A 42 CONSULTING.AI WHITEPAPER

# The Regulatory Readiness Gap

*The Regulators Have Set Their Expectations. The Question Is No Longer Whether You Comply — but Whether You Can Prove It Under Examination.*

**Dr M Maruf Hossain, PhD, GAICD**
Chief AI Strategist · 42 Consulting.AI
enquire@42consulting.ai · www.42consulting.ai

42 Consulting.AI
Scaling Intelligence

CONTENTS

# Table of Contents

FOREWORD

# From the Author.

There is a fair objection to any call for stronger AI governance: *we have a policy for that.* Most regulated organisations do.

This paper is my answer to that objection. Having a policy and being able to prove, under examination, that the policy operates are two entirely different things — and the distance between them is precisely where regulatory failure now occurs.

The regulators have stopped asking whether you have governance on paper. They are asking whether you can *demonstrate* it: produce the record, show the monitoring, and provide evidence of the board's challenge. A control you cannot evidence is, for supervisory purposes, a control you do not have.

This whitepaper sets out what readiness actually requires — obligation by obligation, against the standards that already apply. It is written for the directors and accountable executives who will, personally, answer for the gap between what their organisation says and what it can show.

*M Maruf Hossain*

SECTION ONE

# Executive Summary.

The regulatory posture toward enterprise AI has shifted from setting expectations to testing them. This whitepaper establishes three propositions for boards and accountable executives in regulated industries.

**First, the expectations are now explicit and testable — not aspirational.** APRA's April 2026 letter, ASIC's enforcement posture, the Financial Accountability Regime, and the EU AI Act's high-risk obligations do not, in general terms, ask organisations to take AI seriously. They specify what must exist, what must be monitored, and what must be producible. Each expectation implies a test, and the tests are no longer hypothetical.

**Second, readiness means demonstrable operation under examination — not compliance on paper.** A policy describes intent. Readiness is the capacity to produce evidence that the intended controls actually operate: the inventory that shows what is deployed, the record that shows what was authorised, the monitoring that shows the system still behaves as approved. Under supervisory examination, an undocumented control and an absent control are indistinguishable.

**Third, the gap between stated policy and evidence-based readiness is where examination failure occurs.** Organisations rarely fail an examination because they lack a policy. They fail because they could not evidence that the policy operated — because the monitoring was point-in-time, the authorisation was reconstructed after the fact, or the board could not show it had provided effective challenge. This paper maps each regulatory obligation to the specific evidence that demonstrates readiness and sets out five questions that test whether an organisation can produce it.

The argument is built on the Australian regulatory stack as its spine — APRA, ASIC, and the Financial Accountability Regime — with the EU AI Act as a second jurisdiction and ISO/IEC 42001 as the management-system baseline beneath both.

SECTION TWO

# From Expectation *to Examination.*

For most of the period in which enterprises adopted AI, regulatory expectations were expressed in principle. Entities were expected to manage AI risk, govern it appropriately, and use it responsibly. Those expressions left wide latitude for how and whether the principle was operationalised.

APRA's letter of 30 April 2026 changed the register. It did not announce new principles; it reported what a targeted supervisory review had *found* — and framed its observations as areas where governance, risk management, and assurance are failing to keep pace with AI adoption. An observation drawn from examination is not advice. It is a finding against which future examinations will be measured.

The distinction matters because of what it implies about evidence. When a regulator observes that governance maturity is lagging, that boards rely too heavily on vendor presentations, and that point-in-time assurance is ill-suited to systems that learn and degrade, it is describing things it sought but did not adequately find. The next supervisory engagement will look for the same things — and the entity will need to *show* them, not assert them.

This is the shift from expectation to examination. Readiness, in this environment, is not the existence of a governance framework. It is the demonstrable operation of that framework, evidenced to a supervisor who is no longer satisfied by a description. The remainder of this paper treats every regulatory obligation as a question of evidence: what would the organisation produce if asked to prove the control operates?

> *A control you cannot evidence is, for supervisory purposes, a control you do not have.*

SECTION THREE

# The Australian *Obligation Stack.*

AI does not arrive in a regulatory vacuum in Australia. A stack of prudential standards already applies to it — and the April 2026 AI expectations sit on top of that stack rather than replacing it. Understanding which instrument governs which obligation is the first step toward readiness, because each carries its own evidentiary demand.

