply equally to traditional IT and AI workflows, including AI-generated code and agentic systems.
- Unified data inventories, logging and monitoring enable faster detection of cyber incidents, AI model failures or misuse, and data breaches, allowing coordinated incident response, regulatory reporting and board oversight.

More information on cyber security and data governance controls

Further information on the board overseeing effective cyber security and data governance resilience is available from the AICD resources:

- Cyber Security Governance Principles (in partnership with the CSCRC)
- Data Governance Foundations for Boards (in partnership with Allens and MBS)

Workforce impacts

As AI agents become more widespread, a rising area of concern for directors and senior leaders is the potential of AI to result in significant workforce impacts. The board has a role in understanding these impacts and challenging management to ensure adverse effects are managed responsibly and ethically, recognising the potential for disruption and reputational risk.

The board should understand whether management is assessing impacts on different segments of the workforce, monitoring psychosocial risks (such as job insecurity or deskilling), and maintaining compliance with workplace and industrial relations obligations.

Directors should also consider broader societal expectations around fair transition, including engagement with regulators, unions and broader stakeholders where relevant. By embedding these considerations within governance frameworks alongside financial and operational metrics, boards can help in supporting the adoption of AI in a way that sustains organisational performance while maintaining trust and 'social licence'.

Key vendors and suppliers

External vendors are integral to AI systems. Many of the risks organisations face in respect to AI are linked to the products, services and infrastructure provided by these vendors, and the controls surrounding them. For example, a model selected by the organisation may contain embedded biases or incomplete data, resulting in unfair and discriminatory outcomes for customers or clients.

The risks can extend beyond the primary vendor (e.g. AI model provider) and reach a fourth layer as the vendor's models are hosted on particular 4th party cloud computing infrastructure. There can be pronounced concentration risks in AI supply chains due to commercial relationships and shareholdings between AI model providers and large or 'hyper scaler' cloud infrastructure companes.

While these risks may originate with vendors, accountability remains with the organisation. Boards should therefore oversee robust due diligence, ongoing monitoring, and a clear understanding of how vendor capabilities and weaknesses affect the organisation's risk profile.

The following steps can assist the board in overseeing these vendors:

1. Understand what testing and trialling of the vendor's products has been undertaken prior to adoption, and whether the results are consistent with the organisation's AI risk appetite.
2. Understand the provider's location and ownership structure, including interdependencies with other IT systems and infrastructure providers, and what fallback arrangements are available in the event of a system failure.

3. Monitor the provider's data governance, cyber security posture and security settings.
4. Confirm that performance and data security considerations are reflected in contractual obligations and oversight arrangements, including reporting by the provider, incident notification requirements, and rights to undertake testing.

Assurance

Boards should have confidence that the organisation has appropriate assurance in place for how AI risks are controlled within broader risk management settings.

External assurance or audit can play a crucial role in strengthening an organisation's management of AI risks through providing independent verification of control effectiveness. Increasingly, this assurance and testing will cut across key digital risk areas, such as cyber security and data governance, given the interconnected nature of the digital systems and infrastructure. When qualified third parties evaluate security protocols, access controls, compliance measures, and disclosure practices, they bring both objectivity and specialised expertise that internal teams may lack.

Boards of smaller organisations can consider how external expertise and assurance can be used in a targeted and cost-effective manner, for example by testing how sensitive datasets are used within internal AI systems.

AI policies and processes

The board should have confidence that the organisation has the key policies and processes needed for effective AI governance.

AI inventory or register

The AI inventory or register should track all approved AI systems or tools, together with key information about them. This provides a comprehensive view of AI use across the organisation and will assist with board monitoring and oversight. There should be a process to regularly update the inventory or register.

As AI features and tools are increasingly embedded within broader software and digital platforms, it may be appropriate for the AI functionality to be highlighted or accounted for in the organisational digital and data inventory or register.

AI policy

The AI policy should provide an overview of key elements of the organisation's AI governance framework. Generally, it will be the central policy document setting out the rules and processes relating to the use of AI across the organisation.

The AI policy should guide staff on the acceptable use of AI systems, including which AI systems are approved for use, the rules applying to them, and which systems have been prohibited. In some organisations, these requirements may be included in a separate AI acceptable use policy.

Depending on the size and complexity of the organisation, it may be appropriate for the AI policy to be approved by executive-level AI governance structures. The AI policy will generally be 'owned' by, or be the responsibility of a senior manager, or managers, accountable for AI within the organisation.

The NAIC's AI policy guide and template provides a practical starting point for organisations when developing an AI policy.

Boards of SMEs and NFPs – AI Governance Practices

- Oversee the development of a practical, accessible AI policy that sets clear boundaries for AI use, including restricting the use of unapproved tools with sensitive organisational data.
- Discuss potential risks for your organisation in relation to AI, including the process for patching cyber and data vulnerabilities.
- Ensure existing risk controls are reviewed to assess whether they effectively manage AI-related risks, such as privacy, data governance and cyber security.

Privacy and transparency

Organisations should review their privacy policies and processes to ensure they address specific risks relating to the use of AI systems. For example, it is increasingly common for customers and clients to provide personal information when interacting with AI systems, such as AI agents, including information such as their address. Separately, organisations may use personal information to train in-house AI models. The use of personal information in these ways may have privacy implications.

As highlighted in the section The role of the board and the regulatory landscape, from December 2026, organisations subject to the Privacy Act that use automated decision-making (ADM) systems, including AI systems, must meet new disclosure requirements. Broadly, organisations will need to explain when ADM systems are used, what decisions those systems make, and what personal information is used.

For some organisations, it may be appropriate to go beyond minimum regulatory disclosure requirements and provide a more complete picture of how AI is being used. For example, CBA's 2025 report Our Approach to Adopting AI is a notable example of transparency that provides detail on AI adoption, governance and risk management within CBA.

