Songlines Control: Enterprise Architecture Brief

Cetus AI | The Runtime Control Layer for Enterprise AI

1. Executive Summary

As artificial intelligence scales across the enterprise, organizations face a critical governance gap. AI models and autonomous agents operate largely outside traditional software controls, leading to unpredictable costs, compliance exposure, and opaque decision-making.

Songlines Control is a purpose-built runtime control layer that sits between enterprise applications and AI models. By intercepting every AI interaction at the point of execution, Songlines Control enforces governance policies, optimises token consumption, and provides complete observability—before the request ever reaches the model.

2. Core Architectural Principles

The Songlines Control architecture is built on four foundational principles designed for Australian enterprise and government environments:

Principle	Description
Zero-Trust Execution	No prompt, agent, or model interaction is trusted by default. Every request is evaluated against policy at runtime.
Sovereign-First Deployment	Designed to operate entirely within customer-controlled environments (on-premise or private cloud) to ensure strict data sovereignty.
Sub-5ms Enforcement	Policy evaluation and routing occur with less than 5 milliseconds of overhead, ensuring no degradation to user experience.
Agnostic Interoperability	Built to integrate seamlessly across the entire AI stack, supporting OpenAI, Anthropic, Azure, LangChain, AutoGen, and custom models.

3. How Songlines Control Works

The control layer operates through a high-throughput, low-latency proxy architecture that intercepts, evaluates, and routes AI traffic.

Phase 1: Interception

Every AI request originating from a user application, internal system, or autonomous agent is routed through the Songlines gateway rather than directly to the model provider.

Phase 2: Policy Evaluation (Runtime)

Before execution, the request is evaluated against centrally managed enterprise policies:

- **Data Loss Prevention (DLP):** PII, PHI, and sensitive corporate data are redacted or masked.
- **Agent Guardrails:** Autonomous agent intents are validated against permitted actions.
- **Budget Controls:** Token limits and budget caps are checked per user, team, or application.

Phase 3: Optimisation & Routing

Requests that pass policy checks are optimised to reduce token consumption (typically yielding a 30–60% cost reduction) and routed to the most cost-effective model capable of handling the specific task.

Phase 4: Governed Response

The model's response is intercepted on the return path, checked for hallucinations or policy violations, logged for compliance auditing, and delivered back to the requesting application.

4. Security & Compliance Architecture

Songlines Control is designed for organisations where trust is an operational requirement.

- **Auditability:** A cryptographically verifiable log is generated for every interaction, capturing the original prompt, applied policies, routing decisions, and final output.
- **Least-Privilege Access:** Integrations with enterprise identity providers (IdP) ensure that users and agents only have access to permitted models and data sources.
- **Human-in-the-Loop (HITL):** High-risk agent actions or sensitive prompts can trigger automated approval workflows, pausing execution until human authorization is granted.

5. Deployment Models

Cetus AI offers flexible deployment options to match organisational risk profiles:

Deployment Model	Target Environment	Key Benefit
Managed SaaS	Public Cloud	Rapid deployment and zero infrastructure overhead.
Private Cloud (VPC)	AWS, Azure, GCP	Dedicated tenant isolation within existing enterprise cloud boundaries.
On-Premise / Air-Gapped	Sovereign Data Centres	Complete control for defence, critical infrastructure, and high-security government use cases.

6. Integration Ecosystem

Songlines Control acts as the central nervous system for the enterprise AI stack, requiring no “rip and replace” of existing tools.

- **Models:** OpenAI, Anthropic Claude, Google Gemini, Azure OpenAI, Local/Open-Source models.
- **Agent Frameworks:** LangChain, AutoGen, CrewAI.
- **Enterprise Systems:** Active Directory / Entra ID, Splunk, Datadog, ServiceNow.

For technical documentation or to arrange a proof-of-concept deployment, please contact the Cetus AI Enterprise Team at contact@cetusai.com.au.