Why Cetus AI?

Operational AI governance for complex enterprise environments

White Paper | September 2026 | Version 2.0 · Release 2026.09.4

Cetus AI Pty Ltd | Brookwater, Queensland

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

CUSTOMER RESOURCE

Contents

Executive summary 2
1. AI governance must move from policy to operation 2
2. The Cetus AI platform 3
3. An additive architecture for mature enterprises 3
4. Evidence Fabric: from distributed logs to a governance record 4
5. One platform, connected components and modules 5
6. Deployment should follow the workload's trust boundary 5
7. The value depends on the stakeholder 6
Recommended next steps 6
References 6

Executive summary

Artificial intelligence is already distributed across the enterprise. It is present in cloud platforms, business applications, copilots, software-development tools, internally developed models, automated workflows, and emerging agent systems. Most organisations have established security, identity, integration, monitoring, and risk platforms surrounding these environments. Yet when an AI-assisted action must be explained, the evidence is often fragmented across several systems.

This is the operational AI governance gap. Cloud logs can show execution. Identity platforms can show authentication. API gateways can show routing. Data loss prevention tools can show data movement. Governance, risk, and compliance systems can record policy and risk acceptance. None of those systems, operating alone, necessarily creates one AI-specific record that proves who had authority, which policy applied, which model acted, what data was involved, what conditions were enforced, and where human oversight occurred.

Cetus AI was established to close that gap without asking mature organisations to discard the platforms they already trust. The Cetus AI platform adds AI-specific policy decisions, read-only telemetry ingestion, cross-stack evidence correlation, economic controls, agent governance, and review-ready outputs to the existing enterprise environment. The platform is built and available for controlled enterprise deployment. Public Release 2026.09.4 establishes the current platform position and published capability boundaries used throughout this paper. Customer-specific integrations and deployment patterns still require validation against the customer's architecture, sources, identity, security, performance, failure behaviour, and operating model.

The principle is straightforward: work with the enterprise stack, do not replace it. Cetus AI enables organisations to begin with visibility and evidence, introduce runtime control where justified, and scale toward broader orchestration without forcing every workload through the same architecture. This paper explains the customer problem, the additive architecture, the Evidence Fabric, staged deployment patterns, and why organisations should start with the controls they already trust.

1. AI governance must move from policy to operation

Many organisations have already created responsible AI principles, acceptable-use policies, model-risk processes, and executive oversight forums. These are necessary governance foundations. They are not, by themselves, operational control.

Operational governance answers a more immediate set of questions. Was the person, team, or agent authorised to perform the action? Was the selected model approved for the workload and data classification? Did a policy require redaction, model substitution, a spending limit, or human approval? Can the organisation reconstruct the decision without manually joining records from multiple platforms?

The challenge becomes more significant as AI moves beyond isolated chat interactions. Modern enterprise AI may involve an employee identity, an orchestration layer, several APIs, a model endpoint, an agent toolchain, a business application, and a human approver. Each component can generate valid telemetry while the complete governance story remains unresolved. Regulatory scrutiny is increasing; for example, the Australian Securities and Investments Commission (ASIC) found in its Report 798 that AI adoption is outpacing governance, warning of a widening "governance gap" that could lead to consumer harm and data privacy failures [1]. Similarly, the Australian Prudential Regulation Authority (APRA) has observed that governance, risk management, assurance, and operational resilience practices are failing to keep pace with the scale and complexity of AI adoption across regulated entities [2].

Existing enterprise capability	What it does well	The unresolved AI-governance question
Identity and access management	Authenticates people and services, assigns roles, and manages entitlements.	Was this identity authorised for this AI use case, model, tool, and data classification?
API gateway or orchestration	Routes traffic, applies general API policy, and supports service integration.	Which AI-specific policy decision applied, and what evidence links it to the resulting model action?
Cloud and application monitoring	Records infrastructure, application, and model-service events.	How do those events connect to business context, policy, human oversight, and accountability?
DLP and security controls	Detects sensitive data movement and enforces security rules.	What governance condition was triggered, and how was the resulting AI decision recorded?
GRC and model-risk systems	Maintains policies, risks, controls, registers, and attestations.	Can documented policy influence runtime behaviour and be evidenced at the transaction level?

The answer is not to replace these systems. It is to connect them through an AI-specific governance and evidence layer.

2. The Cetus AI platform

Cetus AI is an Australian enterprise AI governance company. The company's focus is the practical infrastructure required to connect policy intent with technical execution and reviewable evidence.