A broader approach to transparency can have the advantage of building stakeholder trust and strengthening reputation. Given broad societal misgivings about AI, a more fulsome approach to transparency can convey credibility and ethical leadership.

Regulatory obligations in respect of AI and privacy are expected to continue to evolve and it is imprtant that organisations monitor developments.

Questions for directors to ask

1. What steps are we taking to be confident that we are meeting our legal and regulatory obligations for the use of AI and associated data collection, storage, and use?
2. Does our risk management framework adequately address our AI-related risks? Does it differentiate between high-risk and low-risk AI applications?
3. What controls are in place to manage agentic risks?
4. Do our existing privacy, data governance, cyber and procurement policies address AI?

Governance red flags

1. The organisation's risk management framework does not address AI risks.
2. There is no process to assess and evaluate the risks of any AI enabled application or system.
3. Agentic AI systems have been adopted without rigorous risk identification and controls.
4. Data governance and privacy controls and policies have not been reviewed to ensure they are fit for purpose for AI systems.

Case study: Atlassian's Responsible Tech Review

Atlassian, a global software company founded in Australia, uses a Responsible Tech Review Template that is completed by teams prior to the deployment of AI use cases. Reflecting Atlassian's iterative approach to to internal management-led governance, the template has undergone several updates. Updates were driven by staff feedback and an examination of how staff were completing the template in practice. Two lessons from that work stand out.

First, lengthy and complex forms were a barrier to staff completion. A shorter, simpler template was more likely to be completed accurately and without requiring significant assistance from the responsible tech team.

Second, as the use of AI expanded across the organisation, responsibility for completing these assessments extended to a broad range of teams and business functions. This required the template to be accessible and understandable to a wider audience that may not have deep familiarity with the underlying AI technology.

In response, Atlassian refined the questions to use clear, simple language, and shortened the template to improve usability. Importantly, Atlassian was confident that simplifying the template would not compromise the quality of the assessment. This reflected a high level of trust in staff capability, supported by prior investment in AI literacy and the development of a strong culture of responsible technology use.

AI governance enablers

The board should have confidence that the organisation has the data governance, technology platforms and necessary organisational culture to support effective and responsible AI use.

Key points

- A pre-condition of significant organisational investments in AI is enhancing the underlying digital, data and technology infrastructure.
- An organisational culture that embraces AI knowledge, literacy and responsible use is critical to successful AI implementation and adoption.
- Organisations need to be alive to the impact of AI on both internal and external stakeholders, with the board playing a particularly key role in overseeing the impact on employees and the associated support structures.

AI technology and supporting infrastructure

Successful AI enablement depends on robust underlying digital infrastructure, including computing capacity, data storage, and network capabilities, which boards should understand at a strategic level.

Management should provide to the board assurance that the organisation's infrastructure can support current and planned AI workloads. It is important that management assesses how demand for computing capacity from AI workloads, or 'compute demand', can scale rapidly, and non-linearly, and the implications for capital allocation and operating expenditure.

Boards should also oversee how data storage and management practices support effective AI use, including the availability, accessibility and security of high-quality datasets required to train and operate AI systems.

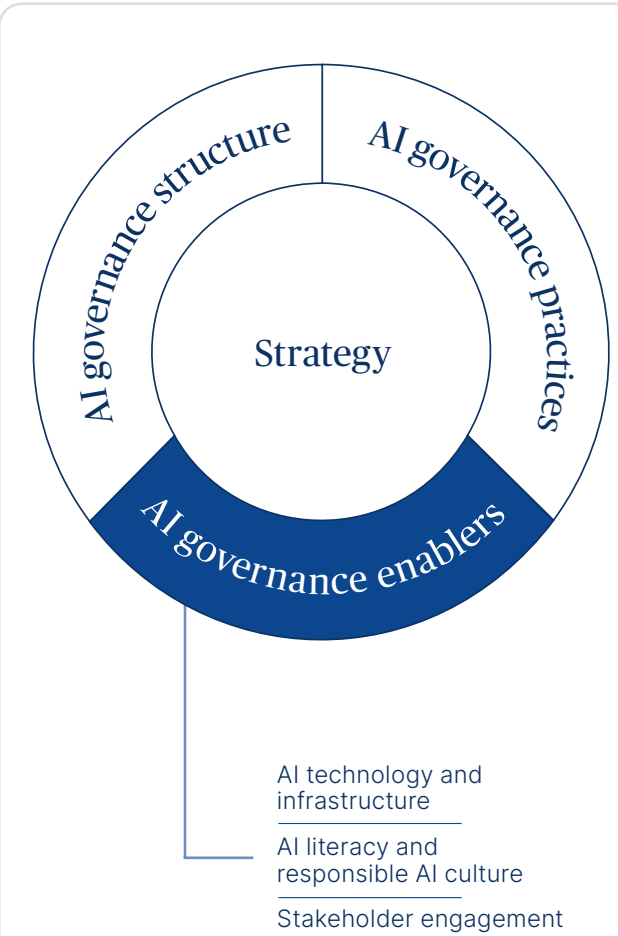

Figure 4: AI Technology and Supporting Infrastructure

AI literacy and a responsible AI culture

AI literacy and a responsible AI culture refer to organisational steps taken to promote an understanding of AI technology and support its effective and responsible use. In any organisation, regardless of size or type, where there is a significant AI adoption, the board should gain comfort that employees are being supported and trained to use the new AI tools in ways which are consistent with the organisation's policies.

A responsible AI culture in an organisation is one in which all employees (and volunteers in the case of many NFPs) feel personally committed to the responsible use of AI and reinforce this among their colleagues. This is particularly important given the high incidence of 'shadow AI' within organisations.

Depending on the organisaton, the board may have a role in setting a 'tone from the top' in building AI literacy and adopting AI systems to improve governance practices.

