APRA AI GOVERNANCE RESPONSE

From Observation to Action: How Songlines Control® Addresses the AI Governance Gaps Identified by APRA

Version 1.0 — May 2026

Cetus AI Pty Ltd | cetusai.com.au | contact@cetusai.com.au
Microsoft Azure Marketplace Partner | MACC Eligible | IRAP Aligned

In response to: APRA Letter to Industry on Artificial Intelligence (AI), 30 April 2026

Foreword

On 30 April 2026, the Australian Prudential Regulation Authority (APRA) published a letter addressed to all APRA-regulated entities — banks, insurers and superannuation trustees — titled *"Summary of common weaknesses and expectations for regulated entities."* The letter was signed by APRA Member Therese McCarthy Hockey and accompanied by an executive-level debrief attachment directed at CROs, CTOs and CISOs.

The letter was not a consultation. It was a direct statement of observed failures and a clear signal of enforcement intent. APRA stated explicitly that *"where entities fail to adequately identify, manage or control AI risks in a manner proportionate to their size, scale and complexity, we will take stronger supervisory action and, where appropriate, pursue enforcement."*

This white paper reviews each of the four major problem areas APRA identified, explains the specific risks and consequences of failing to address them, and demonstrates how Songlines Control® — Cetus AI's enterprise AI governance platform — provides the technical infrastructure required to meet APRA's expectations.

The APRA Letter: Context and Significance

APRA's April 2026 letter was the product of a targeted engagement exercise conducted in late 2025 across a group of selected large banks, insurers and superannuation trustees. The findings were published not just for the entities reviewed, but as a signal to the entire regulated sector.

APRA's key overarching observation was stark:

> *"Assurance practices are not keeping pace with the scale, speed and complexity of AI. Governance has not matured at the same pace [as AI adoption]."*

APRA identified four distinct problem areas, each with specific observations and minimum expectations. These are examined in detail below.

Problem Area 1: Information Security Practices Are Not Keeping Pace with AI Threats

What APRA Observed

APRA found that AI adoption is materially changing the cyber threat landscape for regulated entities. Common attack pathways identified include prompt injection, data leakage, insecure integrations, exploit injection, and the manipulation or misuse of autonomous AI agents. AI is shortening the attack cycle and increasing the speed, coordination and impact of cyber attacks.

APRA specifically called out that *"the use of enterprise AI tools by staff outside approved control frameworks is also a concern."* In most entities reviewed, preventative controls were lacking — organisations were relying on policy direction or after-the-fact detective measures rather than enforceable technical restrictions.

APRA also observed that identity and access management capabilities had not adjusted to non-human actors such as AI agents, and that security testing programmes had significant gaps in coverage of AI implementations.

The Business Consequence of Inaction

An uncontrolled AI environment creates a direct pathway for data exfiltration, model manipulation, and regulatory breach. Prompt injection attacks — where malicious instructions are embedded in user inputs to manipulate AI behaviour — are now a primary attack vector. An employee using an unapproved AI tool can inadvertently send sensitive customer data to an external model with no audit trail, no redaction, and no visibility for the security team. Under CPS 234, the entity is responsible for that breach regardless of whether the tool was officially sanctioned.

The consequences extend beyond security. An AI agent operating outside approved control frameworks can make decisions — in claims triage, loan processing, or customer interaction — that create regulatory exposure under the Privacy Act 1988, the Australian Consumer Law, and APRA's own conduct risk expectations.

What APRA Expects

APRA Expectation	Minimum Requirement
Operational resilience	Assess AI reliance implications; credible fallback processes where AI supports critical operations
Security controls	Address AI-specific threats including prompt injection, privileged access management, hardened configurations, and controls over agentic workflows
Security testing	Robust testing across AI-generated code, software components and libraries
Third-party risk	Ongoing consideration of third-party and concentration implications

How Songlines Control® Responds

Capability	How It Addresses the APRA Expectation	Business Impact
Prompt Injection Prevention	Real-time detection of prompt injection attempts at the proxy layer, with configurable warn/ redact/block responses and a full injection audit log	Security teams have immediate visibility of injection attempts; malicious instructions are blocked before reaching the model
Shadow AI Prevention	Model access controls enforce allowlists and blocklists per team and workflow; unapproved models are blocked at the proxy layer, not just by policy	Staff cannot use unapproved AI tools without the attempt being logged and blocked — preventative control, not detective
Zero-Trust Architecture	All integrations are deliberate, observable, and governed; SHA-256 hashed API key validation and TLS 1.2+ in transit	Every AI connection is authenticated and logged; no open-ended connections that create shadow AI risk
RBAC for Non-Human Actors	Role-based access controls enforced at the system, user, and agent level — including autonomous AI agents	Identity and access management explicitly covers AI agents, directly addressing APRA's observation about non-human actors
Immutable Audit Trail	Every prompt, model selection, and response cryptographically signed and logged; records cannot be modified or deleted	Security testing and incident response have a complete, tamper-proof record of all AI activity