The Cetus AI platform provides one governance evidence layer across an organisation's existing AI estate. It can observe approved telemetry, evaluate policy, govern identities and agents, attribute cost and sustainability impact, record incidents, and produce evidence for risk, legal, board, audit, and operational review. Customer-specific integrations and deployment patterns still require validation against the customer's architecture, sources, identity, security, performance, failure behaviour, and operating model.

The platform supports three levels of governance maturity.

Level	Purpose	Representative capabilities
Observe	Establish visibility and produce evidence without changing runtime decisions.	Discover Shadow AI, ingest approved telemetry read-only, monitor workflows, attribute model spend, and estimate AI-related emissions.
Enforce	Apply AI-specific policy to selected workloads.	Return approve, deny, or conditional decisions; require PII redaction, approved models, budget controls, or human review.
Orchestrate	Coordinate governed services and identities across a broader AI estate.	Govern model routing, agent permissions, simulation, shared control services, and customer-defined deployment patterns.

This model avoids a false choice between doing nothing and introducing a new central control plane across the entire organisation. A customer can begin with evidence visibility, validate the value on one real workflow, and introduce deeper controls only where the risk, architecture, and operating model justify them.

3. An additive architecture for mature enterprises

The architectural position of Cetus AI is deliberately different from a mandatory middleware proposition. Source platforms retain the responsibilities they already perform well. MuleSoft, as an implementation pattern, can

continue to orchestrate. Entra ID can continue to manage identity. Azure and AWS can continue to execute and record workloads. SAP and Salesforce can retain business context. SIEM and DLP platforms can continue to provide enterprise security controls.
Cetus AI adds the missing AI-specific functions across those boundaries.

Architecture layer	Representative systems and functions
Existing enterprise stack	MuleSoft, SAP, Salesforce, Azure, AWS, internal ML, Entra ID, SIEM, and DLP.
Cetus AI governance layer	Policy Decision Point, read-only log ingestion, Evidence Fabric correlation, and lightweight SDKs.
Governance outputs	Automated-decision register, individual decision records, policy evidence, tamper-evident audit chain, cost attribution, and sustainability evidence.

The Policy Decision Point is designed to receive selected metadata describing the user or agent, model, use case, workflow, and data classification. It can return a deterministic decision—approve, deny, or approve with conditions—while the source platform retains the payload and request path. This enables an organisation to add AI-specific policy without making Cetus AI the owner of every underlying transaction.
Conditions can include actions such as PII redaction, model substitution, a human-review requirement, a spending threshold, or an approved endpoint. The existing orchestration or application layer enforces the returned decision. The resulting event is then available to the evidence layer for correlation with identity, execution, and business context.

4. Evidence Fabric: from distributed logs to a governance record

Enterprise systems generate large volumes of valid telemetry. The challenge is that a regulator, auditor, board committee, customer, or incident team rarely needs a collection of unrelated log files. They need a coherent record of the decision.
Cetus AI's Evidence Fabric correlates approved source records and makes gaps visible across the customer environment using shared identifiers, timestamps, model identity, policy references, data classification, and human-review context. Correlation does not establish causation, but it connects the events to support review. Potential sources include Entra ID, MuleSoft, Azure Monitor, AWS CloudTrail, SAP BTP, Salesforce Event Monitoring, internal AI systems, and Policy Decision Point decisions.
The result is a set of structured outputs.

Evidence output	Customer purpose
Automated-decision register	Maintains an inventory of systems that make or materially assist decisions affecting people, customers, employees, or other stakeholders.
Individual decision record	Links identity, model, data category, policy, conditions, execution, and human oversight for a defined transaction or workflow.
Policy-enforcement log	Records approvals, denials, conditions, overrides, and review events against the relevant enterprise policy.
Tamper-evident audit chain	Cryptographically links governance events to support integrity and later verification.
Economic and sustainability evidence	Attributes model consumption, spend, and estimated emissions to the relevant team, workflow, or business activity.

This evidence supports compliance, risk management, and assurance activities. It does not replace legal interpretation, governance ownership, risk acceptance, or human decision-making. Capabilities can support a customer process within configured coverage; they do not make the customer compliant, prove causation, replace an auditor,

or guarantee completeness. The practical value is operational: establishing a continuous evidence-production process as part of normal AI operations.

5. One platform, connected components and modules

The Cetus AI platform is the core enterprise AI governance platform, structured around integrated components that ensure each capability contributes to a common policy, identity, evidence, and audit foundation.

Teams is the enterprise visibility and governance-intake component. It supports shadow AI discovery, user declarations, inventory, and risk scoring, helping organisations understand which AI tools are being used across the business.

