



# Cetus AI Platform Architecture Brief

Enterprise control and evidence architecture for Release 2026.09.4

---

Architecture Brief | September 2026 | Release 2026.09.4

Cetus AI Pty Ltd | Brookwater, Queensland

[www.cetusai.com.au](http://www.cetusai.com.au) | [hello@cetusai.com.au](mailto:hello@cetusai.com.au)

CUSTOMER RESOURCE

# Contents

---

- Executive Summary ..... 2
- The Customer Problem: Fragmented Governance ..... 2
- An Additive Enterprise Architecture ..... 2
- Application to the Cetus AI Platform ..... 2
  - Teams: Enterprise Visibility and Intake ..... 2
  - Control: Policy, Identity, and Runtime Guardrails ..... 3
  - Evidence Fabric: Cross-Stack Correlation ..... 3
  - Value Control: Integrated Financial Validation ..... 3
- Limitations and Customer Responsibilities ..... 3
- Recommended Next Steps ..... 3
- References ..... 3
- Architecture Review Checklist ..... 4

# Executive Summary

Artificial intelligence has become deeply embedded in the enterprise, distributed across cloud platforms, business applications, copilots, internal models, and autonomous agents. The defining challenge for technology leaders is no longer acquiring AI capability, but rather governing its execution safely, transparently, and economically across a complex environment. The operational AI governance gap emerges when execution logs, identity platforms, API gateways, and risk systems each hold a piece of the puzzle, but no single record proves who had authority, which policy applied, and what conditions were enforced during an AI transaction.

Cetus AI addresses this gap through an additive architecture designed for mature enterprises. Rather than replacing existing orchestration or security platforms, the Cetus AI platform introduces an AI-specific governance and evidence layer. This Architecture Brief outlines the technical foundation of the Cetus AI platform (Release 2026.09.4), detailing how the core components—Teams, Control, Evidence Fabric, and Value Control—integrate with the existing enterprise stack to provide policy enforcement, cross-stack evidence correlation, and financial validation.

## The Customer Problem: Fragmented Governance

Organizations have established security, identity, and integration systems, yet these platforms were designed for deterministic software, not probabilistic AI systems. When an AI-assisted action occurs, the evidence is scattered: Identity systems show authentication but not AI-specific authorization. Orchestration platforms route traffic but do not capture the AI policy decision. Security tools detect data movement but do not record the AI context. GRC systems hold policy documentation but cannot influence runtime model behavior.

Without a unified governance architecture, enterprises face significant compliance, security, and economic risks, especially as they adopt agentic AI and confront shadow AI usage.

## An Additive Enterprise Architecture

The architectural philosophy of Cetus AI is to work with the enterprise stack, not replace it. The platform adds AI-specific policy decisions, telemetry ingestion, and evidence correlation across existing boundaries.

| Architecture Layer        | Representative Systems                                             | Cetus AI Governance Layer                                                                        |
|---------------------------|--------------------------------------------------------------------|--------------------------------------------------------------------------------------------------|
| Existing Enterprise Stack | Identity systems, SIEM, DLP                                        | Source of telemetry, identity, and business context. Enforces policy decisions.                  |
| Governance Integration    | Policy Decision Point (PDP), SDKs, API integrations                | Evaluates metadata against policy, returns deterministic decisions (approve, deny, conditional). |
| Evidence & Control        | Cetus AI Platform (Teams, Control, Evidence Fabric, Value Control) | Correlates evidence, records policy enforcement, attributes cost, maintains audit chains.        |

The **Policy Decision Point (PDP)** receives metadata about the user, model, workflow, and data classification. It evaluates this against enterprise policy and returns a decision, while the customer’s existing orchestration platform retains the payload and enforces the action. This separation of concerns allows organizations to implement AI governance without forcing every workload through a mandatory new transit layer.

## Application to the Cetus AI Platform

The Cetus AI platform is the core enterprise AI governance solution, structured around four primary components that operate together to provide end-to-end control and visibility.