Problem Area 2: Governance Maturity Is Lagging AI Adoption

What APRA Observed

APRA found that while most entities recognise that existing prudential standards apply to AI risk, *"few have operationalised governance in practice."* APRA observed a tendency to treat AI risk as *"just another technology"* — missing key differences such as the distinct characteristics of predictive systems, adaptive model behaviour, ethical considerations including inherent bias, and privacy and data risks.

Specific gaps identified include weak controls over post-deployment monitoring, weak model behaviour monitoring, change management, and decommissioning of AI capabilities. APRA also observed that Boards are still developing the technical literacy required to provide effective challenge on AI-related risks, with an *"overreliance on vendor presentations and summaries without sufficient examination of key AI risks."*

The Business Consequence of Inaction

A Board that cannot see its AI estate cannot govern it. Without an inventory of AI tooling and use cases, directors cannot meet their obligations under the Corporations Act 2001 to exercise reasonable care and diligence. If an AI system causes customer harm — through biased decisions, incorrect outputs, or privacy breaches — and the Board cannot demonstrate it had visibility of that system and its risks, individual directors face potential personal liability.

For CROs and CISOs, the absence of operationalised governance means that AI risk cannot be reported to the Board in a meaningful way. Risk appetite settings are theoretical rather than enforced. Post-deployment monitoring is manual, intermittent, or absent.

What APRA Expects

APRA Expectation	Minimum Requirement
Governance frameworks	Policy, standard, guidance and reporting lines to promote safe, responsible and sustainable AI adoption
Lifecycle accountability	Ownership and accountability from design through deployment, monitoring and decommissioning
AI inventory	A maintained inventory of AI tooling and AI use cases
Human oversight	Human involvement for high-risk decisions and accountability
Staff training	Training and education on AI use, misuse, limitations and secure practices

How Songlines Control® Responds

Capability	How It Addresses the APRA Expectation	Business Impact
Centralised Model Registry	Every AI model in use across the organisation is registered and tracked in real time — including third-party models embedded in software and platforms	Boards and executives have a live, accurate AI inventory at all times; the "shadow AI" problem is eliminated at the infrastructure layer
Executive Dashboard	Real-time board-level visibility of AI spend, model usage, policy compliance status, and risk indicators across the entire organisation	Board Directors can view the status of organisational AI governance in an instant — in real time — ensuring they are compliant with their director responsibilities under the Corporations Act
Policy-First Architecture	Governance policies are evaluated at the point of execution, before any request reaches a model — not after the fact	Risk appetite is enforced technically, not just documented in policy; post-deployment monitoring is continuous and automated
Human-in-the-Loop (HITL) Gates	Configurable approval checkpoints for high-risk AI actions, with multi-channel routing, timeout rules, and a full approval audit trail	High-risk decisions — in claims triage, loan processing, customer interaction — require human sign-off before AI output is acted upon
Model Behaviour Monitoring	Continuous monitoring of model performance, response patterns, and anomalous behaviour across the AI lifecycle	Model drift, bias indicators, and failure modes are detected continuously — not through point-in-time sampling

Problem Area 3: Supplier Risk Management Is Insufficient

What APRA Observed

APRA observed some entities *"heavily dependent on a single provider for multiple AI use cases"* with few demonstrating robust contingency planning or tested exit and substitution strategies. Contractual arrangements often lagged practice, with limited evidence of specific provisions addressing audit rights, model updates and deviations, incident notification, or changes to data handling.

APRA also noted that AI capabilities are increasingly embedded within software, platforms and developer tools, making upstream dependencies — foundation models, training data sources, and fourth-party service providers — opaque. This limits entities' ability to independently assess model performance, bias, resilience and security.

The Business Consequence of Inaction

Provider concentration risk in AI is qualitatively different from traditional IT concentration risk. When an AI provider changes a model — updates weights, adjusts behaviour, or modifies training data — the entity's AI outputs can change materially without warning. A claims triage model that was performing within acceptable parameters last month may be producing systematically different outcomes this month, with no notification from the vendor and no mechanism for the entity to detect the change.

