

WHITE PAPER

Progressive AI Governance

A Practical Framework for Australian Enterprise

Published by: Cetus AI Pty Ltd

Date: July 2026

Reference: CETUS-WP-006

Classification: Public

Audience: CFOs, CISOs, Chief Sustainability Officers, Board Directors

Executive Summary

Australian enterprises are deploying AI faster than they can govern it. The result is uncontrolled spend, unmanaged regulatory exposure, invisible Scope 3 emissions, and a growing population of unsanctioned AI tools operating outside organisational oversight.

The regulatory environment is accelerating this urgency. ASIC REP 798 requires demonstrable AI governance for financial services. AASB S2 mandates climate-related financial disclosures — including Scope 3 emissions from AI compute. The APS AI Policy 2024 sets governance expectations for government suppliers. The Privacy Act 1988 amendments demand evidence of data handling controls. ISO/IEC 42001 establishes the international standard for AI management systems.

The instinct of many organisations is to respond with policy documents. But policy without runtime enforcement is aspiration, not governance. The gap between what an organisation says it does and what its AI systems actually do at runtime is where regulatory risk accumulates.

"Governance should be a progression, not a project. Start with visibility. Expand to control. Lead with governance."

This white paper introduces the **Progressive AI Governance Framework** — a three-stage maturity model that allows organisations to start with immediate visibility (Observe), expand to runtime enforcement (Enforce), and scale to full sovereign orchestration (Orchestrate) — without disrupting production workloads at any stage.

Songlines Control® by Cetus AI is the only platform that delivers all three stages from a single runtime layer, a single audit trail, and a single policy engine. Nine purpose-built modules activate progressively as the organisation advances through the framework — ensuring that every investment in governance compounds rather than expires.

The Problem: AI Without Governance

The Visibility Gap

In most Australian enterprises, AI spend is invisible. It is buried in cloud infrastructure line items, distributed across departmental credit cards, and attributed to generic "compute" categories that make it impossible to determine what AI is costing the organisation, who is consuming it, and what value it is delivering. Gartner's 2025 AI Cost Management Survey found that 73% of enterprise AI spend is unattributed to a specific business outcome. Deloitte's 2025 Tech Value Survey reported that organisations allocate an average of 36% of their digital budgets to AI — yet fewer than one in five can produce a complete AI cost attribution report for their board.

Without visibility, governance is impossible. You cannot govern what you cannot see.

Shadow AI Proliferation

The average Australian enterprise has 30 or more unsanctioned AI tools in active use. Employees adopt ChatGPT, Claude, Gemini, Copilot, and dozens of vertical AI applications without IT approval, security review, or data classification assessment. Each tool represents a potential data leakage vector, a compliance exposure, and an uncontrolled cost centre.

Shadow AI is not a technology problem — it is a governance gap. Employees adopt these tools because they are productive. The solution is not to block AI adoption; it is to bring it under governance without destroying the productivity benefit.

Regulatory Acceleration

Australian regulators are moving faster than most organisations realise:

Regulation / Standard	Requirement	Implication for AI Governance
ASIC REP 798	Demonstrable governance of AI in financial services	Audit trail of every AI decision, model inventory, risk assessment
AASB S2	Climate-related financial disclosures including Scope 3	AI compute emissions must be measured and disclosed
APS AI Policy 2024	AI governance framework for government agencies	Suppliers must demonstrate governance controls
Privacy Act 1988	Protection of personal information	PII must not be sent to external AI models without controls
ISO/IEC 42001	AI management system standard	Systematic evidence of AI governance processes
Corporations Act s299(1)(f)	Director liability for material risk disclosure	AI-related climate and operational risks must be disclosed

The common thread across all of these is evidence. Regulators do not accept policy documents as proof of governance. They require evidence that governance is being enforced at runtime — immutable logs, policy enforcement records, and audit-ready evidence packages.

