

What Does Your Governance Infrastructure Need to Look Like Before AGI Arrives?

Published by: Cetus AI
Subtitle: A Strategic Brief for Australian Enterprise and Government Leaders
Date: June 2026
Reference: CETUS-WP-003 | Version 2.0

2–5 Years	40%	\$100K	5 Pillars
Expert median AGI timeline	Agentic AI projects cancelled by 2027	Per developer/year in AI tokens	Of AGI-ready governance

Executive Summary

The question is no longer whether Artificial General Intelligence will arrive. The question is whether your organisation will be ready to govern it when it does.

Leading AI researchers and the CEOs of the world's most advanced AI laboratories now broadly agree that AGI — AI systems capable of performing any intellectual task that a human can perform — is likely to emerge within the next two to five years. Anthropic's CEO Dario Amodei has stated publicly that powerful AI systems will emerge in late 2026 or early 2027. OpenAI's Sam Altman has described AGI as arriving 'sooner than most people think.' As of early 2026, forecasters on the prediction market Metaculus assign a 25% probability of AGI arriving by 2029 and a 50% probability by 2031.

The governance infrastructure that organisations need for AGI is the same infrastructure they need right now for the agentic AI systems already operating inside their enterprises today. The window to build it is open now — and it will not remain open indefinitely.

This white paper argues that the organisations that build the right governance foundations in the next 12 to 18 months will be positioned to deploy AGI safely, compliantly, and competitively when it arrives. Those that do not will face the same outcome that Gartner predicts for more

than 40% of agentic AI projects by 2027: cancellation, not because the technology failed, but because the governance did.

Product Architecture Note: The Three Modes of Songlines Control®

To provide clarity on how the capabilities described in this paper are implemented, it is important to understand the deployment architecture of Songlines Control®. The platform operates as a single unified control plane with three distinct operating modes, reflecting an organisation's AI maturity journey:

Operating Mode	Primary Focus	Key Capabilities Described in this Paper
Observe Mode	Visibility & Telemetry	Shadow AI Detection, Model Inventory, Usage Reporting
Enforce Mode	Inline Governance & Routing	Inline PII Redaction, Policy Enforcement, Immutable Audit Trail
Orchestrate Mode	Agentic Workflow Management	State Externalisation, Context Window Management

When this paper refers to specific capabilities such as inline redaction or shadow AI detection, they are delivered through the Enforce and Observe modes of the platform.

Part One: The AGI Timeline Is Closer Than Most Boards Realise

What the Experts Are Saying

The AGI timeline debate has shifted dramatically in the past 18 months. What was once a fringe position — that AGI could arrive within a decade — is now the mainstream view among the researchers and executives closest to the technology.

Dario Amodei, CEO of Anthropic, wrote in his widely-read essay *Machines of Loving Grace* that he expects powerful AI systems to emerge in late 2026 or early 2027. Sam Altman of OpenAI

has described the path to AGI as 'closer than most people think.' The AI 2027 research project, which extrapolates from compute scale-ups, algorithmic improvements, and benchmark performance, projects AGI-level capabilities arriving within the next 12 to 24 months for specific domains.

Source	Prediction	Probability / Confidence
Dario Amodei, Anthropic CEO	Powerful AI in late 2026 or early 2027	High confidence (public statement)
Sam Altman, OpenAI CEO	'Sooner than most people think'	High confidence (public statement)
Metaculus Forecasters (Feb 2026)	AGI by 2029	25% probability
Metaculus Forecasters (Feb 2026)	AGI by 2031	50% probability
80,000 Hours Expert Survey (2025)	AGI before 2030	Within range of expert opinion
AI 2027 Research Project	AGI-level in specific domains	12–24 months

The Bridge Between Now and AGI: Agentic AI

Boards and executive teams do not need to wait for AGI to arrive to face AGI-level governance challenges. They are already here, in the form of agentic AI.

Agentic AI systems — AI that takes multi-step autonomous actions, calls external tools, manages workflows, and makes decisions without human approval at each step — are already operating inside enterprise environments at scale. Gartner projects that 40% of enterprise applications will embed task-specific AI agents by the end of 2026, up from less than 5% in 2025. That is an eight-fold increase in a single year.

The organisations that build the governance infrastructure to manage agentic AI today are building the infrastructure they will need for AGI tomorrow. The technical challenges are the same — only the stakes are higher.