The opacity of embedded AI is equally dangerous. An entity using a software platform with an embedded AI component may have no visibility of which foundation model is being used, what data it was trained on, or how it handles sensitive customer information. Under CPS 234 and the Privacy Act, the entity cannot outsource its accountability for that data.

What APRA Expects

APRA Expectation	Minimum Requirement
Supply chain visibility	Map and maintain visibility over the full AI supply chain including fourth-party dependencies
Contractual governance	Arrangements providing transparency, auditability and assurance over AI services
Model behaviour visibility	Ability to understand model behaviour, material changes, performance issues and outcomes
Concentration risk management	Plausible failure scenarios and credible substitution, portability or exit arrangements

How Songlines Control® Responds

Capability	How It Addresses the APRA Expectation	Business Impact
Provider-Agnostic Governance	The same policy engine applies to every model provider — Azure OpenAI, Anthropic, AWS Bedrock, Google Gemini, and local sovereign models — regardless of where the model is hosted	Entities are not locked into a single provider's governance framework; APRA-required controls apply uniformly across the entire AI supply chain
Intelligent Model Routing	Songlines Gateway enables dynamic routing between providers based on cost, latency, policy, or data residency — with no code changes required	Substitution of a critical AI provider becomes an operational decision, not a multi-month integration project; concentration risk is actively managed at the infrastructure layer
Third-Party Model Monitoring	Model behaviour, response patterns, and performance metrics are tracked continuously across all providers, including embedded and third-party models	Material changes in model behaviour — including silent model updates by vendors — are detected and logged; entities have the audit trail APRA requires
Immutable Audit Trail	Every interaction with every model is cryptographically signed and logged, regardless of provider	Entities can independently demonstrate model performance, bias indicators, and data handling to APRA supervisors — without relying on vendor-provided summaries
Sovereign Data Routing	PII and sensitive data is routed only to approved sovereign or on-premise models, enforced at the proxy layer	Fourth-party opacity is addressed at the infrastructure layer; sensitive customer data cannot reach an unapproved model regardless of what the calling application requests

Problem Area 4: Change Management and Assurance Are Not Sufficient for Dynamic AI

What APRA Observed

APRA observed that AI risks cut across multiple domains simultaneously — operational risk, cyber and information security, data governance, model risk, change control, legal and regulatory compliance, privacy, conduct risk, and procurement. Existing change and assurance management approaches are often fragmented and do not provide sufficient integrated assurance.

APRA specifically noted *"reliance on point in time and sample based assurance methods, despite these methods being ill suited to probabilistic models that learn, adapt and degrade over time."* Few entities had continuous validation or monitoring in place to detect model drift, bias, failure modes, or control breakdowns

in a timely manner. Internal audit and risk functions were found to lack the specialist skills and tools required to engage in AI assessment, particularly where agentic behaviour, automated decision making, or AI-assisted code generation were involved.

The Business Consequence of Inaction

A probabilistic AI model is not a static system. It can degrade, drift, or produce systematically biased outputs over time — and point-in-time audits will not detect this. An annual model review that finds a claims triage model performing within parameters provides no assurance about what the model was doing in the eleven months between reviews.

For regulated entities, this creates a specific liability: if a model produces discriminatory outcomes — in lending decisions, insurance pricing, or claims assessment — and the entity cannot demonstrate continuous monitoring and timely detection, the regulatory and legal exposure is significant. APRA, ASIC, and the OAIC all have jurisdiction over different dimensions of this failure.

The fragmentation of assurance across multiple teams — risk, audit, cyber, data governance, legal — means that no single function has a complete picture. Each team sees its own slice of the AI risk landscape, and the gaps between those slices are where the most significant risks accumulate.

What APRA Expects

APRA Expectation	Minimum Requirement
Control frameworks	Globally recognised control libraries and change control for AI implementations
Integrated assurance	Assurance across cyber security, data governance, model performance, operational resilience, privacy and conduct risks
Technical capability	Second line risk and internal audit functions with technical capability to independently assess AI systems including agentic workflows
Continuous monitoring	Monitoring proportionate to criticality, including model purpose, limitations, explainability and potential customer impacts

