AZUREIQ ADVISORY

WHITE PAPER

AI Governance & the Australian Regulatory Horizon

What Every Board and C-Suite Leader Must Act On Now

IN THIS PAPER

AUGUST 2026

Executive Summary

Australia's AI (artificial intelligence) governance landscape has shifted decisively. What was, until recently, a patchwork of voluntary guidelines and aspirational frameworks is now becoming binding law, with personal consequences for those at the top of organisations who fail to take it seriously.

Three changes are arriving in close succession:

- The Privacy Act's automated decision-making transparency obligations take effect on 10 December 2026, requiring APP (Australian Privacy Principles) entities to disclose where algorithms make or substantially influence decisions about individuals.
- AS (Australian Standard) ISO 42001, the Australian Standard for AI Management Systems, is already in force and is rapidly becoming the expected baseline for responsible AI governance, including in government procurement and regulatory assessments.
- Mandatory AI legislation is expected to be introduced to Parliament in early 2027, following the Prime Minister's announcement of a National AI Framework and the establishment of the new Office of AI in July 2026.

Against this backdrop, courts and regulators are making clear that existing directors' duties under the Corporations Act 2001 (Cth) already apply to AI governance, and that the traditional safe harbours provide less protection than boards have assumed. The Star Entertainment judgment (March 2026) and a series of

ASIC (Australian Securities and Investments Commission) enforcement actions have demonstrated that AI-related governance failures carry real personal liability.

> This paper is for boards and C-suite leaders who need to understand what is now required, not just what is coming. Each section ends with a practical board agenda item.

Part 1: AS ISO 42001 – The Standard That Is Already Here

What Is AS ISO 42001?

AS ISO 42001:2023 is the Australian adoption of the international standard ISO/IEC (International Organization for Standardization / International Electrotechnical Commission) 42001:2023, published by Standards Australia in February 2024. It specifies requirements for establishing, implementing, maintaining, and continually improving an Artificial Intelligence Management System (AIMS) within organisations.

The standard was developed through ISO's Technical Committee 42 (Artificial Intelligence), with direct input from Australia's mirror committee IT-043. It follows the same high-level structure (Annex SL) used by ISO 27001 (information security) and ISO 9001 (quality management), which means organisations already certified to those standards will find the architecture familiar.

Crucially, the standard is not limited to technology companies. Any organisation that provides or uses products and services that incorporate AI systems is within scope. That encompasses banks, insurers, health providers, government agencies, retailers, and professional services firms, in short, the overwhelming majority of Australia's large enterprises.

What Does It Require?

AS ISO 42001 is a process standard, not a technical checklist. It does not mandate specific AI architectures or algorithms. Instead, it requires organisations to establish documented, governed, and auditable management processes across the AI lifecycle. The core requirements span:

- **Risk.** AI Risk Assessment and Treatment
- **System Impact Assessments** before deployment of AI that may significantly affect individuals
- **Data Quality and Provenance Controls:** understanding where training and operational data comes from and how it may introduce bias
- **Human Oversight Mechanisms:** ensuring appropriate human review is built into AI-supported decision processes
- **Supplier and Customer Responsibility Allocation:** defining who is accountable when third-party AI tools are embedded in products
- **Performance Monitoring and Logging:** ongoing tracking of AI system behaviour, drift, and outcomes
- **Internal Audits and Management Reviews:** periodic assessment of the AIMS by qualified parties

Certification and Assurance

Certification to AS ISO 42001 is voluntary at present. Organisations can adopt the standard internally as a governance framework, or seek third-party assessment through a conformity assessment body. However, "voluntary" is doing less work than it used to.

EMERGING EXPECTATION

Australian Government agencies are increasingly referencing AS ISO 42001 in procurement requirements for AI-enabled services. Regulated industries including financial services and healthcare can expect APRA (Australian Prudential Regulation Authority) and AHPRA (Australian Health Practitioner Regulation Agency) to treat compliance with the standard as evidence of reasonable governance. Organisations that cannot demonstrate an AIMS will face growing disadvantage in tenders, regulatory interactions, and insurance assessments.

The standard also includes audit requirements that support external assessment, meaning the management review records organisations create for internal purposes can and will be examined by regulators, litigation counterparties, and third-party auditors. Getting the paper trail right matters.