These systems are not passive tools that respond to queries. They are autonomous actors that can trigger workflows, move data, interact with external systems, execute financial transactions, and influence outcomes independently. The governance challenges they create — accountability, traceability, cost control, regulatory compliance — are the same challenges that AGI will create, at a larger scale and with higher stakes.

Part Two: Why Your Current Governance Infrastructure Is Not Enough

The Six Governance Gaps That Will Break Under AGI

Research from Gartner, Harvard Business Review, and the American Bankers Association points to a consistent pattern: enterprises are deploying AI agents faster than they can control, explain, or audit them. The following six governance gaps are the most critical — and the most dangerous as AI systems become more capable.

Gap 1: No Centralised AI Inventory. Most organisations do not have a complete, current record of the AI systems operating inside their enterprise. AI tools are adopted at the team level, embedded in SaaS products, and deployed by individual developers without central visibility. Under AGI, this gap becomes catastrophic. An AGI-level system operating without visibility is an ungoverned actor with potentially unlimited capability.

Gap 2: Governance Introduced After Deployment. In most enterprises today, AI systems are built first and reviewed later. Legal, risk, and compliance teams are brought in only when pilots are nearly complete. Gartner's prediction that more than 40% of agentic AI projects will be cancelled by 2027 is largely attributable to this pattern. Under AGI, post-deployment governance is not just inefficient — it is dangerous.

Gap 3: No Decision Traceability. When an AI system makes a decision, most organisations cannot reconstruct the reasoning behind it. There is no record of what data was used, which model was involved, what prompt was submitted, or what alternatives were considered. Under AGI, decision traceability is a legal and regulatory necessity. ASIC's REP 798 identified this gap explicitly. APRA's April 2026 letter called for board-level accountability for AI decisions.

Gap 4: Policy on Paper, Not in Code. Most organisations document their AI governance policies in Word documents and compliance manuals. These documents describe what AI systems should do. They do not actually control what AI systems do. Under AGI, this gap is existential. An AGI-level system that can reason about and potentially circumvent documented-but-unenforced policies is a governance failure waiting to happen.

Gap 5: Undefined Human-in-the-Loop Thresholds. Most organisations have not defined — in technical, enforceable terms — which AI decisions require human approval before execution. Under AGI, the human-in-the-loop threshold is one of the most important governance decisions an organisation will make. Regulators in Australia, the EU, and the United States are already moving toward mandatory human oversight requirements for high-risk AI decisions.

Gap 6: No Economic Governance Layer. Agentic AI systems can run continuously, retry tasks, call multiple tools, and scale usage without warning. One large enterprise recently disclosed that they are budgeting \$20,000 per developer per year in AI tokens today, with projections of \$100,000 per developer per year within two years. Under AGI, the economic governance challenge is orders of magnitude larger.

Part Three: The Australian Regulatory Context

ASIC and APRA Are Already Moving

Australian regulators are not waiting for AGI to arrive before acting. The regulatory signals from ASIC and APRA in 2026 make clear that AI governance is an immediate enforcement priority, not a future consideration.

Regulator	Document	Date	Key Requirement
ASIC	REP 798: Beware the Gap	Oct 2024	AI inventory, consumer harm assessment, governance frameworks
ASIC	Open Letter to AFS Licensees	May 2026	Board-level AI risk discussion; urgency, focus and accountability
APRA	Letter to Industry on AI	Apr 2026	Step-change in AI risk management; information security; supplier risk
EU	EU AI Act (2024/1689)	Aug 2024	Risk-based framework; human oversight; audit trails; technical documentation
ISO	ISO 42001:2023	Dec 2023	AI management system standard; benchmark for governance maturity

ASIC published Report 798 (Beware the Gap) in October 2024, identifying critical governance gaps in how Australian financial services firms are managing AI risk. In May 2026, ASIC

published an open letter to all AFS licensees calling for 'urgency, focus and accountability' in AI governance, explicitly requiring board-level discussion of AI risk.

APRA published its own letter to all regulated entities on 30 April 2026, calling for a 'step-change' in AI risk management and governance. Both regulators have signalled that enforcement action will follow for entities that do not demonstrate material improvement.

For APRA-regulated entities — banks, insurers, and superannuation funds — AI governance is now a prudential risk management obligation, not a discretionary best practice. The enforcement clock is already running.

Part Four: What AGI-Ready Governance Infrastructure Looks Like

The Five Pillars of AGI-Ready Governance

Based on the regulatory requirements from ASIC, APRA, and the EU AI Act, the governance gaps identified in enterprise AI deployments, and the specific challenges that AGI will create, we identify five pillars of AGI-ready governance infrastructure.