Stakeholder engagement

Engaging stakeholders is important to understanding the potential risks and harms associated with AI systems. The board should be comfortable that management has appropriately engaged with stakeholders when deploying AI systems and tools. Depending on the organisation and the nature of the AI deployment, this may include:

- **Customers and clients:** Transparency about how AI will be deployed, including the use of personal information. This should be undertaken in a manner consistent with Privacy Act requirements. In some instances, it may be appropriate for an organisation to undertake targeted consultation with customers or clients where there are concerns about AI systems and outputs. For example, a human services or care-sector NFP on how AI may be deployed in service delivery. For organisations with large customer bases, this is unlikely to be feasible.
- **Employees:** Widespread adoption of AI systems should be accompanied by consultation, training and support for employees. Employees play a critical role in successful AI implementation and value realisation, yet HTI's Invisible Bystanders research found that worker engagement is often limited. Organisations may also have obligations under relevant industrial relations laws to consult with employees on widespread AI implementation that materially affects the nature of work or employment.
- **Suppliers**: Depending on the commercial arrangements with suppliers, it may be appropriate to consult with these stakeholders on AI implementation. Where AI systems alter the nature of the relationship or use supplier data, transparency may be important.

- **Broader stakeholders**: Many larger organisations have stakeholders that extend beyond customers and employees to include the broader community, government and regulators. These stakeholders can be key to reputation and social licence. It may be appropriate for the organisation to take proactive steps to consult with these stakeholders on AI system trials and implementation, particularly where they may affect a large number of customers and employees.

Boards of SMEs and NFPs – AI Governance Enablers

- Consider whether existing technology and data systems effectively support AI deployment, including whether improvements to data collection and quality are needed.
- Form a view of internal capability to effectively use AI and invest in foundational AI training for leadership, staff, contractors and volunteers, where appropriate.
- Engage with employees, volunteers, clients and stakeholders to understand how AI use may affect them, and identify internal champions who can advocate for effective, responsible AI use across the organisation.

Questions for directors to ask

1. Do we have a strategy or process for understanding the impact of our AI implementation on internal and external stakeholders and have we consulted where necessary?
2. Do we have a comprehensive view of our underlying data volume and quality and how that will influence our AI effectiveness?
3. Have our cyber security and data governance processes, practices and policies been updated for AI implementation?
4. Are our current technology platforms and data flows fit-for-purpose to enable safe and responsible data and AI governance? I.e. Does the system architecture enable transparency or explanation of decisions?

Governance red flags

1. Management is unable to provide an account of how AI implementation and use will impact our key stakeholders.
2. There is no data stocktake or inventory and limited understanding of internal data quality.
3. No or limited training and support for employees on AI system use and risks.
4. Weak understanding of the necessary technological infrastructure to support AI system implementation and use.

Case study: Canteen - listening to the voice of members

As a national not-for-profit organisation supporting young people affected by cancer, Canteen ensures that young people drive their organisation. The majority of Canteen's Board of Directors are young leaders, with five Member Directors (young people) sitting alongside four Associate Directors who are volunteers bringing relevant expertise to the board.

Listening to the voices of its members is at the heart of Canteen's work. In response to calls from members to embrace AI to improve services for their community, Canteen established a time-bound management-led AI Sprint Committee as a board-linked governance and innovation body. For 12 months, the committee was tasked with setting guardrails, clarifying accountabilities, and integrating AI oversight into Canteen's existing governance practices, while identifying appropriate opportunities for innovation.

Structured, frequent consultation with staff and young people is built into the committee's work, so that members continue to have a voice in shaping how AI is used. Together, these mechanisms keep AI use aligned with Canteen's mission.

Case study: Building Westpac senior leadership's AI literacy

Westpac, one of Australia's largest banks, has placed executive AI literacy at the centre of its AI governance efforts, recognising that effective oversight depends on informed leadership.

A key initiative is its Executive AI Council, a strategic management forum designed to build awareness and fluency in AI across the executive team. Meeting monthly, the council brings together the CEO and the executive team but is deliberately not a risk or governance forum. The council exists to demonstrate, share and accelerate cultural change. Each session features practitioners from across the business, including developers and product owners, presenting AI use cases directly to the executive team. The format is intentionally hands-on, with live demonstrations rather than slide decks, and questions from the executive team focused on the practitioner's journey rather than just the technical detail.

The board is engaged with equal deliberation. The board receives a quarterly update on AI but is also building its AI literacy alongside the executive team through a structured AI education series. The sessions have been sequenced to build understanding cumulatively, starting with sessions on culture and transformation, then responsible AI, and then a technical grounding in AI itself. Bringing board and executive education together allows both groups to develop a shared baseline understanding of the technology and the opportunities and risks for the bank.

This top-down investment in capability signals that AI is a strategic priority and supports more informed decision-making on strategy, risk and governance. By building AI literacy at the board and executive level, Westpac is strengthening its ability to oversee AI use and reinforcing a culture where responsible AI is embedded across the organisation.

Case study: Telstra - AI Governance Operating Model in action

This case study outlines how Telstra, Australia's largest telecommunications company, has applied the four components of HTI's AI Governance Operating Model.

Strategy

Telstra's approach to AI is not guided by a standalone AI strategy separate from the rest of the business. Instead, AI is treated as a core enabler of the organisation's corporate strategy. AI is included in Telstra's strategy as a means of achieving its overarching goal of meeting customers' connectivity needs.

AI Governance Structures

Telstra has a dedicated management-led AI governance body responsible for overseeing and approving high-impact AI use cases. Previously, this was the Risk Council for Data and AI (RCAID). However, concerns emerged regarding RCAID's overlap with existing functions and its scalability, given the speed and consistency of its processes.

In response, RCAID was redesigned as the AI Risk Oversight Council (AIROC). AIROC has simplified and faster processes. By removing overlapping checks with existing privacy and cyber security controls, the council focuses solely on AI-specific risks.

AI Governance Practices

Telstra has focused on embedding governance into everyday practice. As AI adoption expanded across the organisation, Telstra simplified the language used in its policy, templates, and processes to make them more accessible to non-technical audiences. It removed ambiguous terminology, rewrote templates in plain English, and developed a single platform and end to end workflow to guide staff through AI use case approval.

