Shankar Siva | Governance and risk adviser | Author, Warranted AI

Submission to the Australian Prudential Regulation Authority consultation on Draft Prudential Standard CPS 510 Governance (June 2026)

26 August 2026

To: General Manager, Policy and Frameworks, Policy and Advice Division
Australian Prudential Regulation Authority

Submitted via: PolicyDevelopment@apra.gov.au

About this submission

Shankar Siva is a governance and risk adviser and the author of *Warranted AI: Seven Conditions for Defensible Reliance on Consequential AI-Enabled Decisions and Actions in Regulated Industries*.

Warranted AI defines a Warranting Standard for organisational reliance on consequential AI-enabled decisions and actions. It is a threshold test for when reliance can be justified, defended and stood behind. It is not an AI ethics framework, a model explainability method or a compliance checklist. Warranted AI is not proof of legality. It is the evidentiary condition for defensible organisational reliance. The examples and implementation observations in this submission draw on direct experience working with APRA-regulated entities navigating governance challenges in operational settings.

This submission addresses Question 6 only (artificial intelligence governance). Questions 1 to 5 are not addressed.

This submission builds on, and does not repeat, two prior contributions: a June 2026 submission to the Office of the Australian Information Commissioner on automated decision-making transparency, which addressed the transparency obligations an entity owes to the individuals affected by its decisions; and a July 2026 submission to the Financial Stability Board's consultation on sound practices for responsible AI adoption, which addressed oversight practices at the level of the financial institution. This submission addresses a third and distinct level: the board governance mechanics that CPS 510 governs. Its contribution is a single question - what must be demonstrably true for a board to warrant reliance on AI-mediated information and AI-shaped decisions - applied to the draft standard's own provisions.

Question 6 - Do the draft CPS 510 requirements adequately support the oversight and use of AI and related technologies? Are there areas where additional expectations or guidance would assist?

Technology neutrality is the right design. It holds only if the mechanics the standard relies on - information, delegation, challenge and skills - remain effective once what they act on becomes AI-mediated.

Draft CPS 510 contains no AI-specific requirements. The first point of this submission is that this is a strength to be preserved, not a gap to be filled. A standard that named particular technologies would date quickly and invite compliance-by-category. The draft instead governs through technology-neutral mechanics: the quality of information provided to the Board (paragraphs 20 to 22), the guardrails around delegation (paragraphs 15 to 17), the Board's non-delegable duty of constructive challenge (paragraph 13(h)), and the skills the Board must collectively hold (paragraphs 13(g) and 47 to 49). Question 6 therefore does not ask whether the standard should mention AI. It asks whether those mechanics remain effective as the substance they act on changes character. In three places they require reinforcement - largely in CPG 510, with one candidate refinement to the standard itself.

1. Management information: Quality has a provenance dimension the draft standard does not yet reach

Paragraph 20 requires a policy that sets the Board's expectations for management information, and covers its expected content, quality and frequency. Paragraph 22 requires senior managers to brief the Board with succinct and relevant information. Neither paragraph reaches provenance. Neither paragraph asks whether the information put to the Board was generated, filtered, summarised or prioritised by an AI system, or what human validation stood between those systems and the Board pack. As regulated entities adopt AI in their reporting chains - and they are doing so now - a growing share of what boards read will be machine-produced before it is management-endorsed.

The prudential significance is that provenance-blind information quietly undermines the Board's non-delegable responsibilities. Constructive challenge under paragraph 13(h) presupposes that the Board can interrogate what it is told. A board cannot meaningfully challenge information whose generation it cannot see. This is the concern the Warranted AI framework addresses through Oversight Independence: oversight requires the capacity to observe, challenge, interrupt or override through means that do not depend on the overseen system's own self-assessment.[1] Transposed to Board level - where the triggers for the Board's attention, the information available to it, and its practical capacity to intervene all run through the same AI-mediated reporting chain, the Board holds oversight in form but not in effect. That is the gap between governance intent - what the entity approved and documented - and Governance in Effect - what demonstrably operated as decisions were shaped.