At the foundation is **CPS 220 Risk Management**, which requires a framework that spans all material risks. Built upon it, **CPS 230 Operational Risk Management** — effective 1 July 2025 — requires entities to govern operational risk, maintain business continuity, manage service providers, and, critically, to assess the impact of new technologies on their operational risk profile. AI is precisely such a technology. **CPS 234 Information Security** makes the board ultimately responsible for information security and requires the classification of information assets and the assessment of third-party security capabilities. **CPS 510 Governance** sets expectations for governance and remuneration policy, and **CPS 511 Remuneration** aligns with the accountability regime addressed in the next section.

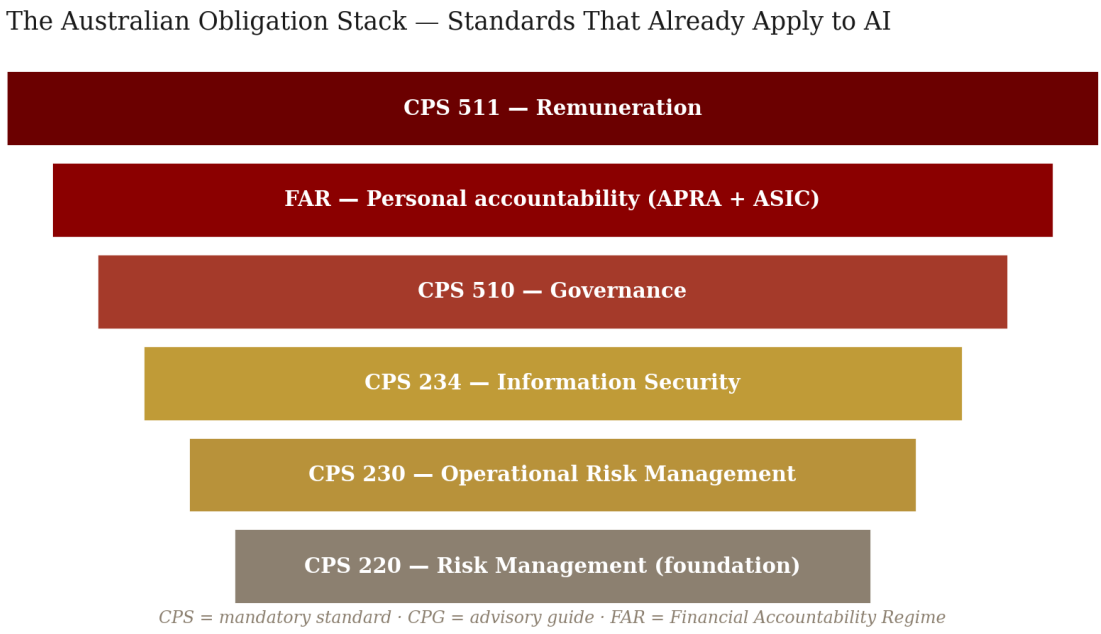


*Figure 1 — The Australian obligation stack. The April 2026 AI expectations layer onto standards that already apply.*

A precise distinction must be held throughout. Instruments carrying the *CPS* prefix are Prudential *Standards* — mandatory and enforceable. Instruments carrying the *CPG* prefix, such as **CPG 235 Managing Data Risk**, are Prudential Practice *Guides* — advisory, influential, but not in themselves

enforceable obligations. Conflating the two overstates some obligations and understates others; readiness depends on knowing which is which.

### WHICH INSTRUMENT GOVERNS WHICH AI OBLIGATION

**CPS 220 — Risk Management**
The obligation to bring AI within the entity's risk management framework as a material risk, not a side technology.

**CPS 230 — Operational Risk**
The obligation to assess AI's impact on the operational risk profile, manage AI service providers, and maintain resilience — effective 1 July 2025.

**CPS 234 — Information Security**
Board-level responsibility for information security, asset classification, and assessment of third-party (including AI vendor) security capability.

