# Artificial Intelligence and the Australian Boardroom

## Director Duties in the Age of Automated Execution

*An independent briefing for Australian company directors on legislative obligations, regulatory expectations, and governance liabilities regarding Artificial Intelligence.*

White Paper | August 2026

Cetus AI Pty Ltd | ABN 16 695 570 819 | cetusai.com.au

## Executive Summary

Artificial intelligence (AI) has rapidly shifted from a technology novelty to a core driver of enterprise value. For Australian boards, this transition brings a profound shift in risk and liability. As AI systems move from internal productivity tools to autonomous agents executing customer-facing decisions, the regulatory landscape has tightened significantly.

The year 2026 has seen a watershed of regulatory activity. The Australian Prudential Regulation Authority (APRA) issued a formal warning to industry regarding AI risk management [1]. The Federal Parliament passed the Competition and Consumer Amendment (Unfair Trading Practices) Act 2026, capturing algorithmic manipulation [2]. Crucially, sweeping reforms to the Privacy Act 1988 have introduced mandatory transparency obligations for automated decision-making that take effect in December 2026 [3].

For company directors, the message from regulators and the courts is unambiguous: the statutory duty of care and diligence under section 180 of the Corporations Act 2001 cannot be outsourced to an algorithm. Boards are expected to maintain sufficient AI literacy to challenge executive management, oversee AI strategy, and ensure that governance frameworks move beyond aspirational policies to enforceable technical controls.

This briefing outlines the current legislative and regulatory landscape affecting Australian boards, detailing the specific obligations directors must meet to discharge their duties in the age of artificial intelligence.

## The Corporations Act and Director Duties

The foundation of board liability regarding AI rests upon the Corporations Act 2001 (Cth). While the Act is technology-neutral, its application to AI governance is becoming increasingly specific as regulators focus on digital oversight.

### The Duty of Care and Diligence (Section 180)

Under section 180(1) of the Corporations Act, directors must exercise their powers and discharge their duties with the degree of care and diligence that a reasonable person would exercise. In the context of AI, this duty requires active, informed oversight rather than passive reliance on management or vendor assurances.

Recent jurisprudence, including the Federal Court's decision in ASIC v Bekier (2026), provides critical guidance on how section 180 applies to risk management oversight [4]. The Australian Securities and Investments Commission (ASIC) has explicitly identified both the use of AI and directors' conduct as primary focus areas for 2025–26 [5].

To discharge their duty of care regarding AI, directors must ensure they possess adequate technical literacy. Ignorance of how AI models function, or the risks they introduce, is not a legally defensible position. Boards must actively interrogate management on how AI is deployed, what data it accesses, and what safeguards prevent algorithmic failure.

### The Fettering of Discretion

A fundamental principle of corporate law is that directors must not fetter their discretion. They cannot delegate their decision-making authority to a third party—or, increasingly, to an autonomous system—without retaining ultimate oversight.

If an enterprise deploys an AI system that makes autonomous decisions affecting capital allocation, customer outcomes, or regulatory compliance, the board remains legally responsible for those outcomes. Directors cannot claim the legal protections of the business judgment rule if they have blindly deferred to an algorithm's output without understanding its parameters or ensuring adequate human-in-the-loop controls for high-risk decisions.

## The Shifting Regulatory Landscape (2026 Updates)

The Australian regulatory environment for AI matured rapidly throughout 2025 and 2026, shifting from voluntary guidelines to enforceable statutory obligations. Boards must ensure their compliance frameworks are updated to reflect these new realities.

## Privacy Act Reforms and Automated Decision-Making

The most immediate compliance challenge for Australian boards stems from reforms to the Privacy Act 1988. The Privacy and Other Legislation Amendment Act 2024 introduced sweeping changes that take effect on 10 December 2026 [3].

From this date, Australian Privacy Principles (APP) entities face new obligations regarding automated decision-making (ADM). Organisations must disclose in their privacy policies:

1. The types of personal information used in substantially automated decisions.
2. The nature of decisions made solely or significantly by computer programs.
3. Instances where those decisions could reasonably be expected to significantly affect individual rights or interests.

The Office of the Australian Information Commissioner (OAIC) is currently finalising formal guidance on these transparency requirements [6]. Boards must ensure management is conducting comprehensive audits of all AI systems to identify which processes trigger these new disclosure obligations before the December deadline.

*Furthermore, penalties for serious or repeated privacy breaches remain severe, reaching up to the greater of $50 million, three times the benefit obtained, or 30 per cent of adjusted turnover.*

## APRA's Formal Warning on AI Governance

For the financial services sector, the regulatory expectations are explicitly defined. On 30 April 2026, APRA issued a formal letter to all regulated entities following a targeted deep-dive into AI adoption among large banks, insurers, and superannuation trustees [1].

The findings present a stark warning for boards. APRA observed that while AI adoption is accelerating rapidly, governance, risk management, and assurance practices are failing to keep pace. Specifically, APRA noted an overreliance by boards on vendor presentations without sufficient examination of key AI risks, such as unpredictable model behaviour.

APRA has set clear expectations that boards must:

- Maintain sufficient understanding and literacy to set strategic direction and provide effective challenge.
- Oversee an AI strategy consistent with the entity's risk appetite.
- Ensure effective monitoring and reporting, particularly for third-party dependencies.