Pillar 1: Real-Time AI Inventory and Observability. Every AI system operating inside the enterprise — every model, every agent, every API call — must be registered, monitored, and attributable to a business owner. This is not a spreadsheet exercise; it requires a technical control layer that intercepts and logs AI interactions in real time. The AI inventory is the foundation of every other governance control. You cannot enforce policy on a system you cannot see.

Pillar 2: Policy-as-Code Enforcement. Governance policies must be technically enforced at the point of execution, not documented in manuals and hoped to be followed. This means that every AI request — every prompt, every model call, every agent action — is evaluated against a set of machine-readable policies before it is executed. Policy-as-code enforcement enables organisations to prevent PII from being sent to unauthorised models, restrict model use to approved cases, require human approval for high-risk decisions, and enforce budget caps in real time.

Pillar 3: Immutable Decision Audit Trail. Every AI decision — every prompt, every response, every model used, every policy applied — must be recorded in an immutable, timestamped audit trail that can be produced on demand for regulatory review, legal proceedings, or internal investigation. The audit trail must be comprehensive enough to reconstruct the full context of

any AI decision. For AGI-level systems making consequential decisions autonomously, this audit trail is the primary mechanism of accountability.

Pillar 4: Human-in-the-Loop Architecture. Organisations must define, in technical and enforceable terms, the risk thresholds above which AI decisions require human review and approval before execution. These thresholds must be embedded in the governance infrastructure, not left to individual discretion. The human-in-the-loop architecture must be proportionate to risk: low-risk decisions can proceed autonomously; high-risk decisions affecting customers, finances, or compliance must be escalated for human review.

Pillar 5: Economic Governance and Cost Attribution. Every AI interaction must be attributed to a cost centre, project, team, or individual, with budget caps that enforce automatically in real time. As AI systems become more autonomous and their usage more difficult to predict, economic governance becomes a risk management function, not just a finance function. An AGI-level system with access to enterprise systems and no cost controls is a financial risk as well as an operational one.

Part Five: The AGI Governance Readiness Assessment

The following table provides a practical self-assessment framework for boards and executive teams to evaluate their organisation's AGI governance readiness. Each pillar is assessed across three maturity levels.

Pillar	Level 1: Unprepared	Level 2: Developing	Level 3: AGI-Ready
AI Inventory	No central record; self-reported at best	Partial inventory; some systems registered	Real-time, technically enforced inventory of all AI interactions
Policy Enforcement	Policies documented but not enforced	Some technical controls; gaps exist	Policy-as-code enforced at infrastructure layer for all AI interactions
Audit Trail	No systematic record of AI decisions	Partial logging; not immutable or complete	Immutable, timestamped, complete audit trail for every AI interaction
Human-in-the-Loop	No defined thresholds; ad hoc	Some thresholds defined but not enforced	Risk-based thresholds technically enforced; escalation paths defined
Economic Governance	No cost attribution; monthly invoice surprise	Some attribution; no real-time enforcement	Real-time attribution and budget caps enforced at execution

Most Australian enterprises are currently at Level 1 or Level 2 across all five pillars. The regulatory expectation from ASIC and APRA is Level 2 minimum, with a clear trajectory toward Level 3. AGI will require Level 3 across all pillars.

Part Six: How Songlines Control® Delivers AGI-Ready Governance

Songlines Control® is the AI governance platform built specifically to address these five pillars. It operates as a provider-agnostic governance layer that sits between your organisation and every AI model you use — OpenAI, Anthropic, Azure OpenAI, AWS Bedrock, Google Gemini, and any other provider — enforcing governance policies in real time, building the audit trail automatically, and providing complete visibility into AI spend and usage.

Pillar	Songlines Capability	Business Impact
AI Inventory	Real-time executive dashboard showing every model, team, and request across the organisation (Observe Mode)	Board directors can answer 'What AI are we running?' in seconds. Shadow AI is visible and governable from first use.
Policy-as-Code	Pre-built policy templates for PII, sovereignty, model restriction, approval gates, and cost caps — enforced at infrastructure layer (Enforce Mode)	Policies cannot be bypassed by individual developers or applications. Compliance is technical, not aspirational.
Audit Trail	Immutable, timestamped log of every prompt, model, policy, and response — exportable for regulatory review (Enforce Mode)	Directly addresses ASIC REP 798 and APRA April 2026 requirements. Evidence base for board accountability.
Human-in-the-Loop	Configurable approval gates that hold requests pending human review when risk thresholds are triggered (Enforce Mode)	Implements APRA and ASIC human oversight requirements without requiring changes to developer workflows.
Economic Governance	Real-time cost attribution by team, project, and user; budget caps that enforce automatically before limits are exceeded (Orchestrate Mode)	Eliminates end-of-month invoice surprises. Orchestrate Mode adds intelligent routing that reduces token spend by 30–60%.

