Operationalising ITIL AI Governance

How Cetus AI supports operational application of what ITIL defines

Technical Alignment Paper | September 2026 | Version 2.0

Cetus AI Pty Ltd | Brookwater, Queensland

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

CUSTOMER RESOURCE

Contents

Executive Summary

Artificial Intelligence (AI) introduces unique governance challenges, including autonomy in decision-making, lack of transparency, and fast-changing regulatory requirements. ITIL v5 defines what organisations should govern, establishing a common language and practical approach to responsible AI adoption [1]. However, traditional IT governance is often insufficient for AI, and the framework does not inherently enforce policy at the point of AI execution.

This technical alignment paper describes how configured Cetus AI capabilities can support the operationalisation of ITIL AI Governance principles. It explores the alignment across the ITIL 6C Capability Model, the four governance perspectives, and the ITIL AI Governance Improvement Model, subject to the selected deployment and connected evidence sources. Cetus AI enables organisations to transition from defining responsible AI principles to enforcing operational control across their existing enterprise environments.

The Operational AI Governance Gap

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. As highlighted by ITIL AI Governance, which emphasises the need to balance AI opportunity with accountability, trust, and risk management, knowing how to use AI is only half the story; governing it responsibly is equally critical [2]. Yet, a significant gap remains between policy definition and operational enforcement.

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.

Existing enterprise capabilities perform their specific functions well. Identity and access management systems authenticate people and services. API gateways route traffic and apply general policies. Cloud monitoring records infrastructure events. Data Loss Prevention (DLP) tools detect sensitive data movement. Governance, Risk, and Compliance (GRC) systems maintain policies and risk registers. However, 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.

The answer is not to replace these systems. It is to connect them through an AI-specific governance and evidence layer. This is the operational AI governance gap that Cetus AI addresses, enabling organisations to move from theoretical compliance to demonstrable, operational control.

Aligning with the ITIL 6C Capability Model

The ITIL 6C Capability Model provides a structured approach for classifying AI capabilities into six distinct dimensions: Creation, Curation, Clarification, Cognition, Communication, and Coordination [3]. This model is fundamental because it helps organisations view AI as a multifaceted enterprise capability rather than a monolithic technology. By breaking down AI into these specific functions, organisations can more accurately assess risk and determine the appropriate level of governance required for each use case. Treating all AI uniformly often leads to either excessive restriction that stifles innovation or insufficient control that exposes the organisation to significant risk.

Cetus AI is designed to support the governance of these diverse AI capabilities through its core components, each serving a specific role in the operational governance lifecycle:

- **Teams**: This component serves as the enterprise visibility and governance-intake layer. It allows organisations to track AI initiatives, manage approvals, and ensure that new AI projects align with strategic objectives and risk appetites before they are deployed into production environments.

- **Control**: The core enforcement engine, providing the policy, identity, runtime guardrail, exception, and human-oversight capabilities. It translates documented governance policies into actionable rules that are applied at the point of AI execution, ensuring that AI systems operate within defined boundaries.
- **Evidence Fabric**: The shared evidence capability that connects approved records across the complex customer environment. It correlates telemetry from various systems to create a unified, reviewable record of AI decisions and actions, providing the transparency required for effective governance.
- **Value Control**: An integrated module within the platform that focuses on the financial and operational impact of AI. It distinguishes forecast, observed, attributed, finance-validated, and risk-adjusted value, ensuring that AI investments deliver measurable returns while managing associated costs.

The platform's additive architecture is a key differentiator. It allows existing source platforms—such as API gateways, identity providers, and cloud infrastructure—to retain their responsibilities. Cetus AI does not seek to replace these systems but rather adds the missing AI-specific functions, such as policy decision APIs and read-only log ingestion, to create a comprehensive governance layer.

ITIL 6C Dimension	AI Function	Cetus AI Support
Creation	Generating new content or data.	Enforces data-handling requirements and PII controls via Runtime Guardrails.
Curation	Organising and managing data.	Supports data classification and policy evaluation through the Policy Engine.
Clarification	Extracting insights from data.	Provides visibility into model usage and performance through Observability.
Cognition	Supporting decision-making.	Correlates evidence and human-review context via the Evidence Fabric.
Communication	Interacting with users or systems.	Manages agent permissions and session policies through Agent Governance.
Coordination	Orchestrating tasks across systems.	Integrates with existing orchestration layers via the PDP sidecar.

Enforcing Governance Perspectives

ITIL AI Governance extends the concept of governance across four critical perspectives, recognising that effective AI oversight requires a multidimensional approach: Decision Authority & Risk Management, Ethical Principles & Responsible AI, Data Governance & Performance Management, and Regulatory Compliance & Operational Standards [3]. These perspectives provide a comprehensive framework for ensuring that AI is used responsibly, securely, and effectively.

Cetus AI capabilities are specifically designed to map to these perspectives, supporting operational enforcement within configured coverage and customer-approved operating responsibilities:

1. **Decision Authority & Risk Management**: Establishing clear authority and managing risk are paramount. The Policy Decision API supports this by returning deterministic decisions—approve, deny, or approve with conditions—based on the specific context of the AI request. Crucially, the source platform retains the payload, ensuring that Cetus AI does not become an unnecessary bottleneck or data repository. Furthermore, the Model Risk Register and Shadow AI discovery capabilities address the wider inventory of models, providing visibility into unsanctioned AI use and enabling proactive risk management across the enterprise.
2. **Ethical Principles & Responsible AI**: Translating ethical principles into operational reality is a significant challenge. Runtime Guardrails enforce specific conditions, such as requiring model substitution for sensitive tasks or mandating human-review gates before critical decisions are executed. This ensures that ethical guidelines are not just documented but actively enforced. The Evidence Fabric plays a vital role here by recording the context of human oversight, providing a clear trail of accountability for AI-assisted decisions.