Control is the policy, identity, runtime guardrail, exception, and human-oversight component. The Policy Decision Point translates enterprise rules into operational decisions, supporting model restrictions, data-handling requirements, allowlists, PII controls, spending conditions, and human-review gates. As AI systems become more autonomous, identity is no longer limited to human users; APRA has explicitly highlighted the challenge of nonhuman actors such as AI agents [2]. Control governs these agents by defining permission boundaries, managing credentials, controlling sessions, and recording actions.

Evidence Fabric is the shared evidence capability that connects approved records across the customer environment. It correlates distributed logs into coherent governance records, compliance reports, and tamper-evident audit records.

Value Control is an integrated module within the Cetus AI platform. It provides visibility into AI economics and sustainability. Value Control distinguishes forecast, observed, attributed, finance-validated, and risk-adjusted value. It attributes model consumption and spend to specific business owners and workflows, and provides estimated AI-related emissions to support sustainability reporting requirements.

6. Deployment should follow the workload's trust boundary

Enterprise AI estates contain workloads with different security, privacy, resilience, and operational requirements. A public-information assistant, an internal productivity tool, a customer decision workflow, and an industrial agent should not automatically use the same deployment model.

Cetus AI supports several patterns so that control and evidence can be matched to the workload.

Deployment pattern	How it works	Appropriate starting point
Read-only evidence	Ingests approved telemetry and produces correlated evidence without intervening in runtime decisions.	Organisations beginning with visibility, baselining, or disclosure readiness.
Policy Decision Point sidecar	Evaluates metadata inside or close to the customer boundary; the existing orchestration layer enforces the decision.	Mature enterprises requiring runtime policy without a new mandatory transit layer.
Dedicated instance	Runs the complete platform within a customer tenant or private-cloud boundary.	Regulated workloads, customer isolation, and organisation-specific operating requirements.
Managed or customer-hosted control plane	Centralises selected AI services or runs within customer infrastructure where broader orchestration is appropriate.	Greenfield AI services, consolidated platforms, and environments with defined sovereignty requirements.

The deployment decision should consider which data is required, where policy evaluation occurs, whether the control is advisory or mandatory, how the workload behaves during service interruption, how evidence is retained, and who operates each component. Customer responsibilities include validating these patterns against their own architecture and operating model.

7. The value depends on the stakeholder

AI governance initiatives often fail when they are presented as a technology platform with one generic value proposition. The benefit changes according to the stakeholder's accountability.

Stakeholder	Primary concern	Cetus AI value
Board, risk, and legal	Can the organisation demonstrate oversight, authority, exceptions, and human intervention?	Automated-decision inventories, policy records, review evidence, decision history, and sustainability reporting support.
CIO and enterprise architecture	Can governance be added without disrupting existing platforms and delivery programs?	Additive integration, explicit trust boundaries, consistent policy decisions, and staged adoption.
CISO, privacy, and compliance	Can the organisation control sensitive data, unsanctioned AI, privileged agents, and incident evidence?	PII controls, least privilege, Shadow AI governance, incident traceability, and tamper-evident records.
Finance, operations, and sustainability	Can AI consumption be attributed and governed before it becomes an uncontrolled line item?	Spend attribution, workflow economics, budget guardrails, model-selection evidence, and AI emissions reporting.
Engineering and AI teams	Can governance be integrated without stopping useful experimentation?	APIs, SDKs, clear policy outcomes, simulation, and phased control by workload.

The common outcome is not "more governance". It is faster adoption with clearer authority, measurable controls, and evidence that supports review.

Recommended next steps

Organisations should begin their AI governance journey by establishing visibility over their current AI estate. The recommended approach is to identify a single, high-priority AI workflow—such as a customer-facing application or a critical internal process—and apply read-only evidence collection to understand its behaviour and telemetry. From there, organisations can introduce the Policy Decision Point to apply necessary guardrails and scale governance controls across broader enterprise architectures as required.

References

[1] Australian Securities and Investments Commission. (2024). *Report 798 Beware the gap: Governance arrangements in the face of AI innovation.* https://www.asic.gov.au/about-asic/news-centre/find-a-media-release/2024-releases/24-238mr-asic-warns-governance-gap-could-emerge-in-first-report-on-ai-adoption-by-licensees

[2] Australian Prudential Regulation Authority. (2026). *APRA Letter to Industry on Artificial Intelligence (AI).* https://www.apra.gov.au/news-and-publications/apra-letter-industry-artificial-intelligence-ai