APRA's AI governance wake-up call for financial services leaders

Letter to Industry on AI

Oversight must catch up and then keep pace with AI adoption

The better the question. The better the answer.
The better the world works.

EY

Shape the future with confidence

What does the Letter say?

APRA has delivered a clear wake-up call to banks, insurers and super funds via its Letter to Industry on Artificial Intelligence (30 April 2026). The regulator's targeted review found AI adoption accelerating across APRA-regulated entities.

As institutions move to harness AI to create productivity and customer benefits, AI is fundamentally changing the way the workforce operates and services are delivered. The frenetic pace of this change is putting tried-and-tested governance, risk management, security and assurance systems under pressure.

APRA says that treating AI risk as "just another technology" results in gaps in lifecycle management and oversight. And the regulator expects urgent action to close these governance gaps. If AI risks aren't managed proportionately, institutions can expect increased supervisory focus and enforcement actions.

Boards are now squarely accountable for ensuring AI is governed and controlled under existing standards, particularly CPS 220, CPS 230 and CPS 234.

What opportunity does it present?

APRA's letter highlights a strategic opportunity for forward-looking boards and executives. AI undeniably brings significant benefits in efficiency and customer experience, and APRA itself cautions that failing to embrace AI could leave firms at a strategic disadvantage. The opportunity is to close the AI oversight gap – control risk and unlock value simultaneously by embedding AI risk management as an integral part of AI adoption.

Experience across EY client engagements shows organisations that proactively strengthen AI governance can unlock innovation safely and build stakeholder trust, experiencing around 30% fewer AI-related risks through stronger oversight and controls.

What issues should you be aware of?

APRA's review identified several common weaknesses that boards and senior executives should address:

- Board literacy and oversight gap: While boards are focused on AI's potential, many lack the technical literacy to challenge how models behave in practice, leading to over-reliance on vendor and management assurances. Board cycles and point-in-time reporting are no longer suitable for managing AI risks, which can emerge and be exploited in near real-time, requiring continuous visibility and faster escalation than occurs in traditional board processes.
- AI-driven cyber threats: AI is reshaping the cyber threat landscape, expanding increasing attack pathways and compressing response times. Frontier AI models in the hands of attackers will enable accelerated identification and exploitation of vulnerabilities (e.g., Anthropic's Mythos model). Entities must continually uplift security capabilities in this evolving threat environment. Legacy cyber controls are struggling to keep pace, requiring faster, more adaptive defence and response capabilities.
- Governance maturity lag: AI deployment is outpacing governance. Many firms still treat AI as conventional IT, overlooking unique risks (e.g., adaptive models, bias, privacy) and leaving gaps in AI lifecycle management.
- Concentration and opacity risk: Institutions are increasingly dependent on a small number of AI providers, often without credible exit or fallback plans. At the same time, limited visibility into upstream models and data heightens systemic risk. The danger is that failures in opaque third-party AI could cascade into critical operations.
- Assurance limitations: Point-in-time assurance is ill-suited for dynamic, self-learning AI systems. Few organisations have continuous monitoring in place, leaving risks such as model drift or bias undetected and boards without timely insight.

What this means for boards and senior executives

APRA's letter signals the need to:

- **Understand where the organisation sits on the AI adoption curve**
  Not all AI is equal. The type, scale and criticality of AI in use determines the required control environment. As organisations progress from productivity tools and copilot to embedded, agentic and self-directing systems, the risk profile changes materially. Boards should have clear visibility of this progression and ensure governance, cyber security, third-party and resilience controls scale proportionately, avoiding blind spots as AI becomes more autonomous and operationally critical infrastructure.
- **Treat AI as operationally critical — even where adoption is cautious**
  The accelerating use of advanced AI by threat actors is raising the threat level for all organisations. Even firms in the early stages of AI adoption are exposed through vendor-embedded AI and AI-enabled attack techniques. Institutions will need to prioritise and strengthen cyber security to respond to evolving threats, which may include using AI to more rapidly identify and resolve threats.
- **Shift to continuous, lifecycle-based oversight and assurance**
  Approving individual AI use cases and relying on point-in-time reviews or traditional change controls is no longer sufficient for adaptive AI systems. Boards should expect ongoing insight aligned with risk appetite, including continuous monitoring of AI behaviour and performance (to avoid drift and bias), security weaknesses, dependencies, third-party reliance and concentration risk. Security needs strong foundational controls and basic cyber hygiene to improve containment speed and recovery time, and minimise operational impact.
- **Manage AI risk holistically across the enterprise**
  The use of AI intensifies a wide range of risks, including data, model, privacy, compliance, conduct, technology, cyber and third-party risks. Institutions need a holistic AI standard and oversight requires monitoring and reporting processes that tie together these individual risk classes in aggregate.