How Songlines Control® Responds

Capability	How It Addresses the APRA Expectation	Business Impact
Continuous Model Monitoring	Real-time monitoring of model performance, response patterns, anomalous behaviour, and policy compliance — not point-in-time sampling	Model drift, bias indicators, and control breakdowns are detected continuously; audit and risk functions have live data, not historical snapshots
Integrated Compliance Reporting	A single pane of glass across cyber security controls, data governance, model performance, PII handling, and policy compliance	Risk, audit, cyber, and legal teams all draw from the same authoritative data source; fragmentation is eliminated
Exportable Audit Trail	The immutable audit trail is exportable in formats suitable for APRA supervisory review, internal audit, and second-line risk assessment	Internal audit and risk functions have the technical evidence base required to independently assess AI systems — without requiring specialist AI engineering skills
Agentic Workflow Controls	Governance policies apply to autonomous AI agents and agentic workflows, not just individual user prompts	Agentic behaviour is governed at the infrastructure layer; second-line risk functions can assess and report on agent activity through the same dashboard as all other AI activity
PII Auto-Redaction with Audit	Automated detection and redaction of 20+ PII entity types, with a full redaction audit trail showing what was detected, what was redacted, and when	Privacy and conduct risk assurance is automated and continuous; the entity can demonstrate to the OAIC and APRA that PII controls are operating as intended

The Regulatory Trajectory: What Comes Next

APRA's April 2026 letter is not the end of the regulatory journey — it is the beginning of a more active supervisory phase. APRA stated explicitly that it is *"currently finalising its forward plan with regards to supervision of AI risks, taking a proportional approach to entity prudential reviews, thematic activities and AI supplier engagement."*

The direction of travel is clear. APRA has moved from observation to expectation, and the next step is enforcement. Entities that cannot demonstrate operationalised AI governance — not just documented policy — will face stronger supervisory action.

This trajectory is reinforced by parallel regulatory activity. ASIC's October 2024 report REP 798 *Beware the Gap* identified the same governance failures across AFS licensees. ASIC's 2026 Key Issues Outlook named agentic AI as a priority enforcement area. ASIC's May 2026 open letter to all AFS licensees required board-level AI risk discussions to be documented. The Australian Cyber Security Centre has published specific guidance on frontier AI model risks. The Council of Financial Regulators is actively coordinating across APRA, ASIC, the RBA, and the Treasury on AI policy.

For APRA-regulated entities that are also AFS licensees — which includes most banks, insurers, and superannuation trustees — the combined regulatory pressure from APRA and ASIC creates a compliance imperative that cannot be deferred.

Getting Started with Songlines Control®

Songlines Control® is available today on the Microsoft Azure Marketplace with a **30-day free trial**. Enterprise customers with a Microsoft Azure Consumption Commitment (MACC) can apply their pre-committed Azure spend directly to Songlines subscriptions.

For APRA-regulated entities, Cetus AI offers a **two-week Proof of Concept (POC)** engagement: deploy Songlines Control® in your environment, instrument your key AI workloads, and receive a concrete data report and recommended governance strategy for your leadership team — at no risk, within two weeks.

Product	Core Capability	Pricing
Songlines Control®	AI inventory, cost attribution, policy enforcement, PII redaction, immutable audit trail, continuous monitoring	AUD $9,995/ month
Songlines Gateway	Everything in Control, plus intelligent model routing across all providers to reduce token spend and manage concentration risk	AUD $19,995/ month
Songlines Platform	Full sovereign AI stack — on-premise or private cloud, complete data residency, air-gap capable	Contact for pricing

Both Songlines Control® and Songlines Gateway are IRAP-aligned, Australian data residency capable, and designed specifically for the regulatory environment that APRA-regulated entities operate in.

To arrange a demonstration or POC engagement, contact **contact@cetusai.com.au** or visit **cetusai.com.au**.

Summary: APRA's Expectations and Songlines Control® Capabilities

APRA Problem Area	Key APRA Expectation	Songlines Control® Capability
Information Security	Controls over agentic workflows; shadow AI prevention; prompt injection defence	Prompt injection prevention; shadow AI blocking; zero-trust architecture; RBAC for AI agents
Governance Maturity	AI inventory; board-level oversight; human-in-the-loop; policy frameworks	Centralised model registry; executive dashboard; HITL gates; policy-first architecture
Supplier Risk	Supply chain visibility; concentration risk management; model behaviour monitoring	Provider-agnostic governance; intelligent routing for substitution; third-party model monitoring
Change Management & Assurance	Continuous monitoring; integrated assurance; technical capability for audit	Continuous model monitoring; integrated compliance reporting; exportable audit trail; agentic workflow controls

Version 1.0 — May 2026 | © 2026 Cetus AI Pty Ltd | cetusai.com.au