3. **Data Governance & Performance Management**: Understanding the cost and performance of AI is essential for sustainable adoption. The Value Control module attributes model consumption and spend directly to specific teams, workflows, or business activities. It provides a rigorous methodology for distinguishing between forecast value and finance-validated outcomes, preventing the overstatement of AI ROI. Guardrail Analytics connect day-to-day AI operations with performance reporting, giving leaders the insights needed to optimise AI usage and manage costs effectively.
4. **Regulatory Compliance & Operational Standards**: As regulatory scrutiny of AI increases, the ability to demonstrate compliance is critical. The Evidence Fabric correlates events from across the enterprise stack to produce automated-decision registers, individual decision records, and tamper-evident audit chains. This technical evidence supports compliance activities by providing a verifiable record of AI operations. However, it is important to note that this evidence does not replace legal interpretation or guarantee compliance; it provides the factual foundation upon which compliance assessments can be made.

The ITIL AI Governance Improvement Model

The ITIL AI Governance Improvement Model outlines a structured approach for adapting governance: stress-testing current frameworks, defining AI-specific requirements, designing governance adjustments, and operating governance dynamically [3].

Cetus AI supports this model through staged deployment patterns that match control and evidence to the workload's trust boundary:

- **Observe**: Establish visibility and produce evidence without changing runtime decisions. This aligns with stress-testing and defining requirements by discovering Shadow AI and baselining usage.
- **Enforce**: Apply AI-specific policy to selected workloads using the PDP sidecar or dedicated instances. This supports designing and implementing governance adjustments.
- **Orchestrate**: Coordinate governed services across a broader AI estate using a managed or customer-hosted control plane, enabling dynamic governance operations.

Limitations and Customer Responsibilities

While Cetus AI provides a robust framework for operationalising ITIL AI Governance, it is essential to delineate the boundaries of the platform's capabilities and the ongoing responsibilities of the customer. Cetus AI is built and available for controlled enterprise deployment. The platform's architecture is designed to integrate with existing enterprise systems, but this additive approach requires careful planning and execution.

Customer-specific integrations and deployment patterns still require rigorous validation against the customer's architecture, data sources, identity management systems, security protocols, performance requirements, failure behaviour, and overall operating model. 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.

It is crucial to use careful language when describing the platform's regulatory and compliance impact. Cetus AI 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 platform provides technical evidence and operational controls, but it does not replace legal interpretation, governance ownership, risk acceptance, or human decision-making.

The Evidence Fabric correlates approved source records and makes governance gaps visible. However, correlation does not establish causation. The presence of a correlated record indicates that a sequence of events occurred according to the ingested telemetry, but it does not independently verify the underlying intent or the external consequences of those events.

Furthermore, Value Control is an integrated module within the Cetus AI platform, not a separate platform. It distinguishes forecast, observed, attributed, finance-validated, and risk-adjusted value. The accuracy of these outputs relies heavily on the quality of the underlying data and the customer-defined attribution controls. The objective is

not to promise a universal percentage saving, but to provide the attribution and policy controls required to identify avoidable cost, manage budgets, and make model-selection decisions using evidence.

Recommended Next Steps

Organisations seeking to operationalise ITIL AI Governance and move from policy definition to operational enforcement should consider a phased approach, prioritising visibility and evidence before introducing mandatory runtime controls.

1. **Review Capability Status and Evidence Boundaries:** Review the approved public position for Cetus AI capability areas, including availability, deployment modes, dependencies, known boundaries, reviewer roles, and evidence destinations. Understanding these boundaries is the first step in aligning platform capabilities with organisational governance requirements.
2. **Conduct an Architecture Review:** Evaluate the existing enterprise stack to determine the appropriate deployment pattern for specific AI workloads. Consider whether a read-only evidence approach is sufficient for initial baselining, or if a PDP sidecar is required for runtime policy enforcement without disrupting existing orchestration layers.
3. **Establish Visibility with Evidence Fabric:** Begin by ingesting approved telemetry and producing correlated evidence without intervening in runtime decisions. This "Observe" phase allows organisations to discover Shadow AI, monitor workflows, attribute model spend, and estimate AI-related emissions, providing a factual basis for subsequent governance decisions.
4. **Introduce Targeted Runtime Controls:** Once visibility is established, apply AI-specific policy to selected workloads where the risk, architecture, and operating model justify intervention. Use the Policy Engine and Runtime Guardrails to enforce conditions such as PII redaction, model substitution, or human-review gates.
5. **Scale and Orchestrate:** As governance maturity increases, coordinate governed services and identities across a broader AI estate. This may involve governing model routing, agent permissions, and shared control services, ensuring that the organisation's AI governance posture evolves in step with its AI capabilities.

References

[1] PeopleCert. "ITIL AI Governance (Version 5)." https://www.peoplecert.org/browse-certifications/it-governance-and-service-management/ITIL-1/itil-ai-governance-version-5-4234 [2] ITIL. "ITIL AI Governance (Version 5)." https://www.itil.com/professionals/certifications/ITIL-AI-Governance-Version-5 [3] ITIL. "AI Governance White Paper | Responsible AI with ITIL." https://www.itil.com/Itil-News-and-Announcements/ai-governance-white-paper