The Relationship Between AS ISO 42001 and the Incoming Legislation

The mandatory AI standards framework expected in 2027 is widely anticipated to use AS ISO 42001 as its technical reference. Organisations that implement a compliant AIMS now will be positioned for mandatory requirements rather than scrambling after legislation passes.

The parallel is instructive: when the GDPR (General Data Protection Regulation) came into force in Europe, organisations with existing ISO 27001 programs adapted significantly more quickly and at lower cost than those starting from scratch. The same dynamic is expected in Australia's mandatory AI transition.

BOARD AGENDA ITEM

Ask management: Do we have an AI Management System in place that would withstand an external audit? If not, what is the roadmap to establish one, and what is the timeline? Request a gap assessment against AS ISO 42001 within the next board cycle.

Part 2: The Privacy Act's Automated Decision-Making Obligations

The Obligation and Its Deadline

The Privacy and Other Legislation Amendment Act 2024 (Cth) introduced new obligations under Australian Privacy Principles 1.7 and 1.8. These provisions take effect on 10 December 2026, less than four months away at the time of writing, and apply to all APP entities, which includes most private sector organisations with an annual turnover exceeding $3 million, as well as all federal government agencies.

The obligation is triggered when an APP entity "arranges a computer program, using personal information about an individual, to make or directly support a decision that could reasonably be expected to significantly affect the individual's rights or interests."

The Office of the Australian Information Commissioner (OAIC) has signalled a broad reading of this trigger. Automated assessment tools, scoring engines, recommendation systems, and AI-enabled workflows that influence staffing, lending, insurance, healthcare, or benefits decisions are all likely within scope, not just fully automated decisions, but those where automation performs a "substantially related function" to the decision.

What Organisations Must Do

By 10 December 2026, APP entities must update their Privacy Policy to disclose:

- The types of personal information used in automated decision-making technology operations
- The categories of decisions made solely by automated means
- The categories of decisions in which automation performs a substantially related function to the decision process

This sounds straightforward. In practice, it requires a systematic audit of every AI and rules-based system that touches personal information, an exercise that typically reveals far more automated decision-making than organisations realised they had. Credit scoring, fraud detection, customer triage, staff rostering, claims processing, and student assessment tools may all be in scope.

Enforcement Context

The OAIC is not waiting for December 2026 to act. In January 2026, the regulator initiated a privacy compliance sweep targeting approximately 60 organisations across six sectors, specifically examining automated decision-making practices. This is proactive enforcement: organisations can be in the crosshairs before the deadline has passed.

The penalty regime is significant. Under reforms effective since late 2024, maximum civil penalties for serious or repeated privacy interferences reach $50 million for bodies corporate. Importantly, the Privacy and Other Legislation Amendment Act 2024 also introduced a statutory tort for serious invasions of privacy, meaning individuals can now sue organisations directly, creating dual exposure (regulatory penalty plus civil litigation) from a single incident.

The Board's Role Under Privacy Law

Privacy law formally sits with the organisation as a legal entity, not individual directors. However, the governance obligations are increasingly board-level. Regulators expect boards to:

- Receive regular reporting on privacy and cyber risk as a standing agenda item, not just in response to incidents
- Approve and own the organisation's privacy policy and governance framework, not merely note it
- Ensure accountability structures for privacy are clearly defined, with named executive owners
- Require management to demonstrate readiness for the December 2026 ADM (automated decision-making) disclosure obligations, with evidence, not just assertions

There is a particular obligation on boards to understand the automated decision-making landscape within their organisation. Boards that have approved AI deployments without ensuring the associated privacy disclosures are accurate and current will find themselves poorly positioned if a regulator or class action draws scrutiny to those systems.

RISK ALERT

The statutory tort for serious privacy invasion, in force since June 2025, means that a single high-profile AI-related privacy failure can now generate both regulatory proceedings and mass litigation simultaneously. D&O (directors' and officers') coverage should be reviewed in light of this.

BOARD AGENDA ITEM

Ask management: Have we identified all systems that use personal information to make or support significant decisions about individuals? Does our current Privacy Policy accurately describe those systems? What work is underway to meet the 10 December 2026 disclosure deadline, and is it on track?

