Shoring up AI Governance: How Songlines Control® Closes Critical Gaps Identified in ASIC's REP 798 Report

Published by: Cetus AI
Date: May 2026
Reference: ASIC Report 798 — Beware the Gap: Governance Arrangements in the Face of AI Innovation (October 2024) [1]

Executive Summary

In October 2024, the Australian Securities and Investments Commission published one of the most significant regulatory warnings in the history of Australian financial services technology governance. Report 798 — Beware the Gap — reviewed 23 AFS and credit licensees and found a consistent and dangerous pattern: the pace of AI adoption was outrunning the governance and risk management frameworks designed to control it.

ASIC's executive summary was unambiguous about the stakes:

> "Potential harms [from unmanaged AI] include bias and discrimination, provision of false information, exploitation of consumer vulnerabilities and behavioural biases, and the erosion of consumer trust."
> — ASIC REP 798, Executive Summary

These are not theoretical risks. They are the direct consequences of deploying AI without adequate governance — consequences that expose consumers to harm and expose licensees to regulatory enforcement, reputational damage, and director liability.

Eighteen months on, the regulatory environment has intensified. ASIC's 2026 Key Issues Outlook [2] identifies "advanced technology harming consumers (including agentic AI)" as a priority concern, noting "variable maturity in how businesses manage AI governance risks." In April 2026, APRA issued a formal letter to all regulated entities [3] calling for "a step-change in AI-related risk management and governance," warning that "where entities fail to adequately identify, manage or control AI risks... we will take stronger supervisory action and, where appropriate, pursue enforcement." In May 2026, ASIC issued an open letter to all AFS licensees and market participants [4] requiring boards and senior executives to table and discuss AI-related cyber risks at their ultimate board and risk governance committees.

The governance gap ASIC identified in 2024 is now a governance chasm — and regulators are actively closing in on organisations that have not addressed it.

Songlines Control® was built specifically to close that gap. This paper maps each of ASIC's eight findings to the specific capabilities Songlines Control® provides, and explains what those capabilities mean in practice for boards, risk officers, and compliance teams.

What ASIC Found: Eight Findings, Four Structural Failures

ASIC's review of 624 AI use cases across 23 licensees produced eight findings that can be grouped into four interconnected structural governance failures.

Gap 1: No Inventory, No Visibility

ASIC Findings 1 & 6

Many licensees could not readily identify all AI use cases operating within their organisation. ASIC noted that "a lack of an AI inventory, or the recording of models in several dispersed model registers" meant that boards and senior management did not have clear visibility of their AI use. One licensee, when asked to respond to ASIC's information request, discovered models missing from its own register.

> "Licensees and their boards may not have clear visibility of their AI use."
> — ASIC REP 798, Finding 6

Without a complete, real-time inventory of every AI model in use — including third-party models — it is impossible to govern what you cannot see. For board directors, this is not merely an operational inconvenience. Under the Corporations Act 2001, directors have a duty to exercise reasonable care and diligence. A director who cannot answer the question "what AI is our organisation using and how is it being governed?" is exposed to personal liability if that AI causes consumer harm.

ASIC Concern	Songlines Control® Response	Business Impact
No AI inventory or dispersed model registers	Centralised model and provider registry with real-time status, sovereignty flags, and viability classification	Compliance teams can produce a complete AI asset register for a regulator within minutes, not weeks
Boards lack visibility of AI use	Executive dashboard with request volumes, cost, and policy events by model — updated in real time	Board Directors can view the status of organisational AI governance in an instant, ensuring they are meeting their director responsibilities under the Corporations Act
Models missing from registers	Every request automatically attributed to a registered model; unregistered models cannot route through the platform	Eliminates the risk of shadow AI — models deployed without governance oversight — which is the single most common cause of the governance gap ASIC identified
No consistent AI strategy	Centralised policy framework applied across all models and providers	Demonstrates to ASIC that AI governance leads deployment, not the reverse

Gap 2: Governance Lagging Deployment — The 'Secret Sauce'

ASIC Finding 7

ASIC's central case study described a licensee that deployed a credit scoring model with "limited understanding" of the third-party platform used, "incomplete model documentation with missing critical elements," and "poor governance and a lack of a monitoring process." An internal report described the model as a "black box with no ability to explain the variables in the scorecard." The licensee continued using the model for months before replacing it.