The Emissions Blind Spot

Every AI token consumed is a Scope 3 Category 1 emission under the GHG Protocol. Every model inference requires compute, and that compute has a carbon footprint. Under AASB S2, organisations with material AI usage are required to measure and disclose these emissions as part of their climate-related financial disclosures.

Yet most organisations cannot measure their AI emissions because they cannot measure their AI usage. The visibility gap is also a sustainability gap. Directors face personal liability under Corporations Act 2001 s299(1)(f) for failing to disclose material climate-related risks — and for organisations with significant AI estates, AI emissions are increasingly material.

The Progressive Governance Framework

The Progressive Governance Framework is built on a simple principle: **governance should match maturity**. An organisation that cannot yet see its AI estate should not be forced to implement full enforcement. Equally, an organisation that has complete visibility should not remain passive when it has the information needed to enforce policy.

The framework defines three stages, each building on the previous one:

Stage	Mode	Question It Answers	Time to Deploy	Governance Posture
1	Observe	"What AI do we have and what is it costing us?"	Hours	Passive — visibility without disruption
2	Enforce	"How do we control AI behaviour at runtime?"	Days	Active — policy enforcement before response
3	Orchestrate	"How do we optimise AI across the entire estate?"	Weeks	Strategic — multi-model optimisation

Stage 1: Observe Mode — Immediate Visibility

Observe Mode is the entry point. It deploys in hours with zero infrastructure changes — no proxy, no code modifications, no disruption to production workloads. It sits alongside existing AI workflows and collects complete telemetry passively.

What the organisation gains:

- Complete AI inventory across all models, providers, and endpoints
- Token-level cost attribution by model, workflow, team, and business owner
- Shadow AI discovery — automated scanning of 30+ known AI vendors with risk scoring
- Sustainability baseline — CO₂e measurement for every AI interaction
- Immutable audit trail — cryptographically signed records of every AI request
- Usage telemetry — latency, token counts, request frequency, error rates

Regulatory obligations satisfied:

- ASIC REP 798 — AI inventory and risk register
- AASB S2 — emissions baseline measurement

- APS AI Policy 2024 — AI asset register and usage reporting
- ISO 42001 — AI management system evidence

Scenario: A mid-tier Australian bank deploys Observe Mode on Monday morning. By Friday, the CFO has a complete AI cost attribution report for the first time — revealing \$2.4M in annual AI spend that was previously invisible in cloud infrastructure invoices. The CISO has identified 34 unsanctioned AI tools in use across the organisation. The Chief Sustainability Officer has a baseline CO₂e measurement for AI compute emissions.

Stage 2: Enforce Mode — Runtime Policy Enforcement

Enforce Mode adds an inline enforcement layer that intercepts every AI request before it reaches the model. This is where governance moves from observation to action. Policy is enforced at the point of execution — not retrospectively reviewed after the fact.

What the organisation gains (in addition to Observe):

- Inline PII redaction — sensitive data is detected and removed before it leaves the organisation
- Model access controls — block, allow, or route by user, team, or policy
- Intelligent model routing — route requests to the optimal model based on cost, latency, and data residency requirements
- Budget enforcement — hard caps, soft alerts, and automated throttling
- Runtime guardrails — prompt filtering, content controls, and topic blocklists enforced in under 5ms
- Policy-as-code — governance rules that execute automatically at the infrastructure layer

Regulatory obligations satisfied (in addition to Observe):

- Privacy Act 1988 — evidence that PII is not sent to external AI models
- ASIC REP 798 — evidence of runtime governance enforcement
- APS AI Policy 2024 — demonstrable policy enforcement for government suppliers
- ISO 42001 — operational controls with enforcement evidence

Scenario: The same bank, now four weeks into deployment, activates Enforce Mode. PII redaction is applied to all AI requests — ensuring that customer data never reaches external model providers. The Model Routing Engine reduces AI spend by 40% by routing routine summarisation tasks to cost-optimised models while preserving frontier models for complex analysis. Budget alerts fire automatically when any department exceeds its monthly AI allocation.