AI Governance Enablers

As AI capability has spread across Telstra, the organisation has strengthened the enablers that support responsible AI at scale, including training, staff engagement and culture. For example, Telstra made completion of its Copilot training and Responsible AI module a prerequisite for Microsoft 365 Copilot access. This resulted in more than 18,000 completions, helping build responsible AI capability across the organisation.

For the full case, see HTI's AI Governance Lighthouse Case Study: Telstra (2026)

Part 3 - Measuring AI Returns

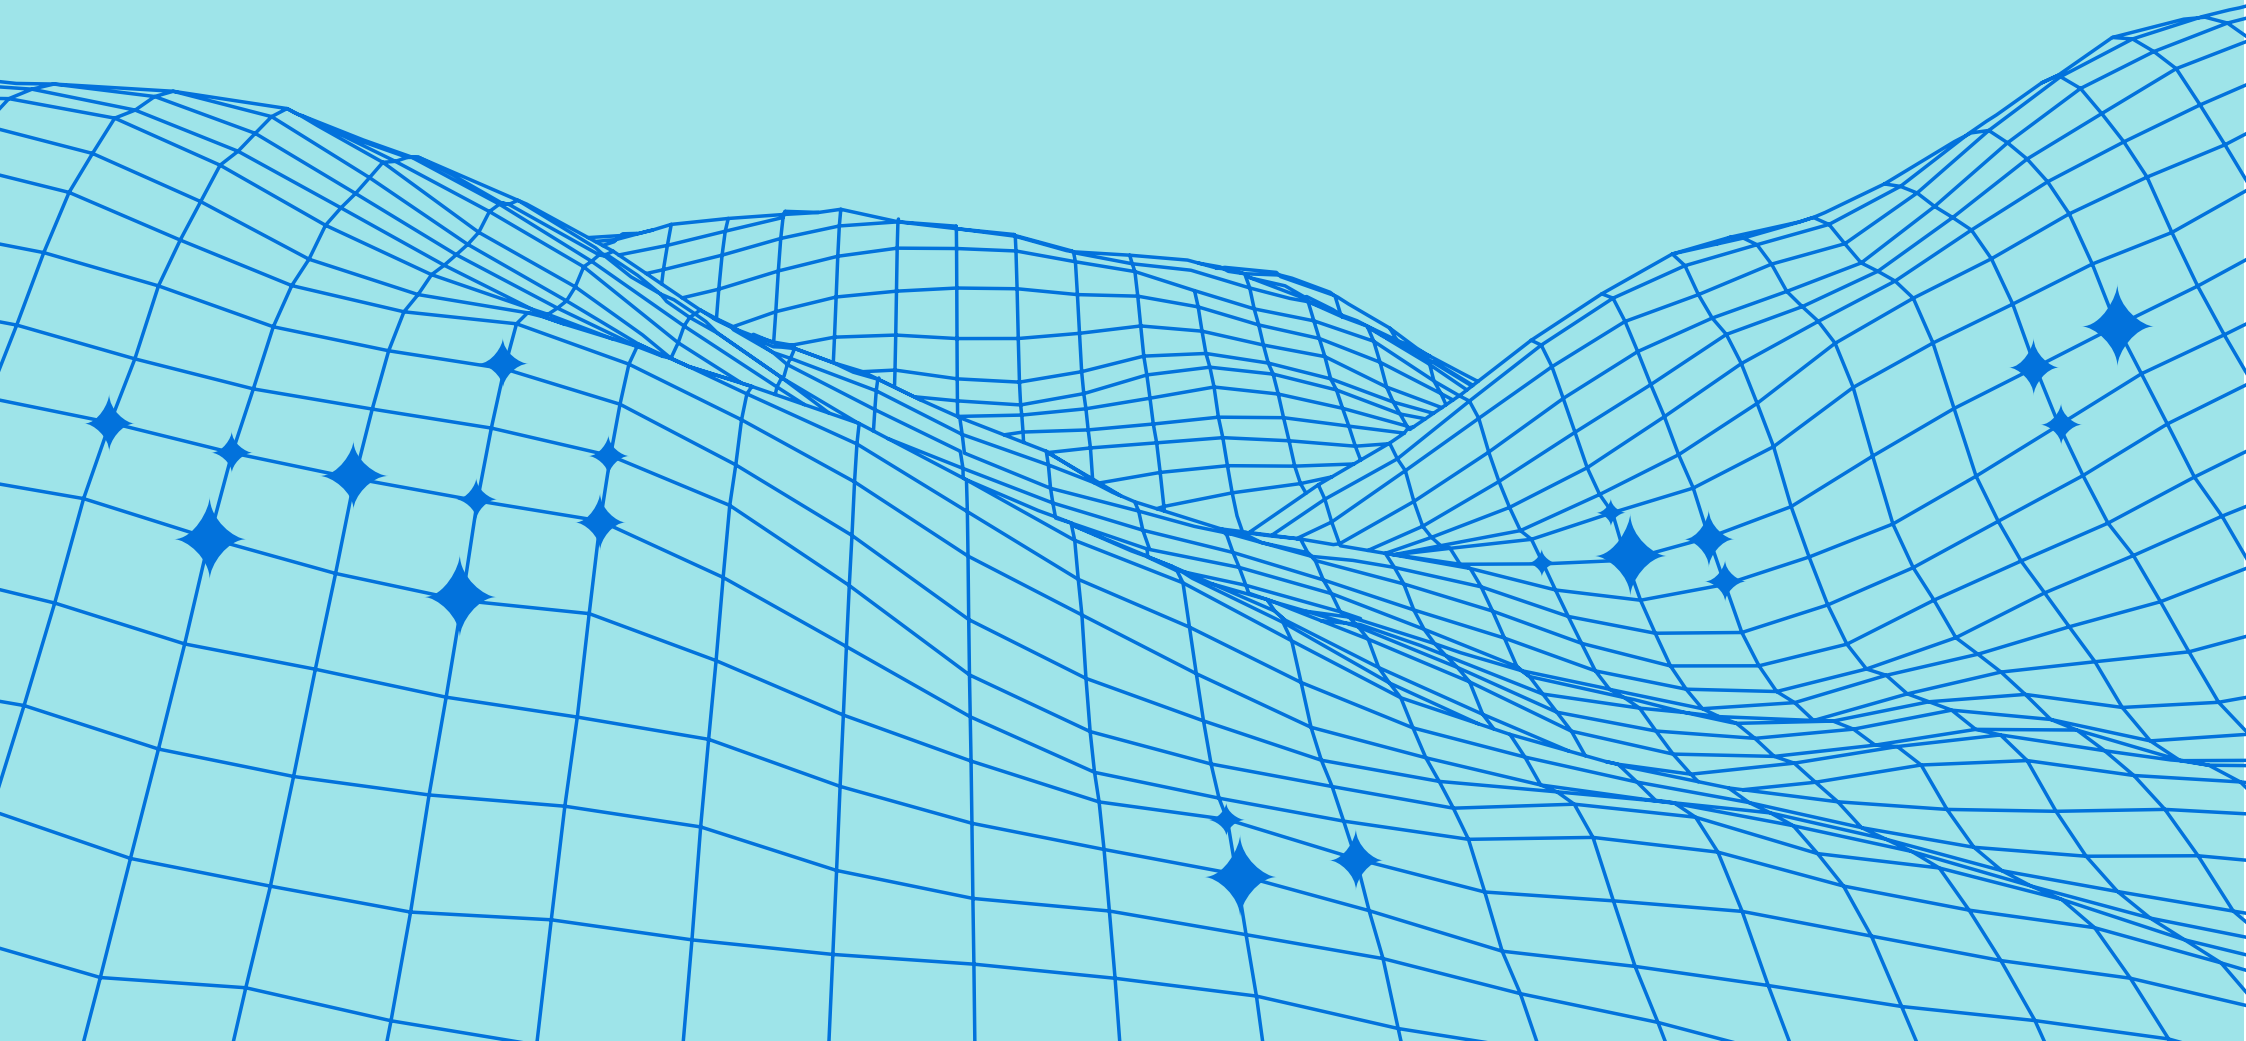

Measuring value and returns from AI

Key points

- Greater value is realised where AI is used to transform processes, systems or business models, yet relatively few organisations have moved beyond the pilot phase in these strategies.
- AI ROI assessment should focus on specific, measurable use cases and test if the investment has been realised in terms of organisational efficiency and improved outputs.

AI has the potential to deliver significant benefits for organisations, but realising those benefits is not automatic. Boards therefore have an important role in overseeing whether investments in AI are generating sustainable value.

The role of the board in AI ROI

Boards play an important role in overseeing whether the organisation's investments in AI are generating sustainable value. This is important because many organisations are making significant investments in AI capability in an environment where evidence on returns remains mixed.

Board oversight of return on investment from AI serves three important functions:

1. **Supports oversight of strategic and technology investments**. Boards should understand what value AI investments are expected to deliver, over what timeframe, and how success will be measured.
2. **Promotes organisational discipline and focus**. Requiring clear use cases, defined success metrics and structured evaluation processes can help avoid fragmented AI adoption that consumes resources without delivering meaningful outcomes.
3. **Value and risk in AI deployment are closely connected**. AI systems that create significant commercial, regulatory or reputational risks may ultimately erode rather than create organisational value. Boards should consider whether AI benefits are being achieved in a manner that is safe, lawful and aligned with the organisation's risk appetite.

The challenge of measuring AI returns

Evidence on the returns from AI investments remains mixed. Some studies have found that 80% of organisations have not seen a tangible impact on profitability from generative AI investments and that 95% of AI pilots fail.