This is the governance gap in its most dangerous form: AI deployed in production, affecting consumer outcomes, with no controls in place. ASIC identified two licensees in this situation and described them as representing "the greatest source of risk."

The governance gap described by ASIC — where AI is deployed before governance frameworks are in place — is structurally impossible in a Songlines Control® deployment. The platform enforces a **policy-first architecture**: AI requests are evaluated against the organisation's governance policies before they are routed to any model. A request that would violate a policy is blocked, modified, or escalated for human review — it does not reach the model.

Policy Type	What It Does	Business Impact
PII Detection	Detects embedded personally identifiable information and prevents sensitive customer data from being sent to any AI model without authorisation	Boards and Senior Officers know they are protected from potential prosecution for failing to protect consumer PII data under the Privacy Act 1988
Sovereignty Policy	Ensures regulated data does not leave Australian data residency — enforced at the infrastructure layer, not the application layer	Satisfies IRAP, APRA CPS 234, and APS AI Policy requirements for data sovereignty without relying on developer compliance
Model Restriction Policy	Prevents use of unapproved or unvetted AI models across the entire organisation	Eliminates the risk of staff using unauthorised AI tools (shadow AI) that have not been risk-assessed or approved
Approval-Required Policy	Mandates human review and sign-off for AI requests in designated high-risk use cases	Provides the "human in the loop" oversight that ASIC explicitly requires for consumer-impacting decisions
Cost Cap Policy	Prevents runaway AI spend by enforcing per-model, per-workflow, or per-organisation token and cost limits	Protects against AI budget overruns and ensures AI expenditure remains within board-approved risk appetite
Token Limit Policy	Controls the scope of AI interactions to prevent over-reliance on AI for complex decisions	Supports proportionate AI use and reduces the risk of AI outputs exceeding human oversight capacity

Critically, all six policy types apply to the **entire AI estate** — including third-party models — not just internally developed ones. An organisation cannot accidentally deploy a new AI use case outside the governance framework because the framework is enforced at the infrastructure layer, not the application layer.

Gap 3: Risk Assessed Through a Business Lens, Not a Consumer Lens

ASIC Finding 5

ASIC found that many licensees assessed AI risks from the perspective of business efficiency rather than consumer harm. Algorithmic bias was rarely proactively identified or tested for. Transparency and contestability — the ability of consumers to know AI was being used and to

challenge its outputs — were described as "relatively immature." Only 10 of the 23 licensees had documented requirements about disclosure of AI use to consumers. No licensees had implemented specific contestability arrangements.

> "We observed instances where licensees focused on business over consumer risk, or where the use of AI could have implications for licensees' compliance with existing conduct and consumer protection obligations, but this was not identified as a possible risk."
> — ASIC REP 798, Finding 5

Songlines Control® shifts this dynamic by making consumer-protective controls **technical rather than procedural**. PII detection policies, for example, do not rely on a developer remembering to check a policy document — they are enforced inline, on every request, before the request reaches any model.

The platform's immutable audit trail records, for every AI request: the timestamp, request ID, model used, workflow attribution, user attribution, token counts, cost, latency, status, policy decision (allowed, blocked, modified, or escalated), and IP address. Records are cryptographically signed and cannot be altered after the fact. This is the evidentiary foundation that ASIC's transparency and contestability requirements demand.

ASIC Concern	Songlines Control® Response	Business Impact
Risk assessed through business lens only	Inline policy enforcement for consumer-protective controls (PII, bias flags, approval gates)	Compliance and risk teams can demonstrate to ASIC that consumer risk is managed at the technical layer, not just in policy documents
No disclosure arrangements for AI use (only 10 of 23 licensees)	Immutable audit log records every consumer-impacting AI request with cryptographic signing	Provides the documented evidence base for consumer disclosure obligations — every interaction is traceable, timestamped, and exportable
No contestability arrangements (0 of 23 licensees)	Complete request-level audit trail with exportable CSV and PDF compliance reports	If a consumer contests an AI-influenced decision, the organisation can produce a complete, tamper-proof record of exactly what the AI did and why
Algorithmic bias not tested	Anomaly detection monitoring with alert history and severity classification	Ongoing monitoring for unexpected model outputs — the "monitoring process" that ASIC found absent in its case study
Governance fragmented across documents	Single enforcement layer with centralised policy management across entire AI estate	Eliminates the "evolving arrangements lead to complexity and fragmentation" failure pattern ASIC documented in Finding 6

Gap 4: Third-Party Model Risk — The Invisible Threat

ASIC Finding 8

30% of all AI use cases in ASIC's review used models developed by third parties. For 13 of the 23 licensees, more than half of their models were third-party developed. Yet many licensees did not have robust third-party management procedures. In some cases, licensees could not even identify the AI technique used in a third-party model because "vendors are hesitant to provide details beyond standard marketing literature."