**CPS 510 / CPS 511 — Governance & Remuneration**
Governance expectations and the remuneration-policy mechanics that the Financial Accountability Regime relies upon.

**CPG 235 — Managing Data Risk**
Advisory guidance on data-lifecycle risk — directly relevant to AI data governance, but a guide, not a binding standard.

SECTION FOUR

# The Personal Accountability *Mechanism.*

The feature of the Australian regime that generalist AI governance frameworks consistently miss is this: a governance failure involving AI does not stop at the entity. The Financial Accountability Regime attaches to named individuals.

The Financial Accountability Regime commenced for banks on 15 March 2024 and was extended to insurers and superannuation trustees on 15 March 2025. In APRA and ASIC's own words, it imposes a strengthened framework of responsibility and accountability to improve the risk governance cultures of regulated entities, *their directors, and their most senior executives*. It is jointly administered: APRA sets the prudential expectation, and ASIC brings the conduct lens — and both regulators hold enforcement powers, exercised independently except where disqualifying an individual or issuing directions, which require agreement.

The regime operates through named accountability. Entities meeting the enhanced notification threshold must prepare accountability statements — documents that assign specific areas of responsibility to specific accountable persons — and accountability maps showing how responsibility is distributed across the organisation. Where an AI system operates within an accountable person's area of responsibility, the accountability for its governance is, by construction, personal and pre-assigned.

The mechanism has teeth through remuneration. The regime requires entities to defer at least 40 per cent of an accountable person's variable remuneration for a minimum of four years, and to maintain a remuneration policy that reduces variable remuneration where an accountable person fails to meet their accountability obligations. Accountability is therefore not symbolic. It is financial, deferred, and clawback-capable.

Layered above this is the directors' duty of care and diligence under section 180 of the Corporations Act. This is where the AI readiness question becomes acute. A director who accepts a management assurance that 'AI is being used responsibly' — without the evidence that demonstrates it — has not discharged the duty of care. They have accepted a representation they cannot test. APRA's own finding that boards rely too heavily on vendor presentations is, in this light, not merely a comment on competence. It is a description of personal liability being incurred without evidence to discharge it.

> *The accountability for an AI system's governance is personal, pre-assigned, and financially deferred. An assurance that cannot be evidenced does not discharge it.*

SECTION FIVE

# The Second *Jurisdiction.*

For any regulated organisation whose AI systems affect people in the European Union, a second and highly specific obligation applies: the EU AI Act. Its requirements for high-risk systems are the most detailed any jurisdiction has yet imposed, and the penalties for breach are correspondingly severe — up to €35 million or 7 per cent of worldwide annual turnover. Each requirement, unlike a statement of principle, implies evidence an examiner can demand.

AI systems used for credit scoring, risk assessment, and pricing in life and health insurance are classified as high-risk under Annex III. High-risk systems are subject to the Act's most substantive requirements, set out in Articles 9 to 15. These are not aspirational principles; they are a structured compliance burden, each element of which implies evidence an examiner can demand.

The Seven High-Risk Obligations — EU AI Act, Articles 9–15

| | |
|---|---|
| **Art. 9** | Risk Management System |
| **Art. 10** | Data & Data Governance |
| **Art. 11** | Technical Documentation |
| **Art. 12** | Record-Keeping / Logging |
| **Art. 13** | Transparency to Deployers |
| **Art. 14** | Human Oversight |
| **Art. 15** | Accuracy, Robustness & Cybersecurity |

*Figure 2 — The seven high-risk obligations of the EU AI Act, Articles 9–15.*

The seven obligations are: a **risk management system** maintained across the lifecycle (Article 9); **data and data governance** addressing quality and bias (Article 10); **technical documentation** sufficient to demonstrate compliance (Article 11); automatic **record-keeping** and logging (Article 12); **transparency** and provision of information to deployers (Article 13); **human oversight** by natural persons during operation (Article 14); and **accuracy, robustness and cybersecurity** throughout the lifecycle, including resistance to data poisoning and adversarial attack (Article 15).