Stage 3: Orchestrate Mode — Full Sovereign Governance

Orchestrate Mode is the complete platform capability. It adds semantic intelligence, multi-model orchestration, and sovereign deployment options. This is where the platform actively optimises AI operations — reducing cost and emissions simultaneously while governing autonomous agents and enabling air-gapped deployment for classified workloads.

What the organisation gains (in addition to Enforce):

- Semantic caching — deduplicate semantically similar requests to reduce token consumption by 30–60%
- Context window management — intelligent truncation to reduce input tokens for long-running workflows
- Multi-model orchestration — chain models together with governance at every handoff
- Simulation engine — model cost and performance scenarios before production deployment
- Agentic workflow governance — boundaries, approval gates, and escalation paths for autonomous agents
- Sovereign air-gapped deployment — operate entirely within Australian data borders

Regulatory obligations satisfied (in addition to Enforce):

- AASB S2 — full automated disclosure PDF with year-on-year comparison
- Corporations Act s299(1)(f) — director liability coverage through complete AI risk disclosure
- ISO 42001 — complete AI management system with continuous improvement evidence
- Sovereignty requirements — air-gapped deployment for classified workloads

Scenario: At 90 days, the bank activates Orchestrate Mode for its highest-volume agentic workflows. Semantic caching eliminates 45% of redundant model calls. The Simulation Engine models the impact of switching the document analysis pipeline from GPT-4o to Claude 3.5 Haiku — projecting a 62% cost reduction with only a 3% quality impact. The board receives its first automated AASB S2 AI emissions disclosure report.

Module Architecture Deep-Dive

Songlines Control® contains nine purpose-built modules. Each module addresses a distinct dimension of enterprise AI governance. Critically, every module becomes more capable as the organisation progresses through the modes — there is no wasted investment.

Module	Observe Mode	Enforce Mode	Orchestrate Mode
Economic Control	Cost visibility and attribution	Budget enforcement and throttling	Predictive cost modelling
Workflow Economics	Historical cost analysis	—	Pre-production simulation
Model Routing Engine	Routing recommendations	Active routing enforcement	Multi-model orchestration
Observability	Full request/response logging	Policy outcome logging	Agentic workflow tracing
Audit & Compliance	Immutable audit records	Enforcement evidence	Agent decision audit trail
Shadow AI Governance	Discovery and risk scoring	Block/allow enforcement	Automated remediation
Sustainability & AASB S2	CO ₂ e measurement	Emissions reduction via routing	Full AASB S2 disclosure automation
Runtime Guardrail Engine	—	PII redaction, prompt filtering	Adaptive guardrails for agents
Governance & Administration	User roles and departments	Policy-as-code enforcement	Approval workflows, agent boundaries

Economic Control

Economic Control is the financial visibility and enforcement layer. In Observe Mode, it provides real-time month-to-date spend dashboards with token consumption tracking across all vendors — OpenAI, Anthropic, Azure, AWS Bedrock, and Google. Pricing tables for 19+ models are built in, and cost is calculated automatically from token counts on every ingest. In Enforce Mode, it adds hard budget limits that fire before overspend occurs. In Orchestrate Mode, it enables predictive cost modelling and active cost reduction through intelligent routing and semantic caching — delivering 30–60% cost savings.

Model Routing Engine

The Model Routing Engine is the intelligent switchboard that determines which AI model handles each request. Not every task requires a frontier model. Routine summarisation, data extraction, and classification tasks can be handled by cost-optimised models at a fraction of the price — with no degradation in output quality. The routing engine applies four rule types: cost-based routing (route to cheapest capable model), latency-based routing (route time-sensitive requests to faster models), sovereignty routing (ensure regulated data never leaves Australian data residency), and budget threshold routing (automatically fall back to cheaper models when monthly budgets are reached).