So what? What is the actual business consequence of not governing third-party models?

The consequences are direct and serious. When a third-party AI model produces a biased output — denying a consumer credit, generating incorrect insurance advice, or miscalculating a risk score — the licensee is legally responsible for that outcome, regardless of whether the

model was built in-house or procured from a vendor. ASIC is explicit: "existing obligations apply to their use of AI." The vendor's intellectual property concerns do not transfer the licensee's regulatory obligations to the vendor.

In practice, this means:

A financial services organisation using a third-party AI model for credit decisioning that produces discriminatory outcomes can face enforcement action under the Australian Consumer Law, the National Consumer Credit Protection Act, and the Corporations Act — even if the organisation had no visibility into how the model worked. The defence "we didn't know how the vendor's model operated" is not available to a licensee. ASIC's Finding 8 makes clear that "better practices saw licensees setting the same expectations for models developed by third parties as for internally developed models."

The Songlines Control® platform's governance layer is **provider-agnostic**. Whether a model is Azure OpenAI, Anthropic Claude, AWS Bedrock, Google Vertex AI, or a sovereign on-premises deployment, every request passes through the same policy enforcement engine, generates the same telemetry, and contributes to the same audit trail. The licensee does not need to understand the internal operation of a third-party model to govern how it is used — they govern the inputs, outputs, and routing decisions at the infrastructure layer.

ASIC Concern	Songlines Control® Response	Business Impact
Third-party models ungoverned — vendors won't share model details	Provider-agnostic enforcement: same policy engine applies to all models regardless of vendor	Licensees meet their regulatory obligations for third-party AI without needing vendor cooperation or model transparency
No third-party management procedures	Unified governance layer captures telemetry and enforces policies across all providers	Demonstrates to ASIC the "same expectations for third-party models as internally developed models" that the report identifies as better practice
Licensees unaware of AI technique used in third-party models	Songlines Control® governs inputs and outputs — the regulatory obligation — not the internal model architecture	Shifts the governance question from "how does the model work?" (unanswerable) to "what did the model do, and was it within policy?" (fully auditable)
Concentration risk from heavy third-party reliance	Intelligent model routing enables switching between providers when a model underperforms, is unavailable, or fails policy checks	Reduces single-vendor dependency and ensures continuity of governance even when a third-party provider changes their model

The Regulatory Trajectory: Where Does Australia Stand?

ASIC's REP 798 was published in October 2024. In the eighteen months since, the regulatory environment has moved significantly — and the direction of travel is unambiguous.

Current position (as at May 2026): Australia does not yet have AI-specific legislation. In December 2025, the Australian Government accepted industry recommendations to pause mandatory guardrails for AI, instead relying on existing technology-neutral laws. However, this pause does not mean inaction — it means that existing obligations under the Corporations Act, the National Consumer Credit Protection Act, the Privacy Act, and the Australian Consumer Law apply in full to AI use, and regulators are actively enforcing them.

ASIC's current enforcement posture: ASIC's 2026 Key Issues Outlook [2] identifies "advanced technology harming consumers (including agentic AI)" as a named priority. ASIC has stated it will "take enforcement action if licensees' use of AI results in breaches of their obligations." In

May 2026, ASIC issued an open letter to all AFS licensees requiring boards to table AI-related risk discussions at their ultimate governance committees — not as a suggestion, but as an expectation of the regulator [4].

APRA's current enforcement posture: In April 2026, APRA issued a formal letter to all regulated entities — banks, insurers, and superannuation trustees — based on a targeted supervisory review conducted in late 2025 [3]. APRA's findings mirror ASIC's REP 798 almost exactly: "assurance practices are not keeping pace with the scale, speed and complexity of AI," and "many Boards are still developing the technical literacy required to provide effective challenge on AI related risks." APRA stated plainly: "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."