Two structural points govern readiness. First, the Act distinguishes the *provider* who develops or places a system on the market from the *deployer* who uses it under its own authority — and

obligations differ accordingly. A regulated enterprise using a vendor's AI is typically a deployer but cannot assume that the provider has discharged all obligations. Second, the high-risk obligations for Annex III systems have been deferred — under the Digital Omnibus provisional agreement of May 2026, pending formal adoption — from August 2026 to December 2027. That is preparation time, not an exemption.

SECTION SIX

# The Baseline.

Beneath both the Australian stack and the EU obligations sits a cross-jurisdiction baseline: *ISO/IEC 42001*, the world's first AI management system standard, published in December 2023 by ISO and IEC. It applies to any organisation that provides or uses AI systems, of any size or sector, and it provides the management-system substrate on which both regulatory regimes can be operationalised.

Its structure follows the harmonised Annex SL format shared by ISO 27001 and ISO 9001 — Clauses 4 to 10 covering context, leadership, planning, support, operation, performance evaluation, and improvement — and it adds Annex A, a catalogue of thirty-eight AI-specific controls spanning impact assessment, data management, lifecycle, third-party and supply-chain governance, transparency, and monitoring and logging. Controls are selected through a Statement of Applicability, and the system is maintained through the Plan-Do-Check-Act cycle across the full lifecycle, from design to decommissioning.

ISO/IEC 42001 is necessary, but not sufficient, for readiness — and the distinction is the point of this section. In ISO's own framing, an organisation conforming to the standard 'can generate evidence of its responsibility and accountability' regarding its AI systems. It provides the *system* within which evidence is generated. It does not, by itself, constitute the readiness evidence a supervisor demands for a specific obligation. A certificate demonstrates that a management system exists; it does not demonstrate that, for a given deployed model, the pre-authorisation record exists, the monitoring is continuous, or the board provided effective challenge. The baseline gives an organisation the machinery. Readiness is what the machinery is required to produce.

SECTION SEVEN

# What Readiness *Actually Requires.*

Readiness is the capacity to produce specific evidence for specific obligations. The table below maps the regulatory obligation to the evidence that demonstrates readiness and to the MANDATE Suite instrument that produces it. It is the operational heart of this paper: the difference between a policy that asserts a control and an architecture that evidences it.

*Each row answers one question: if a supervisor asked you to prove this control operates, what would you produce?*

| Regulatory obligation | Readiness evidence required | MANDATE instrument |
|---|---|---|
| **Aggregate exposure (CPS 220; board oversight)** | A live inventory of every AI system and the total autonomous authority delegated across them | Autonomy Budget / ADAE scoring |
| **Pre-authorisation (FAR; s180 duty)** | A record, predating first action, of what each system was authorised to do, by whom, within what limits | Pre-authorisation record |
| **Continuous assurance (APRA; CPS 230)** | Continuous monitoring of model behaviour against approved parameters — not point-in-time testing | Continuous Admissibility Model |
| **Decision defensibility (EU Art. 12)** | An audit record committed before the action proceeds, not reconstructed afterward | Write-Before-Execute |
| **Board literacy (APRA finding)** | Evidence the board provided effective challenge, independent of vendor presentations | Board literacy / KRI reporting |
| **Supplier & concentration (CPS 230; CPS 234)** | A map of AI provider dependence, fourth-party exposure, and substitution planning | Supplier concentration mapping |
| **Contestability (ASIC REP 798)** | A mechanism for an affected customer to obtain an explanation and contest a decision | Contestability arrangement |

The pattern across the table is consistent. In every case, the regulatory obligation is satisfied not by a statement of intent but by an artefact that can be produced on demand. Readiness is the sum of those artefacts — and the organisation that holds them has closed the distance between its policy and its provable operation.

SECTION EIGHT

# The Readiness *Diagnostic.*

The most useful test of readiness is not whether an organisation has a governance framework. It is whether it could produce evidence under examination. Each of the following five questions is phrased as a supervisor would put it — and each maps to an obligation set out earlier in this paper. An organisation that can answer all five with an artefact, not an assertion, is examination-ready.