Yet, many organisations are reporting broader benefits. According to KPMG Global AI Pulse Q1 2026 survey, 64% have seen real business value, which rises to 82% among 'AI leaders'. In Australia, research from Deloitte found that 61% of companies reported efficiency improvements from AI use.

Early benefits of generative AI often appear as individual productivity, yet these do not automatically translate into bottom-line impact. Greater value is typically realised where AI is used to transform processes, systems or business models, yet relatively few organisations are currently operating at this level. In Australia, research from Deloitte found that 30% of businesses report using AI to transform their processes, compared to 34% globally.

At the same time, the pressure to demonstrate returns is intensifying. According to the Kyndryl Readiness Report, 61% of business leaders reported increasing pressure to demonstrate returns on AI investments in 2025.

A structured framework for AI return on investment

The following framework outlines key steps organisations, and ultimately boards, can use to assess the return on AI investments.

- **Focus on assessing ROI for a specific, targeted use case**, rather than a general capability or system. For each use case, organisations should identify the specific business improvement being sought and the metrics that will demonstrate success.
- **Approach ROI with the implementation phase in mind**. AI initiatives typically move through three phases: experimentation (can we do this?); integration (does it work?) and scaling (can this be rolled out?). Each phase has a distinct ROI profile with different benefits, costs and risks rising or falling in importance.
- **Attempt to take a broad view of the benefits** so the organisation can gain a clearer understanding of the range of direct and indirect benefits that AI investments achieve over time.
- **Test assumptions** at the integration and scaling phases. Three factors are particularly important to test:
  a. adoption – what proportion of the target user base is actively using the capability?;
  b. translation efficiency – how well does the targeted business improvement translate into broader organisational value?; and
  c. accuracy and performance – how well does the AI use case perform, and can we rely on it to deliver the quality and consistency required?
- **Ensure all relevant costs are modelled**, including token costs associated with usage and process to train staff and provide effective quality assurance.