Is AI governance being audited? Yes. APRA conducted a targeted supervisory review of large banks, insurers, and superannuation trustees in late 2025, the results of which formed the basis of its April 2026 letter. ASIC has indicated it will "continue to monitor how our regulated population uses AI, and the adequacy of their risk management and governance processes." Both regulators have signalled that AI governance is now a standing item in their supervisory programmes — not a one-time review.

The trajectory: Mandatory AI guardrails have been paused, not abandoned. The Government has committed to reviewing the position as AI adoption accelerates. Organisations that invest in governance infrastructure now will be well-positioned to comply with any future AI-specific regulatory obligations — and, more immediately, to demonstrate to ASIC and APRA that their governance arrangements lead their AI use rather than lag it.

What Does IRAP Alignment Mean for Financial Services Organisations?

IRAP — the Information Security Registered Assessors Program — is an Australian Government initiative administered by the Australian Signals Directorate (ASD). It provides a framework for the independent assessment of the security of ICT systems that store, process, or communicate Australian Government information.

Why should a board director in financial services care about IRAP?

The answer lies in what IRAP alignment actually demonstrates. An IRAP-aligned system has been independently assessed against the Australian Government Information Security Manual (ISM) — the same security standard that applies to Defence, intelligence agencies, and critical infrastructure operators. For a financial services organisation, deploying an IRAP-aligned AI governance platform means:

1. Demonstrable security assurance. IRAP provides independent, third-party evidence that security controls are in place and effective — not just claimed in a vendor's marketing materials. This is precisely the kind of evidence-based assurance that APRA's April 2026 letter demands: "Governance should not rely only on assurances. It should be supported by evidence — test results, audit findings, lessons from incidents, and independent validation."

2. Data sovereignty protection. IRAP-aligned deployments in Australian Azure data centres (Australia East and Australia Southeast) ensure that sensitive customer data — including the AI request payloads, audit logs, and policy records managed by Songlines Control® — never leaves Australian jurisdiction. This directly addresses the Privacy Act 1988 requirements for cross-border data disclosure and the APRA CPS 234 requirements for information security.

3. Regulatory credibility. ASIC and APRA both reference ASD guidance as the benchmark for cyber resilience. An organisation that can demonstrate its AI governance platform is IRAP-aligned is demonstrating alignment with the same security framework that Australia's most security-conscious institutions use. In a regulatory examination, this is a materially stronger position than a vendor's self-certification.

4. Future-proofing for government and regulated sector contracts. For financial services organisations that interact with government agencies, superannuation funds, or defence-related entities, IRAP alignment is increasingly a procurement requirement. Deploying an IRAP-aligned AI governance platform positions the organisation to meet these requirements without additional remediation.

In short: IRAP alignment is not a technical checkbox. It is the independent evidence that a board director can point to when asked by ASIC or APRA whether the organisation's AI governance infrastructure meets the security standards expected of a regulated entity.

Getting Started

Songlines Control® is available now on the Microsoft Azure Marketplace with a 14-day free trial. Two plans are available:

- **Songlines Control®** — AI cost intelligence, policy enforcement, observability, and immutable audit trails. AUD $9,995/month.
- **Songlines Gateway** — Everything in Control, plus intelligent model routing across all providers to reduce token spend and eliminate single-vendor dependency. AUD $19,995/month.

Both plans include Australian data residency (Azure Australia East), IRAP-aligned documentation, and an exportable compliance evidence pack designed for ASIC and APRA regulatory examinations.

View Songlines Control® on the Microsoft Azure Marketplace

For enterprise deployments, IRAP assessments, board-level briefings on AI governance obligations, or co-sell enquiries through Microsoft's financial services vertical, contact sales@cetusai.com.au.

References

Cetus AI · cetusai.com.au · sales@cetusai.com.au
Page 1 of 1 — Version 2.0 — May 2026

1. ASIC Report 798 — Beware the Gap: Governance Arrangements in the Face of AI Innovation (October 2024) ↩
2. ASIC Key Issues Outlook 2026 (January 2026) ↩↩
3. APRA Letter to Industry on Artificial Intelligence (AI) (April 2026) ↩↩
4. ASIC Open Letter to AFS Licensees and Market Participants on AI and Cyber Resilience (May 2026) ↩↩