---

[1] Shankar Siva, *Warranted AI: Seven Conditions for Defensible Reliance on Consequential AI-Enabled Decisions and Actions in Regulated Industries* (Working Paper, July 2026), SSRN Abstract ID 7137118, available at https://ssrn.com/abstract=7137118. The Oversight Independence condition is developed in Chapter 4.

Proposed refinement. The management information policy required by paragraph 20 should address provenance: the identification of AI-generated or AI-mediated content that is material to Board decision-making, and the validation applied to it before it reaches the Board. This could be effected by a short addition to paragraph 20 or developed in CPG 510. Paragraph 54(e) already requires the independent performance assessment of a significant financial institution to evaluate the quality of management information. Guidance should confirm that quality includes provenance.

2. Delegation: Paragraph 17's guardrails track the human delegate, not the operative mechanism

Paragraphs 15 to 17 establish a clear and welcome delegation architecture. Paragraph 17(c) requires mechanisms for monitoring the proper and effective exercise of delegated authority. As entities deploy AI - including increasingly agentic systems - in credit, claims, remediation and customer-facing workflows, a delegate will often exercise that authority substantially in reliance on system outputs. The delegation register will then record a compliant chain of human authority while the operative decision mechanism sits partly outside it. Monitoring that tracks the human signature but not the mechanism the human adopted satisfies the form of paragraph 17(c) without satisfying its purpose.

Proposed refinement. CPG 510 should direct entities to identify where delegated authority is materially exercised in reliance on AI systems, and to apply the paragraph 17 guardrails - documentation, risk-based limits, monitoring and review - at that layer. Board oversight of delegation then reaches the mechanism that actually produces outcomes, not only the officer who holds the authority.

3. Skills: The capability boards need is prudential, not technical

Paragraphs 13(g) and 47 to 49 are the right architecture, and APRA is right not to prescribe specific skills. The practical risk is miscalibration in either direction: boards treating AI as a specialist matter that belongs wholly to management or reaching for deep technical specialisation at Board level. Neither serves the standard. The capability the Board needs is a governance capability - the collective ability to interrogate AI-mediated information, to recognise the limits of a system's self-reporting, and to identify when reliance on an AI-shaped decision is not warranted. A confidence score is a system's testimony about itself, not an independent check. That capability belongs in the Board skills matrix (paragraph 47), with assessment criteria under paragraph 48, and it bears on fitness and propriety where a responsible person's role requires them to stand behind reliance on such systems.

Proposed refinement. CPG 510 should characterise AI-related Board capability as the capability to oversee and challenge - not to build or operate - and should frame the skills-matrix assessment criteria in those terms.

One test to anchor the guidance

If APRA develops AI-related expectations in CPG 510, this submission suggests anchoring them in a single test rather than a list of practices: what must be demonstrably true for the Board to warrant reliance on AI-mediated information and AI-shaped decisions? The standard implied is a threshold,

not a ceiling - the minimum that must hold before reliance is warranted, not an ideal of governance perfection. Framed this way, the test gives proportionality, the concern of Question 1, a floor to be proportionate to. And it remains stable as technologies change: the test does not depend on whether the mechanism is a scoring model, a large language model or an agentic system. It is also coherent with the direction of current international work. The Financial Stability Board's consultation on responsible AI adoption engages the same question at institution level, and the author's submission there proposed an equivalent independence test for operational oversight.[2]

Reinforced at these three points, the technology neutrality of draft CPS 510 would position the standard to govern AI well without ever naming it - which is the mark of a governance standard built to last.

The author would be glad to assist APRA's consideration of these points.

Shankar Siva
Governance and risk adviser
Author, *Warranted AI: Seven Conditions for Defensible Reliance on Consequential AI-Enabled Decisions and Actions in Regulated Industries*
ORCID 0009-0009-6920-4697

[2] Financial Stability Board, *Sound Practices for Responsible Adoption of Artificial Intelligence (AI): Consultation Report*, 10 June 2026 (consultation closed 22 July 2026); the author's response is published at https://www.fsb.org/uploads/Shankar-S-independent-response.pdf.