*Crucially, APRA highlighted the risk of "Shadow AI"—the use of enterprise AI tools by staff outside approved control frameworks. The regulator criticised entities that rely primarily on policy direction rather than enforceable technical restrictions. APRA warned that it will take stronger supervisory action, and pursue enforcement where appropriate, against entities that fail to adequately manage AI risks under Prudential Standards CPS 230 (Operational Risk Management) and CPS 234 (Information Security).*

### Australian Consumer Law and Algorithmic Manipulation

The Competition and Consumer Amendment (Unfair Trading Practices) Act 2026, which received Royal Assent in July 2026 and commences on 1 July 2027, introduces a general prohibition on unfair trading conduct [2].

This technology-neutral legislation prohibits conduct that manipulates a consumer or unreasonably distorts their decision-making environment. The Australian Competition and Consumer Commission (ACCC) has flagged that this will capture AI-enabled “dark patterns,” algorithmic manipulation, and deceptive chatbot interactions. Boards must ensure that AI systems deployed in customer-facing roles do not engage in conduct that breaches these new unfair trading provisions, which carry maximum penalties of $100 million.

### Sovereign AI Standards and Data Centres

On 15 July 2026, the Prime Minister announced plans to legislate an Australian Standards for AI framework, expected in early 2027 [7]. Coordinated by the newly established Office of AI within the Department of the Prime Minister and Cabinet, this framework will impose mandatory requirements on large AI data centres regarding energy, water, and grid support. While primarily targeting infrastructure providers, boards of enterprises procuring AI services must monitor these developments, as they will impact the cost and compliance requirements of the sovereign AI supply chain.

## The “Shadow AI” Governance Gap

Perhaps the most significant unmanaged risk facing Australian boards is the proliferation of “Shadow AI.” This occurs when employees utilise unsanctioned, often consumer-grade generative AI tools to process corporate data, write code, or draft documents.

Industry data highlights the scale of this vulnerability. Recent reports indicate that 68 per cent of employees use free-tier AI tools via personal accounts, and massive volumes of corporate data are being transferred to these applications daily [8].

When an employee uploads a confidential board paper, customer financial records, or proprietary source code to a public AI model, the organisation loses control of that data. It may be used to train future iterations of the model, resulting in a data breach that triggers mandatory notification under the Notifiable Data Breaches (NDB) scheme.

As APRA noted in its April 2026 letter, traditional security controls are proving inadequate. Endpoint Data Loss Prevention (DLP) tools, Cloud Access Security Brokers (CASB), and Mobile Device Management (MDM) solutions are entirely bypassed the moment an employee accesses an AI tool via their personal smartphone.

**The Board’s Responsibility: Boards must ask management how they are enforcing AI policy. If the answer relies on employee training, acceptable use policies, or trust, the organisation is exposed.**

**Regulators increasingly expect technical enforcement mechanisms that actively prevent data leakage to unsanctioned AI models, rather than merely documenting that a policy exists.**

## Framework for Board-Level AI Oversight

To discharge their duties effectively, Australian directors should implement a structured framework for AI oversight. The following table outlines the critical domains boards must interrogate.

| Governance Domain | Board-Level Inquiry | Regulatory Driver |
|---|---|---|
| **Visibility & Inventory** | Do we have a comprehensive, real-time register of all AI systems (sanctioned and shadow) operating within the enterprise? | APRA CPS 234; ASIC expectations |
| **Automated Decisions** | Which AI systems are making decisions that significantly affect individuals, and are we prepared for the December 2026 transparency disclosures? | Privacy Act (APP 1.7-1.9) |
| **Technical Enforcement** | How are our AI acceptable use policies technically enforced? Can we cryptographically prove that PII is not being sent to public models? | APRA Letter (April 2026); Privacy Act |
| **Third-Party Risk** | What is our exposure to upstream AI vendors? Do we understand the data handling practices of the foundation models we rely upon? | APRA CPS 230 |
| **Bias & Discrimination** | How do we test our AI models for discriminatory outcomes in recruitment, lending, or service delivery? | Anti-Discrimination Laws |
| **Consumer Protection** | Are our customer-facing AI agents transparent about their non-human nature, and do they avoid manipulative practices? | Australian Consumer Law |
| **Board Literacy** | Does the board possess the requisite technical understanding to challenge management's AI strategy and risk assessments? | Corporations Act (s.180) |

## Conclusion

The era of treating artificial intelligence as a purely operational IT issue has ended. AI is now a central pillar of corporate governance, carrying significant legal, regulatory, and reputational liabilities.

Australian regulators have made their expectations clear: boards must actively govern AI risk, ensure transparent and fair automated decision-making, and implement technical controls that enforce compliance. Directors who fail to elevate their AI literacy, or who allow their organisations to rely on aspirational policies without architectural enforcement, risk breaching their statutory duties.

*The challenge for the modern board is not to stifle AI innovation, but to ensure it is deployed within a secure, sovereign, and strictly governed control framework.*

## References


[1] Australian Prudential Regulation Authority (APRA). "APRA Letter to Industry on Artificial Intelligence (AI)." 30 April 2026. https://www.apra.gov.au/news-and-publications/apra-letter-industry-artificial-intelligence-ai

[2] Parliament of Australia. Competition and Consumer Amendment (Unfair Trading Practices) Act 2026. Assented 6 July 2026.

[3] Parliament of Australia. Privacy and Other Legislation Amendment Act 2024.

[4] Federal Court of Australia. ASIC v Bekier. March 2026.

[5] Australian Securities and Investments Commission (ASIC). "Corporate Plan 2025–26."

[6] Office of the Australian Information Commissioner (OAIC). "Consultation on Guidance for Transparency in Automated Decision Making." May 2026.

[7] Prime Minister of Australia. "AI and Australia’s Interests." Media Release, 15 July 2026.

[8] Menlo Security. "2025 AI Workspace Report."