Part 3: Directors' Duties and AI – The Personal Liability Boards Can't Ignore

The Existing Framework Has Not Changed – But Its Application Has

Australian directors' duties are set out in the Corporations Act 2001 (Cth), and they have not been amended for AI. What has changed is how courts and regulators are applying those duties to AI-related governance failures, and the answer, increasingly, is that existing duties already require competent AI oversight.

The four duties most directly engaged by AI governance are:

- **Section 180, Duty of Care and Diligence:** directors must act with the degree of care and diligence that a reasonable person would exercise. Courts are increasingly treating failure to govern AI risks as a failure of reasonable care.
- **Section 181, Duty of Good Faith:** directors must act in good faith in the best interests of the corporation. Approving AI deployments without understanding their risk profile may breach this duty if the deployment later causes harm.
- **Section 182, Improper Use of Position:** directors must not use their position to gain advantage for themselves or others at the company's expense.
- **Section 183, Improper Use of Information:** AI systems that process inside information or confidential data create additional exposure under this provision.

The Safe Harbours Are Under Strain

The Business Judgment Rule (s 180(2)) protects directors who make an informed, good-faith business judgment from liability for outcomes. This protection has always required that the director genuinely informed themselves before deciding. The Chief Justice of the Federal Court of Australia recently articulated why AI changes this calculus in ways boards have not yet absorbed.

> **"It must be extremely doubtful that a director who simply adopts" an AI recommendation could invoke the business judgment rule. Courts require rational reasoning and conscious decision-making: neither is present when a director blindly adopts machine output.**
>
> CHIEF JUSTICE, FEDERAL COURT OF AUSTRALIA

The Section 189 reliance defence, which allows directors to rely on information provided by others, similarly does not straightforwardly extend to AI. The statute assumes information comes from a person. AI systems do not qualify. More fundamentally, the defence requires independent assessment of the advice, which the "black box" nature of many AI systems makes difficult to demonstrate.

The practical implication: a director who approved a strategy, a loan portfolio decision, or a risk assessment on the basis of AI-generated outputs, without being able to articulate what questions they asked, what limitations they identified, and why they were satisfied, is in a legally precarious position.

The Star Entertainment Judgment and Shadow AI

The March 2026 Star Entertainment judgment (Justice Lee, Federal Court) introduced a specific new risk: shadow AI. The judgment noted that directors had been using personal accounts on consumer AI tools, including to prepare for board meetings, creating discoverable prompts and transcripts that became evidence in proceedings.

This is not a theoretical concern. Personal AI tool usage by executives and directors generates logs, query histories, and output records that are accessible to litigation counterparties under discovery obligations. If a director used a consumer chatbot to draft board papers, summarise financial data, or develop advice on a strategic matter, those interactions may be discoverable.

RISK ALERT

Using personal accounts on consumer AI tools for board-related work is a governance failure with legal consequences. Organisations should provide board members with sanctioned, secure AI platforms, and have a policy that makes personal-tool usage on confidential matters clearly prohibited.

The RI Advice Principle: Inaction Also Creates Liability

The ASIC v RI Advice Group case established a principle that operates symmetrically: just as directors can be liable for deploying AI without adequate governance, they can also be liable for failing to use AI where a reasonable organisation would have done so. As AI expands the range of analysis that is practicably available, the standard of what constitutes being "appropriately informed" before a material decision rises accordingly.

A board that resists AI adoption entirely, or treats it as someone else's concern, may find itself unable to satisfy courts that it exercised the care and diligence expected of a modern Australian director.

Misleading AI Capability Claims

Directors also face exposure under the Australian Consumer Law (ACL) where organisations make misleading statements about their AI capabilities, to customers, investors, or regulators. MinterEllison's analysis of recent enforcement trends identifies AI capability claims as a growing source of ACL liability, with directors potentially personally liable where they approved or were aware of the claims.

Investor relations materials, product disclosures, and regulatory submissions that overstate the sophistication, reliability, or governance of AI systems are high-risk areas. Boards should require management to apply the same legal review to AI-related claims as they would to financial projections.

What Good Governance Looks Like

The emerging judicial and regulatory standard for AI governance at board level has five components:

- **AI Literacy:** each director should have sufficient understanding of the AI systems material to the organisation's operations to ask meaningful questions. This does not require technical expertise, but it does require education.
- **AI Inventory and Risk Classification:** management should maintain and regularly report to the board an inventory of AI systems in use, classified by risk level. High-risk systems warrant board-level oversight; lower-risk systems can be managed at executive level with appropriate escalation triggers.