This framework can be applied by organisations of any size. However, the level of evaluation and governance effort should be proportionate to the opportunity, cost, risk and complexity of the deployment.

Over the horizon

Since the first version of this publication in 2024, the pace of AI technological advancements has accelerated rapidly. AI has evolved from largely standalone tools towards more integrated, autonomous systems that are increasingly embedded within software and digital products and services.

In the coming years, advances in agentic AI, embodied AI and robotics may enable systems and machines to perceive, reason and act in dynamic settings with 'digital workers' operating alongside human teams. Examples may include:

- Widespread use of individual agents to undertake tasks on the behalf of individuals and organisations, for example online shopping or interacting with trusted merchants.
- Significant penetration of autonomous vehicles in Australia.
- Increased prevalence of robotics in sectors of the economy beyond manufacturing.
- AI-driven energy solutions focused on the efficient use of energy generation and infrastructure, including renewable sources.
- An embrace of AI-powered diagnostics, personalised healthcare and rapid advancements in drug development.
- Both an increase in cyber security vulnerabilities and improved defences to combat threats, including identity theft.

These trends may be supercharged by advances in quantum computing.

As AI moves from a decision-support role to one that actively shape outcomes for people and organisations, questions of accountability, oversight and responsibility may become more complex.

The key ongoing role of the board

A central message of this resource is that boards play a key role at all Australian organisations in overseeing the effective, responsible and safe deployment of AI. This role will not change despite the rapid advances in AI that are likely to occur in years to come.

Boards need to continue to work closely with management to carefully consider the stakeholder, regulatory and governance implications of AI systems specific to their organisations and their industry. They also must stay engaged with the shifting landscape around AI regulation and governance.

Whatever the future of AI regulation brings, there are already a broad range of existing legal obligations that apply to an organisation's use of AI systems.

By investing in the governance of their organisations, directors will help them deploy AI systems safely, responsibly and strategically so that Australia will benefit sustainably from this technology.

Appendix A: Regulatory obligations and AI

Current as at June 2026

A range of existing laws are relevant to the development and use of AI systems. The Australian Government has foreshadowed further reform of some of these laws where necessary to address emerging AI-related issues.[9]

Privacy

The *Privacy Act 1988* (Cth) (Privacy Act) applies to the collection, use, storage, disclosure and destruction of personal information. This includes where personal information is used to train, test or use an AI system. Personal information is a broad term that includes information that could identify an individual. The Privacy Act applies to all organisations unless exempt under the small business exemption.

The Office of the Australian Information Commissioner has released guidance on the application of the Privacy Act to AI systems.[10] The Privacy Commissioner has also made a number of rulings on the use of AI in the form of facial recognition technology in the retail sector.[11]

Since 2025, the Privacy Act includes a mechanism to address serious invasions of privacy, in the form of a statutory tort.[12] Depending on the circumstances, a person who alleges that their privacy has been seriously invaded by another person (including an organisation) can seek to bring proceedings for breach of this tort. In practice, this could include a serious invasion of privacy using an AI system.

From 10 December 2026, organisations covered by the Privacy Act will be required in certain circumstances to include details in their privacy policies about the use of computer programs to make, or to do a thing substantially and directly related to making, a decision. This will be required if the decision could significantly affect an individual's rights or interests, and the computer program uses their personal information.[13]

Consumer protection

The provision of AI-enabled goods or services to consumers in trade or commerce is subject to the Australian Consumer Law (ACL), which is in Schedule 2 to the *Competition and Consumer Act 2010* (Cth). The ACL prohibits misleading or deceptive conduct, unconscionable conduct, and false or misleading representations in trade or commerce.
It also contains consumer guarantees and a product liability regime that applies to defective goods.

In 2025, the Treasury reviewed the ACL and found that, when considered in combination with other relevant laws, it is broadly capable of adapting effectively to AI-enabled goods and services.

In October 2025, the ACCC took enforcement action against Microsoft Australia under the ACL. The ACCC alleged that Microsoft had misled about 2.7 million Australian customers when communicating subscription options and price increases after integrating Copilot (its AI assistant) into Microsoft 365 plans.[14]

More recently, the ACCC has observed that AI is increasingly being used in customer service with the retail sector currently using AI at higher rates than other sectors. It advised that businesses using AI chatbots should ensure that the responses they provide to consumers do not provide information that may be contrary to their obligations under the ACL.[15]

Organisations may also mislead consumers if they engage in 'AI-washing'; that is, where they make misleading or overstated claims about the AI capabilities of a product. Both the ACCC and ASIC have drawn attention to the practice of AI-washing.[16]

---

[9] See Department of Industry, Science and Resources, National AI Plan (Document, 2 December 2025) 28 [Action 7, Mitigate Harms].
[10] See OAIC, AI: Guidance on privacy and developing and training generative AI models (Webpage, 23 October 2024); Office of the Australian Information Commissioner, Guidance on privacy and the use of commercially available AI products (Webpage, 17 January 2025).
[11] See *Bunnings Group Limited and Privacy Commissioner* [2026] ARTA 130; Gilbert + Tobin, 'Administrative Review Tribunal upholds Bunnings' use of facial recognition technology' (Webpage, 6 February 2026).
[12] Privacy Act 1988 (Cth), sch 2.
[13] See *Privacy and Other Legislation Amendment Act 2024* (Cth), sch 1 pt 15.
[14] ACCC Recent developments in artificial intelligence: Industry snapshot (Report, December 2025).
[15] Ibid
[16] See ACCC, Recent developments in artificial intelligence: Industry snapshot (Report, December 2025); ASIC, Submission No 67 to Select Committee on Adopting Artificial Intelligence (AI), Inquiry into the uptake of AI technologies in Australia (May 2024).

Anti-discrimination

The outputs of AI systems can directly or indirectly discriminate against individuals on the basis of protected attributes. For example, this may occur where their training data reflects historical biases. Commonwealth and State anti-discrimination laws prohibit direct and indirect discrimination based on protected attributes such as race, sex, age and disability. Organisations should ensure that their use of AI systems does not result in unlawful discrimination.