Part Seven: The Window Is Open Now

Why the Next 12 Months Matter

The governance infrastructure that organisations need for AGI is not something that can be built in a week when AGI arrives. It requires time to deploy, configure, and embed into organisational processes. It requires board-level commitment, cross-functional ownership, and the development of governance muscle memory across the organisation.

The organisations that will be trusted to deploy AGI are the ones that have demonstrated, over time, that they can govern AI responsibly. That track record is built now, with the AI systems that are already running inside the enterprise.

There is also a regulatory imperative. ASIC and APRA have both signalled that enforcement action will follow for entities that do not demonstrate material improvement in AI governance. The organisations that act now will be ahead of the regulatory curve. Those that wait will be responding to enforcement notices.

A Practical 90-Day Action Plan

Timeframe	Action	Owner
Days 1–30	Deploy a real-time AI inventory: instrument your environment to capture all AI interactions across all providers	CTO / CISO
Days 1–30	Commission a board-level AI governance briefing aligned to ASIC and APRA requirements	CEO / Board
Days 30–60	Implement policy-as-code for PII detection, model restriction, and cost attribution	CTO / CISO
Days 30–60	Define human-in-the-loop thresholds for high-risk AI decisions	CRO / Legal
Days 60–90	Produce first board-ready AI governance report with full AI inventory and policy compliance status	CRO / CTO
Days 60–90	Initiate ISO 42001 gap assessment and roadmap to certification	CRO / Compliance

Getting Started with Songlines Control®

Cetus AI offers a 30-day free trial of Songlines Control® through the Microsoft Azure Marketplace. For APRA-regulated entities, we offer a two-week proof-of-concept engagement that produces a complete AI inventory, a policy framework aligned to APRA's April 2026 expectations, and a board-ready governance report.

Product	Description	Price (AUD)	Availability
Songlines Control®	AI governance, cost intelligence, policy enforcement (Observe + Enforce Modes), and immutable audit trail	\$9,995/month	Azure Marketplace Direct
Songlines Control® — Orchestrate Mode	Everything in Control, plus intelligent model routing to reduce token spend by 30–60%	\$19,995/month	Azure Marketplace Direct

Both plans include a 30-day free trial, are MACC eligible (Azure Consumption Commitment), and are audit-ready for Australian government and regulated enterprise with Australian data residency.

Contact Cetus AI:
hello@cetusai.com.au | cetusai.com.au | songlinesai.com
Azure Marketplace: search 'Songlines Control'

References

1. Dario Amodei, *Machines of Loving Grace*, Anthropic, October 2024
2. 80,000 Hours, *Shortening AGI Timelines: A Review of Expert Forecasts*, March 2025 (updated February 2026)
3. AI 2027 Research Project, ai-2027.com
4. Gartner, *Predicts 2026: More Than 40% of Agentic AI Projects Will Be Cancelled by 2027*, 2026
5. ASIC, *Report 798: Beware the Gap — Governance Arrangements in the Face of AI Innovation*, October 2024
6. ASIC, *Open Letter to AFS Licensees and Market Participants*, May 2026
7. APRA, *Letter to Industry on Artificial Intelligence (AI)*, April 30, 2026
8. EU AI Act (Regulation (EU) 2024/1689), August 2024
9. ISO 42001:2023, Information Technology — Artificial Intelligence — Management System
10. Accelirate, *The 2026 Agentic AI Governance Crisis*, January 2026

11. Deloitte, *The Agentic Reality Check: Preparing for a Silicon-Based Workforce*, December 2025
12. World Economic Forum, *Non-Human Identities: Agentic AI's New Frontier of Cybersecurity*, October 2025

© 2026 Cetus AI Pty Ltd. All rights reserved. Songlines Control® is a registered trademark of Cetus AI Pty Ltd. This white paper is provided for informational purposes only and does not constitute legal advice.

Cetus AI · cetusai.com.au · hello@cetusai.com.au · songlinesai.com · Version 2.1 — June 2026