Runtime Guardrail Engine

The Runtime Guardrail Engine is the inline enforcement layer that applies policy at the point of execution — in under 5 milliseconds. It handles PII redaction (detecting and removing sensitive data before it leaves the organisation), model restrictions (blocking access to unapproved models), token limits (preventing runaway consumption), topic blocklists (preventing off-policy usage), and cost caps (hard limits that prevent budget overruns). A guardrail sandbox allows organisations to test new policies against production traffic before activating them — eliminating the risk of false positives disrupting operations.

Audit & Compliance

Every AI request processed through Songlines Control® is logged with a cryptographically signed, immutable record. The record includes the model used, tokens consumed, cost, latency, routing decision, policy outcome, and — in Enforce Mode — the specific guardrails that were applied. The audit trail is full-text searchable, date-range filterable, and exportable as CSV for regulatory submissions. It is designed to satisfy Privacy Act 1988 evidence requirements, APS AI Policy 2024 reporting obligations, ASIC REP 798 governance evidence, and ISO/IEC 42001 management system documentation.

The Sustainability Dimension

Why AI Emissions Are a Board-Level Obligation

Under AASB S2 (effective for reporting periods beginning on or after 1 January 2025 for Group 1 entities), organisations must disclose climate-related financial risks and opportunities — including Scope 3 emissions. AI compute falls squarely within Scope 3 Category 1 (purchased

goods and services) under the GHG Protocol. Every model inference requires compute, and that compute has a measurable carbon footprint.

For organisations with material AI usage, failing to measure and disclose AI emissions creates director liability under Corporations Act 2001 s299(1)(f). The legal position is clear: if AI compute emissions are material to the organisation's climate risk profile, directors have a duty to ensure they are measured and disclosed.

The Dual Mandate: Cost and Emissions

The sustainability dimension of AI governance is uniquely aligned with the economic dimension. Every token saved is simultaneously a cost reduction and an emissions reduction. Intelligent model routing does not just reduce spend — it reduces the carbon footprint of every AI interaction. Semantic caching does not just eliminate redundant model calls — it eliminates the emissions associated with those calls.

This creates a powerful alignment for the CFO and the Chief Sustainability Officer. The same governance platform that delivers 40% cost savings also delivers measurable emissions reductions — from the same data, using the same infrastructure.

How Songlines Control® Delivers AASB S2 Compliance

The Sustainability & AASB S2 module provides:

- **Token-level CO₂e measurement** — every AI interaction is attributed a carbon footprint based on the model used, the compute infrastructure, and the token count
- **Scope 3 Category 1 attribution** — emissions are classified according to GHG Protocol methodology
- **Supplier emissions ledger** — emissions attributed by AI vendor (OpenAI, Anthropic, Google, Azure, AWS)
- **One-click AASB S2 disclosure PDF** — automated generation of the disclosure report with year-on-year comparison
- **Director liability assessment** — materiality threshold analysis for AI emissions
- **Executive board pack** — quarterly sustainability reporting with trend analysis and reduction targets

Shadow AI: The Governance Gap Most Platforms Miss

The Scale of the Problem

Shadow AI is the fastest-growing governance gap in Australian enterprise. Unlike shadow IT of the previous decade — which required procurement, installation, and configuration — shadow AI requires nothing more than a browser and a credit card. Employees adopt AI tools because they are productive, and they do so without IT approval, security review, or data classification assessment.

The risks are significant: sensitive corporate data sent to unvetted AI models, compliance obligations violated without awareness, uncontrolled spend accumulating across dozens of departmental subscriptions, and no audit trail of what data has been processed by which models.