- **Independent Verification of AI Outputs:** for material decisions informed by AI, the board should require evidence that management verified outputs through non-AI benchmarks: historical data, industry norms, external expertise.
- **Documented Decision Processes:** given that prompts and AI interactions are discoverable, organisations should document how AI inputs are used in decision-making. The paper trail should show that humans exercised genuine independent judgment.
- **D&O Insurance Review:** coverage should be assessed for adequacy in the context of AI-related liability, including privacy breaches arising from AI systems and ACL claims related to AI capability statements.

BOARD AGENDA ITEM

Ask management: Do we have an AI inventory classified by risk level? What AI tools are board members and executives currently using, on what platforms, and under what policies? Have we reviewed our D&O coverage for AI-related liability? Are there any material AI capability claims in our investor, customer, or regulatory communications that should be reviewed for accuracy?

Part 4: Mandatory AI Standards – The 2027 Horizon

The Regulatory Pivot

On 15 July 2026, Prime Minister Albanese announced a decisive shift in Australia's approach to AI governance. After five years of voluntary frameworks, guidelines, and principles-based guidance, the government committed to mandatory standards, ending what one commentator called "the era of polite requests and hoping for the best."

The announcement had three substantive elements:

- **Mandatory AI Standards** covering the development and deployment of AI systems by Australian organisations, with a staged commencement expected in 2027-2028
- **An Office of AI** established within the Department of the Prime Minister and Cabinet, giving it cross-portfolio authority over Treasury, Industry, Home Affairs, and the Digital Transformation Agency
- **Creator Copyright Protections** requiring consent and price control for the use of Australian creative works in AI training, affecting organisations fine-tuning or training models on local content

What the Mandatory Framework Is Expected to Cover

Draft standards and implementation guidance are expected in late 2026, with legislation introduced to Parliament in early 2027. Based on signals from the consultation process, regulatory horizon trackers, and international precedent, the framework is expected to include:

- A risk-tiered approach, distinguishing high-risk AI applications (those affecting individual rights, safety, or access to services) from lower-risk uses
- Mandatory conformance with AS ISO 42001 or an equivalent standard for organisations developing or deploying high-risk AI
- Transparency and explainability requirements for consequential AI decisions, building on the Privacy Act's ADM obligations
- Mandatory incident reporting for AI system failures that cause harm or material operational disruption
- Obligations on AI providers and deployers, similar to the EU (European Union) AI Act's supply-chain accountability model, requiring clarity on who bears responsibility for AI-related harms

The Office of AI

The placement of the Office of AI within the Department of Prime Minister and Cabinet is significant. It is not a line regulator like the OAIC or ASIC. Instead, it has a convening and coordination mandate across the Commonwealth, meaning it can direct other agencies on AI-related matters and set whole-of-government expectations that flow through procurement, funding, and regulatory interactions.

For large organisations that interact with government, including those in health, infrastructure, financial services, education, and defence supply chains, the Office of AI will be a consequential relationship. Its early guidance and consultation documents will signal where enforcement attention will focus.

Why the "Wait and See" Strategy Is Over

The instinct to wait for final legislation before acting is understandable but carries significant cost in this instance. Several factors make early preparation the rational choice:

- The mandatory framework will require documentation of AI governance history. Organisations that start now will have evidence of their governance journey; those that start after legislation passes will not.
- AS ISO 42001 implementation takes time, typically 9-18 months to reach a state that would withstand external audit. Beginning after legislation passes means operating out of compliance during the transition.
- The Privacy Act's ADM obligations take effect 10 December 2026, regardless of the 2027 legislation timeline. There is no grace period for the privacy disclosure obligations.
- The OAIC, ASIC, and APRA are already enforcing existing obligations in ways that treat AI governance as a core organisational competence. Regulatory relationships formed now, before mandatory requirements arrive, matter.

> **INTERNATIONAL CONTEXT**
>
> The international precedent is instructive. When the EU AI Act entered force in August 2024, EU organisations with existing ISO 42001 programs adapted at a fraction of the cost of those starting from scratch. Australia's mandatory transition is expected to follow the same pattern.