Industrial relations

The *Fair Work Act 2009* (Cth) is relevant to how AI is deployed and used in Australian workplaces. The legislation regulates the employment lifecycle and imposes obligations that are triggered when new technologies materially affect employees, including requirements to consult on major workplace changes (such as the introduction of AI systems). Organisations are also required to comply with unfair dismissal, redundancy and procedural fairness protections, and adhere to modern award or enterprise agreement consultation clauses. All these requirements are likely to be relevant to considerations on AI deployment and the impact on employees.

Work health and safety

The use of AI systems in the workplace can create risks of physical or psychological harm to employees. Australia has harmonised work health and safety laws that impose a primary duty of care to ensure, so far as is reasonably practicable, the health and safety of workers and other persons.[17] The duty applies when organisations develop, procure or deploy AI systems, and measures may include incorporating AI related risks into WHS risk assessments, consultation and training.

In 2026, the NSW Parliament amended the *Work Health and Safety Act 2011* (NSW) to include a new duty to ensure, so far as is reasonably practicable, that workers are not exposed to health and safety risks arising from the allocation of work by a digital work system.[18]

Duty of care

Organisations may have a duty of care towards people who use or are impacted by an AI system. The law of negligence requires that where an organisation has a duty of care to a class of persons, the organisation must exercise the standard of care of a reasonable person in the circumstances to avoid reasonably foreseeable harm. A failure to do so may result in liability for loss or injury suffered as a result of the breach of that duty. Manufacturers, distributors and retailers may owe duties of care in negligence in relation to AI-enabled products, and may also be subject to the product liability obligations under the ACL.

Environmental laws

Australia's mandatory climate reporting regime is subject to phased implementation and applies to entities that are required to prepare and lodge financial reports under Part 2M of the *Corporations Act 2001* (Cth) and meet the relevant thresholds.[19]

Organisations that are covered by the regime must prepare climate-related disclosures in accordance with AASB S2 Climate-related Disclosures. These must be included in a dedicated Sustainability Report, as part of the organisation's annual report.

AI-related emissions are likely to be captured under Australia's climate reporting regime indirectly through scope 2 and scope 3 disclosures – particularly in relation to electricity consumption from data centres and cloud services. This will require entities subject to the regime to assess and report material emissions associated with their use of computational infrastructure and third-party AI providers.

Cyber security

Cyber security is a key consideration for organisations developing and deploying AI, given AI's reliance on data and the increased risk of cyber security breaches.

The *Security of Critical Infrastructure Act 2018* (Cth) establishes a framework to manage risks to Australia's critical infrastructure by regulating specified assets across key industry sectors. It imposes obligations on responsible entities for critical infrastructure assets in relation to cyber and information security risk management, and cyber incident reporting.

The *Cyber Security Act 2024* (Cth) mandates minimum cyber security standards for certain products that directly or indirectly connect to the internet; introduces mandatory reporting of ransomware payments by specified entities; and establishes a voluntary, limited use information sharing framework in response to significant cyber security incidents.

Copyright

The *Copyright Act 1968* (Cth) protects original works created by human authors. The use of copyright-protected material in connection with the training or use of AI systems may involve acts of reproduction and infringe copyright – unless the organisation has the consent of the copyright holder, relies on a licensing arrangement or the use falls within a fair dealing exception.

---

[17] See, e.g. *Work Health and Safety Act 2011* (NSW).
[18] *Work Health and Safety Amendment (Digital Work Systems) Act* 2026 (NSW).
[19] Reporting is being phased in based on an entity's employee size, consolidated gross assets and consolidated revenue. Group 1 entities (the largest emitters and corporations) have been required to disclose from 1 January 2025. Smaller Group 2 and Group 3 entities are being phased in from 1 July 2026 and 1 July 2027, respectively. Not-for profits (NFPs) that meet the size thresholds are also included.

Appendix B: Resources

Party	Resource
Federal Government	NAIC, Essential AI practices
	NAIC, AI and Australian law
	NAIC, Planning for AI
	ASD, Artificial intelligence for small business: Managing cyber security risk
	OAIC, Guidance on privacy and the use of commercially available AI products
AICD	AI use by directors and boards: Early insights
	Cyber Security Governance Principles Version 2 (in partnership with the CSCRC)
	Data Governance Foundations for Boards (in partnership with Allens and Melbourne Business School)
	Effective board minutes and the use of AI (in partnership with the Governance Institute of Australia)
	AI Fluency for Directors Sprint (short course)
	AI Governance for Directors Webinar Series
	Ethics in the Boardroom 2nd edition (in partnership with The Ethics Centre)
	Cyber Security for Directors (short course)
HTI	AI Governance Lighthouse Case Study: Telstra
	AI Governance Operating Model
	Designing AI Governance Structures
	People, Skills, and Culture for Effective AI Governance
	From Invisible to Involved: A Guide to Worker Engagement on AI
	'Invisible Bystanders': How Australian Workers Experience the Uptake of AI and Automation
Business Council of Australia	Adopting AI in the Workplace
Governance Institute of Australia	Governance in the Age of Agentic AI
Commonwealth Bank of Australia	Our Approach to Adopting AI