Discovery, Declaration, and Governance

The Shadow AI Governance module takes a three-stage approach:

- 1. Automated Discovery (Observe Mode)** — The platform automatically scans for 30+ known AI vendors across the organisation's network traffic, SSO logs, and expense data. Each discovered tool is assigned a risk score from 0 to 100 across four severity levels (Critical, High, Medium, Low) based on data handling practices, compliance certifications, and data residency.
- 2. Employee Declaration Portal (Observe Mode)** — Rather than simply blocking tools, the platform provides a self-service declaration portal where employees can register AI tools they are using. This creates a path to compliance rather than driving usage underground. Declared tools are assessed, risk-scored, and either approved or flagged for review.
- 3. Enforcement (Enforce Mode)** — Once the organisation has a complete picture of its AI estate, Enforce Mode enables block/allow rules for unsanctioned tools. Approved tools are routed through the governance layer. Unapproved tools are blocked at the network level. The AI Risk Policy Builder generates a complete AI acceptable use policy aligned with Privacy Act 1988 and APS AI Policy 2024 requirements.

Implementation Roadmap: 90 Days to Full Governance

The Progressive Governance Framework is designed for rapid deployment. Organisations do not need to spend months in planning before gaining value. The following roadmap shows how a typical enterprise moves from zero governance to full platform capability in 90 days.

Week	Action	Mode	Outcome
1	Deploy Observe Mode	Observe	Complete AI inventory, cost baseline, shadow AI discovery, emissions baseline
2	Board reporting	Observe	First AI cost attribution report, shadow AI risk register, AASB S2 emissions baseline
3–4	Policy definition	Observe	Define routing rules, budget thresholds, PII policies, and model access controls based on observed data
5	Deploy Enforce Mode	Enforce	Inline proxy activated — PII redaction, model routing, budget enforcement live
6–8	Optimise and expand	Enforce	Tune routing rules, expand guardrails, activate shadow AI enforcement, measure cost savings
9–10	Deploy Orchestrate Mode	Orchestrate	Semantic caching, simulation engine, multi-model orchestration activated
11–12	Full governance	Orchestrate	Automated AASB S2 disclosure, agentic workflow governance, sovereign deployment options

Key principle: Value is delivered from Week 1. Observe Mode provides immediate visibility — the organisation does not wait 90 days for its first governance outcome. Each subsequent stage adds capability without disrupting what is already running.

Why Cetus AI

Songlines Control® is not a generic observability tool with AI features bolted on. It is a purpose-built AI governance platform designed from the ground up for Australian enterprise and government requirements.

Differentiator	What It Means
Australian-built, Australian-sovereign	Designed for Australian regulatory requirements from day one — not retrofitted from a US or EU product
Single platform, three modes	One runtime layer, one audit trail, one policy engine — progressive capability without re-architecting
Nine purpose-built modules	Each module addresses a distinct governance dimension — cost, routing, compliance, sustainability, shadow AI, guardrails
Azure Marketplace partner	MACC-eligible, IP Co-sell Ready — procurement through existing Microsoft agreements
Audit-ready evidence packages	ISO 27001 aligned, with exportable evidence for Privacy Act, ASIC REP 798, APS AI Policy, and ISO 42001
Zero-disruption deployment	Observe Mode deploys in hours with no code changes, no proxy, and no infrastructure modifications
Unified cost and sustainability	The only platform that unifies AI cost governance and AASB S2 emissions reporting from the same data layer

Conclusion

AI governance is not optional for Australian enterprise. The regulatory environment demands it. The board expects it. The economics require it. The sustainability obligations mandate it.

But governance does not need to be a multi-year transformation programme that delays AI adoption. The Progressive Governance Framework demonstrates that organisations can start with immediate visibility — deployed in hours — and expand to full sovereign governance as their maturity grows. Each stage is additive. Each stage uses the same platform. Each stage delivers measurable value from day one.

The organisations that build this governance infrastructure now will not just reduce their AI costs and satisfy their regulatory obligations. They will be the organisations that can scale AI confidently — because they can see it, control it, and prove it.

Start with visibility. Expand to control. Lead with governance.

Book a Strategic AI Session to see how Songlines Control® can deliver progressive governance for your organisation.

cetusai.com.au | hello@cetusai.com.au

Cetus AI Pty Ltd | Brookwater, Queensland, Australia