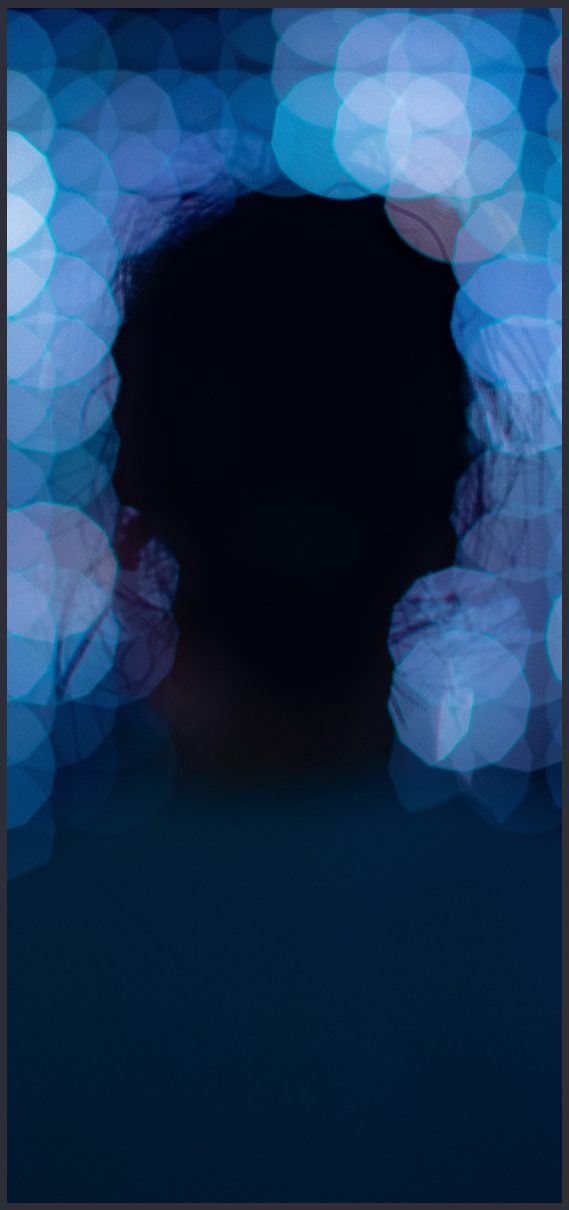

Key questions for boards and executives

To gauge readiness, boards and senior executives should ask themselves:

1. Where is AI already embedded in our critical processes or third-party platforms, and do we have full visibility of these AI uses?
2. Can the board clearly articulate how AI risks are governed, monitored and escalated within our risk framework?
3. Are our assurance and audit approaches continuous and adaptive, or still point-in-time and reactive?
4. Where are we over-reliant on specific AI vendors, and what is our fallback if those providers fail or change?
5. If APRA came knocking, could we demonstrate proportionate controls and evidence under CPS 220 (risk management), CPS 230 (operational risk) and CPS 234 (resilience and security) for our AI deployments?
6. Does the organisation's current cyber risk assessment adequately take into account AI-enabled threats and evolving developments in Frontier AI?

Your immediate priorities

- **Map AI exposure:** Identify all AI use cases and dependencies across critical services, including vendor-embedded AI, to establish a clear inventory and risk profile.
- **Strengthen board capability:** Lift board and executive AI literacy to enable informed oversight, challenge and decision-making.
- **Embed AI into existing frameworks:** Integrate AI risk into your CPS 220, CPS 230 and CPS 234 operational risk and resilience frameworks, rather than inventing parallel structures. Align AI oversight with existing risk appetite, controls and reporting.
- **Redesign assurance for AI:** Move to continuous monitoring and validation of key AI models (e.g., drift, bias and stress testing) and uplift internal audit AI capability to keep assurance aligned with AI's pace of change.