Appendix C: SME and NFP director checklist

Governance element	Board considerations
AI Strategy	• Understand where AI is already being used in the organisation, including AI features in third-party products and services. • Discuss the organisation's risk appetite for AI use with management. • Understand whether new resources and significant investment is required to support AI deployment. • Assess where AI could add value and support strategic objectives, including opportunities to enhance existing processes and resources (such as current SaaS subscriptions) in a cost-effective way.
AI Governance Structure	• Assign accountability for AI oversight to an appropriate senior leader. • Establish a proportionate process for approving the use of AI systems in line with their potential impact on the organisation and stakeholders. • Establish when and how the board will be updated or consulted on AI systems, including working with management to develop targeted metrics on AI use, effectiveness, and risks. • Consider how external assistance may help AI deployment in a cost-effective way.
AI Governance Practices	• Oversee the development of a practical, accessible AI policy that sets clear boundaries for AI use, including restricting the use of unapproved tools with sensitive organisational data. • Discuss potential risks for your organisation in relation to AI, including the process for patching cyber and data vulnerabilities. • Ensure existing risk controls are reviewed to assess whether they effectively manage AI-related risks, such as privacy, data governance and cyber security.
AI Governance Enablers	• Consider whether existing technology and data systems effectively support AI deployment, including whether improvements to data collection and quality are needed. • Form a view of internal capability to effectively use AI and invest in foundational AI training for leadership, staff, contractors and volunteers, where appropriate. • Engage with employees, volunteers, clients and stakeholders to understand how AI use may affect them, and identify internal champions who can advocate for effective, responsible AI use across the organisation.

Appendix D: Glossary

Term	Description
Artificial intelligence (AI)	A set of technologies that enable machines to perform tasks that typically require human intelligence, such as analysing data, generating content, making predictions or supporting decisions.
Agentic AI	AI systems capable of executing tasks with a degree of autonomy, often coordinating actions across systems or workflows to achieve an objective.
ACCC	Australian Competition and Consumer Commission
ACL	Australian Consumer Law
ADM	Automated decision-making
AASB	Australian Accounting Standards Board
AI6 (National AI Centre guidance)	A set of six essential practices developed by the National AI Centre to support the safe and responsible development and use of AI systems. These practices focus on accountability, impact assessment, risk management, transparency, testing and monitoring, and maintaining human control.
AI assurance	Processes that provide confidence that AI systems are operating as intended and within defined risk parameters, including testing, validation and audit.
AI governance	The structures, processes and controls used to oversee the development, deployment and use of AI systems within an organisation.
AI inventory (or AI register)	A consolidated record of the AI systems used within an organisation, including information about their purpose, design, risks and associated controls.
AI literacy	The knowledge and capability required to understand AI systems, their limitations and their potential risks and benefits.
AI risk	The potential for AI systems to create adverse outcomes, including operational, regulatory, ethical, reputational or societal harm, or to amplify existing risks.
AI risk appetite	The level and type of AI-related risk that an organisation is willing to accept in pursuit of its objectives, as set or endorsed by the board.

AI strategy	The organisation's approach to using AI to support its objectives, including how AI investments are prioritised, governed and aligned with broader strategy.
AI system	A machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment.
AI token	A small unit of text – such as a word, part of a word, or punctuation – that an AI model processes and generates to understand and produce a result.
APRA	Australian Prudential Regulation Authority
ASD	Australian Signals Directorate
ASIC	Australian Securities and Investments Commission
Bias (in AI systems)	Systematic and unfair outcomes resulting from skewed data, model design or deployment decisions, potentially disadvantaging individuals or groups.
Data governance	The framework for managing data across its lifecycle, including data quality, security, accessibility and compliance with legal and regulatory requirements.
Digital infrastructure	The computing, data storage, networking and cloud systems that support AI systems and broader digital operations within an organisation.
Explainability	The extent to which the behaviour or outputs of an AI system can be understood and explained to stakeholders and regulators.
Generative AI	AI systems that create new content – such as text, images or code – based on patterns learned from large datasets.
Human oversight	The involvement of people in monitoring, reviewing or intervening in AI system outputs or decisions.
Machine learning (ML)	A subset of AI in which systems are trained on data to recognise patterns and improve performance over time without being explicitly programmed for each task.
NAIC	National AI Centre
Responsible AI	An approach to AI that emphasises lawful, ethical and transparent use, aligned with organisational values and societal expectations.
SaaS	Software-as-a-service
Shadow AI	The use of AI tools by employees without formal approval, governance or oversight by the organisation.
Traditional AI	AI systems that analyse data to identify patterns, make predictions or optimise processes, typically within well-defined tasks. It is also referred to analytical AI.

Acknowledgements

The AICD acknowledges the contribution of the guide's co-authors: Gaby Carney, Llewellyn Spink and Professor Nicholas Davis from HTI, and Simon Mitchell from AICD.

This guide is based on the research and insights of HTI and draws on the experience of AICD members and other subject matter experts who participated in interviews and roundtable conversations.

We thank them for generously sharing their knowledge, experience and reflections.

We particularly thank Jacqueline Chow FAICD, Shirley Chowdhary FAICD, Lyn Cobley AM FAICD, Naomi Edwards FAICD, Steve Ferguson GAICD, Cheng Lim, Kee Wong FAICDLife and Noeline Woof FAICD for their review and input into this publication.

Australian Institute of Company Directors

Acknowledgement of Country

The Australian Institute of Company Directors and the University of Technology Sydney Human Technology Institute (HTI) acknowledge the Traditional Custodians of the Lands on which we are located and pay our respects to the Elders, past and present. We acknowledge the First Nations people across this Country and recognise their unique cultural and spiritual relationships to the Skies, Land, Waters, and Seas and their rich contribution to society.

About AICD

The Australian Institute of Company Directors (AICD) is committed to strengthening society through world-class governance. We aim to be the independent and trusted voice of governance, building the capability of a community of leaders for the benefit of society. Our membership includes directors and senior leaders from business, government and not-for-profit sectors.

About HTI

The Human Technology Institute (HTI) at the University of Technology Sydney 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.

Disclaimer

This document is part of a series of tools and resources provided by the AICD. It is intended as a general guide only and should not be relied upon as a substitute for professional advice. While care has been taken in its preparation, the AICD and HTI do not warrant the accuracy, reliability or completeness of the information contained in this document. To the extent permitted by law, the AICD and HTI exclude all liability for any loss or damage arising out of the use of this document.

For more information:

T: 1300 739 119
E: policy@aicd.com.au

Find us at aicd.com.au