1. **Could you produce, today, a complete inventory of every AI system in operation and the total decision authority delegated across them?** (Aggregate exposure — CPS 220 and board oversight.)
2. **If a regulator asked who authorised a given AI system to act, within what limits, and on what date, could you produce the record — created before the system acted?** (Pre-authorisation — FAR and the s180 duty.)
3. **Can you demonstrate that your assurance detects model drift and behavioural change continuously, rather than at a point in time?** (Continuous assurance — APRA's assurance finding and CPS 230.)
4. **Could the board provide evidence that it provided an effective challenge on an AI risk, independent of management and vendor summaries?** (Board literacy — APRA's board finding.)
5. **If an AI decision materially affected a customer, could that person obtain an explanation and contest it — and could you evidence the arrangement?** (Contestability — the arrangement ASIC found no licensee had implemented.)

> *Readiness is not whether you can describe the control. It is whether you can produce it the moment a supervisor asks.*

ABOUT

# 42 Consulting.AI.

42 Consulting.AI builds constitutional AI governance for regulated enterprises. Its core work, the MANDATE Suite, is a comprehensive governance architecture comprising 64 normative instruments across five frameworks — governing the delegation of machine decision authority, the ethical and rights boundaries on its use, the management system that maintains it, and the regulatory extensions for EU and group structures. The Suite is designed for organisations in financial services, insurance, superannuation, healthcare, government, and other regulated sectors operating autonomous AI systems at scale.

The architecture is supported by an academic preprint series published on Zenodo, which establishes the theoretical basis for each governance mechanism, and by the forthcoming book *Governance-First AI*. The Suite's distinguishing conviction is that AI governance at the depth regulated enterprises require cannot be delivered by a compliance framework alone — it requires a constitutional architecture, adopted by the board, enforced by the infrastructure, and continuously tested against current conditions.

**Dr M Maruf Hossain, PhD, GAICD** · Chief AI Strategist
enquire@42consulting.ai · www.42consulting.ai

The preprint series is available at: zenodo.org/communities/mandate/records

REFERENCES

# References *and Sources.*

All regulatory obligations cited in this paper were verified at the primary source prior to inclusion. Where an obligation derives from a binding standard, it is identified as such; advisory guidance is distinguished from an enforceable requirement.


## AUSTRALIAN REGULATORY SOURCES

APRA (2026). *Letter to Industry on Artificial Intelligence (AI).* Australian Prudential Regulation Authority, 30 April 2026.

APRA. *Prudential Standard CPS 220 Risk Management; CPS 230 Operational Risk Management (effective 1 July 2025); CPS 234 Information Security; CPS 510 Governance; CPS 511 Remuneration; CPG 235 Managing Data Risk.* Australian Prudential Regulation Authority.

APRA & ASIC (2024). *Financial Accountability Regime: Information for Accountable Entities.* Commenced 15 March 2024 (banking); 15 March 2025 (insurance and superannuation).

ASIC (2024). *Report 798: Beware the Gap — Governance Arrangements in the Face of AI Innovation.* Australian Securities and Investments Commission, 29 October 2024.

## EUROPEAN UNION

European Union (2024). *Regulation (EU) 2024/1689 (Artificial Intelligence Act), Articles 9–15 and Article 99; Annex III.* Official Journal of the European Union, 13 June 2024.

## INTERNATIONAL STANDARD

ISO/IEC (2023). *ISO/IEC 42001:2023 — Information technology — Artificial intelligence — Management system.* International Organization for Standardization and International Electrotechnical Commission.

## MANDATE PREPRINT SERIES

Hossain, M. M. (2026a). *The Autonomy Budget: A Portfolio-Level Framework for Governing Delegated Machine Authority in Regulated Enterprises.* Zenodo. https://doi.org/10.5281/zenodo.20480491

Hossain, M. M. (2026b). *Continuous Admissibility: A Temporal Governance Framework for Delegated Machine Authority in Regulated Enterprises.* Zenodo. https://doi.org/10.5281/zenodo.20580716

Hossain, M. M. (2026c). *Write-Before-Execute: A Governance Standard for Evidentially Defensible AI Decision Logging in Regulated Enterprises.* Zenodo. https://doi.org/10.5281/zenodo.20603712

# Consulting.AI

*Scaling Intelligence*

enquire@42consulting.ai · www.42consulting.ai