How EY can help

EY teams support APRA regulated financial institutions to adopt AI with confidence, meeting heightened business and regulatory expectations across governance and assurance.

We support boards and executives to scale AI safely and defensibly, and to accelerate workforce adoption through practical enablement and upskilling so AI benefits are realised without outpacing the control environment.

We also assess AI dependency, concentration and third-party risks across critical operations and service providers, support operational resilience scenario testing, and deliver independent AI assurance, governance maturity assessments and APRA readiness reviews.

Our multidisciplinary approach brings together strategy advice, technology consulting, cybersecurity, data and analytics, risk, compliance and assurance capabilities. This enables us to address AI not as a standalone technology initiative, but as an enterprise-wide risk, resilience and transformation agenda, anchored in board oversight and regulatory confidence.

We bring proven lessons and draw on our own experience as 'Client Zero' for AI transformation to help clients shift skills, ways of working and operating model, while embedding risk management as part of AI adoption rather than as a downstream control.

How we support you across the AI lifecycle

AI strategy and governance foundations

We work with your board and executives to define a clear AI vision aligned to business strategy, customer outcomes and risk appetite. We can help the business to prioritise AI use cases, assess readiness across data, technology and controls, and establish governance frameworks that set clear accountability, decision rights, escalation pathways and board reporting. This ensures AI adoption is valuable, defensible, compliant and scalable from day one.

Data, technology, cyber security and resilience readiness

We help you strengthen the data, technology and cyber security foundations that underpin safe and resilient AI adoption. Our focus is on ensuring AI systems are secure, explainable and resilient by design, including protection against data leakage, model manipulation and adversarial attacks, and the ability to sustain and recover AI-enabled services during disruption.

Responsible and secure AI implementation

We help you embed responsible AI and security-by-design practices across the full AI lifecycle. We support strong model governance, validation and monitoring, clear accountability for AI outcomes, and alignment with regulatory expectations, ethical standards and community trust.

AI governance, third-party risk and operational resilience uplift

We can help you to design, implement and evidence AI governance that stands up to board and regulatory scrutiny, including alignment to CPS 220, CPS 230 and CPS 234. This includes clarifying accountability and decision rights, defining risk appetite and escalation triggers, setting fit-for-purpose board reporting and evidence packs, and uplifting assurance so oversight keeps pace with adaptive AI.

The outcome is clear, defensible assurance that AI risks, including cyber security, third-party and operational resilience impacts, are identified, managed and continuously monitored, giving boards, executives and regulators confidence as AI adoption adapts and scales.

EY Contacts

For further discussion on APRA's AI letter and its implications for your organisation, please contact:

Rody Posthuma
Partner, & AI Oceania Leader, Financial Services Assurance, Ernst & Young, Australia
rody.posthuma@au.ey.com

Goran Stojanoski
Partner, Financial Services AI Risk Management Leader, Ernst & Young, Australia
goran.stojanoski@au.ey.com

Scott Glover
Partner, Strategy & Execution, Ernst & Young, Australia
scott.glover@parthenon.ey.com

John Hare
Associate Partner, Cyber & Technology Consulting, Ernst & Young, Australia
john.c.hare@au.ey.com

Yi Fang
Partner, Cyber & Technology Consulting, Ernst & Young, Australia
yi.fang.chua@au.ey.com

David Millar
Partner, AI Model Governance, Ernst & Young, Australia
david.millar@au.ey.com

Robert Menzies
Partner, AI & Model Risk Management, Ernst & Young, Australia
robert.menzies@au.ey.com

Christina Larkin
EY Oceania Assurance Digital Trust Leader, Ernst & Young, Australia
christina.larkin@au.ey.com

EY | Building a better working world

EY is building a better working world by creating new value for clients, people, society and the planet, while building trust in capital markets.

Enabled by data, AI and advanced technology, EY teams help clients shape the future with confidence and develop answers for the most pressing issues of today and tomorrow.

EY teams work across a full spectrum of services in assurance, consulting, tax, strategy and transactions. Fueled by sector insights, a globally connected, multi-disciplinary network and diverse ecosystem partners, EY teams can provide services in more than 150 countries and territories.

All in to shape the future with confidence.

PH8951949
SCORE 111737-26-AUNZ
ED None

ey.com/au