A Preparation Roadmap for 2026-2027

Based on the available signals, a practical preparation sequence for Australian boards and C-suite leaders is:

- **Now to Q4 2026:** Commission an AI inventory and risk classification exercise. Identify all AI systems in operation, classify by risk level, and map against AS ISO 42001 requirements to identify gaps.
- **By 10 December 2026:** Update privacy policies to meet ADM disclosure obligations under APP 1.7 and 1.8. This is a hard deadline with active enforcement.
- **Q1 2027:** Engage with the Office of AI's consultation process on the mandatory standards framework. Early engagement shapes guidance; late engagement means compliance-taking rather than standard-setting.
- **2027:** Implement an AS ISO 42001-aligned AI Management System, including documented risk assessment processes, human oversight mechanisms, and audit capability.
- **Ongoing:** Build AI governance into board skill matrices, director induction, and executive KPIs (key performance indicators).

> **BOARD AGENDA ITEM**
>
> Ask management: Have we mapped our AI systems against the anticipated mandatory framework requirements? Do we have a budget and timeline for AS ISO 42001 implementation? Are we tracking the Office of AI's consultation process, and do we have representation in those consultations? What is our plan for the 10 December 2026 privacy deadline?

Conclusion: This Is a Governance Imperative, Not a Technology Project

The four themes in this paper (AS ISO 42001, the Privacy Act's ADM obligations, directors' duties, and the incoming mandatory framework) share a common thread. In each case, the regulatory and judicial system is doing something specific: it is moving accountability for AI upward.

AI is no longer something boards can treat as an operational matter delegated to the technology function. The obligations described in this paper sit with the board, in the same way that financial reporting, WHS (work health and safety), and environmental compliance sit with the board. They are not a one-time project; they are an ongoing governance function.

The good news is that the standard for what good looks like is reasonably clear. AS ISO 42001 provides a structured framework. The Privacy Act's disclosure requirements, while operationally complex, have a defined scope. The judicial guidance on directors' duties, while confronting, is consistent: exercise genuine oversight, verify rather than accept, and document the process.

The organisations that will navigate this environment well are those where the board has decided that AI governance is a board-level priority, not because regulators have mandated it yet, but because the risk of not treating it that way has become too large to ignore.

The Australian regulatory landscape for AI has reached an inflection point. The question for boards is no longer whether to govern AI seriously, but whether to begin before or after the consequences of not doing so materialise.

Appendix: Quick Reference – Key Obligations and Timelines

REQUIREMENT	DEADLINE / STATUS	PRIMARY REGULATOR	BOARD ACTION REQUIRED
AS ISO 42001 (AI Management System)	In force (Feb 2024), voluntary but increasingly expected	Standards Australia / Cross-regulator	Gap assessment; implementation roadmap
Privacy Act ADM Disclosure (APP 1.7–1.8)	Mandatory from 10 December 2026	OAIC	Privacy policy audit; system mapping
Statutory Tort: Privacy Invasion	In force since 10 June 2025	Courts	D&O coverage review; incident response plan
OAIC Compliance Sweeps	Ongoing (commenced Jan 2026)	OAIC	Proactive compliance readiness
Office of AI Established	July 2026	PM&C (Department of the Prime Minister and Cabinet)	Consultation engagement
Mandatory AI Standards Legislation	Expected early 2027	Office of AI / Parliament	Draft framework monitoring; AIMS implementation
Directors' Duties: AI Application	Immediate (existing law)	ASIC / Courts	AI literacy; governance framework; D&O review

Key Legislation and Standards Referenced

- Privacy Act 1988 (Cth), as amended by the Privacy and Other Legislation Amendment Act 2024 (Cth)
- Australian Privacy Principles 1.7 and 1.8: Automated Decision-Making Transparency
- Corporations Act 2001 (Cth): ss 180, 181, 182, 183, 189, 190
- AS ISO 42001:2023: Artificial Intelligence Management Systems (Standards Australia, February 2024)
- Australian Consumer Law (Schedule 2, Competition and Consumer Act 2010 (Cth))
- ASIC v RI Advice Group Pty Ltd [2022] FCA (Federal Court of Australia) 496: cyber risk and directors' duties

Disclaimer

This white paper is prepared for general information and educational purposes. It does not constitute legal advice and should not be relied upon as such. Organisations should obtain independent legal advice in relation to their specific circumstances and obligations. The regulatory landscape described is current as at August 2026 and is subject to change.