### Teams: Enterprise Visibility and Intake

**Teams** provides the interface for enterprise visibility, governance intake, and user declarations. It is the component where shadow AI is discovered, risk scoring is initiated, and AI initiatives are registered. By providing a centralized view of AI usage across the organization, Teams enables risk and compliance professionals to understand the true footprint of AI within the enterprise.

## Control: Policy, Identity, and Runtime Guardrails

**Control** is the execution engine for AI governance. It encompasses the policy engine, identity management for both humans and autonomous agents, runtime guardrails, exception handling, and human-oversight workflows. Policy Engine: Translates documented enterprise rules into operational logic. Runtime Guardrails: Enforces conditions such as PII redaction, model substitution, or spending limits before an AI action is executed. Agent Governance: Manages permissions, session policies, and trust boundaries for autonomous AI agents, recognizing that identity is no longer limited to human users.

## Evidence Fabric: Cross-Stack Correlation

The **Evidence Fabric** is the shared capability that connects approved records across the customer environment. It ingests read-only telemetry from diverse sources and correlates it using shared identifiers, timestamps, and policy references.

The resulting structured outputs include: Automated-Decision Register: An inventory of AI systems making or assisting material decisions. Individual Decision Records: Detailed logs linking identity, model, policy, and execution for specific transactions. Tamper-Evident Audit Chain: Integrity-linked records designed to support later verification; this is not independent cryptographic proof of the underlying business event.

## Value Control: Integrated Financial Validation

**Value Control** is an integrated module within the Cetus AI platform, designed to bridge the gap between AI operations and enterprise finance (FinOps). It distinguishes between forecast, observed, attributed, finance-validated, and risk-adjusted value. By acquiring customer-approved cost data and correlating it with exact workflows, Value Control enables organizations to attribute model consumption and spend to specific teams or business activities, moving AI ROI from estimation to evidence-backed validation.

## Limitations and Customer Responsibilities

Cetus AI is built and available for controlled enterprise deployment. However, it operates within a shared responsibility model. Validation: Customer-specific integrations, deployment patterns, and policy configurations must be validated against the customer's unique architecture, security requirements, and operating model. Compliance Guarantee: While Cetus AI capabilities support customer processes, they do not make the customer compliant, replace legal interpretation, or guarantee completeness. Evidence Correlation: The Evidence Fabric correlates approved source records to make governance gaps visible; however, correlation does not establish causation. Integration Patterns: Integrations such as MuleSoft are documented as implementation patterns rather than certified connectors. Deployment: The platform is available for controlled deployment. Air-gapped deployment is currently on the roadmap.

## Recommended Next Steps

Enterprise architecture and technology leaders should begin with a bounded operating path:

1. **Assess visibility:** Identify current AI use, approved telemetry and material evidence gaps.
2. **Define trust boundaries:** Select read-only evidence, policy-decision or combined patterns according to workload sensitivity.
3. **Validate one workflow:** Correlate identity, policy, execution and review records for a defined use case.
4. **Add Value Control where relevant:** Connect customer-approved cost and outcome evidence before wider rollout.

## References

[1] Cetus AI Pty Ltd. (2026). *Why Cetus AI? Operational AI governance for complex enterprise environments*. White Paper. [2] Cetus AI Pty Ltd. (2026). *Platform Overview*.

# Architecture Review Checklist

| Review area      | Customer decision required                                                                        |
|------------------|---------------------------------------------------------------------------------------------------|
| Scope            | Which AI workflows, users, agents and decisions are included in the first controlled boundary?    |
| Evidence         | Which approved systems provide identity, policy, execution, review, cost and outcome records?     |
| Control pattern  | Is the initial path read-only evidence, a Policy Decision Point integration, or a combined model? |
| Responsibilities | Who owns source quality, policy approval, technical enforcement, exceptions and human review?     |
| Validation       | What must be demonstrated before the boundary can be expanded to additional workloads?            |

© 2026 Cetus AI Pty Ltd. This customer resource provides general information and does not constitute legal, accounting, audit, investment or regulatory advice. Any deployment, integration, control, evidence or reporting outcome depends on the customer’s architecture, source coverage, configuration, responsibilities and independent assessment.