ENTERPRISE IT · WHITE PAPER · 2026

From ITIL 5 to AI-Native Operations

Why Existing Frameworks Are No Longer Enough — and How DVSG and the 7-Layer Operational Architecture Close the Gap

Authored by Koen Goos (UllC bv)

An Academic White Paper on Enterprise AI Architecture & Governance

DOCUMENT	White Paper · v1.0
DOMAIN	AI-Native Operations · Governance · Architecture
DATE OF PUBLICATION	May 2026
CLASSIFICATION	Public · For Distribution

Abstract

EXECUTIVE SUMMARY

*Enterprise IT is undergoing a paradigm shift that eclipses prior technological transitions such as cloud computing and virtualization. Whereas legacy frameworks like ITIL, TOGAF, and SIAM were engineered for deterministic, human-centric operations, the rise of Artificial Intelligence introduces probabilistic decision-making, dynamic orchestration, and autonomous reasoning. This white paper examines why existing operational frameworks are insufficient for AI-native enterprises and introduces **Dynamic Value Stream Governance (DVSG)** together with a comprehensive **7-Layer AI-Native Operational Architecture**. By bridging the gap between strategic intent and operational execution, this combined model provides a robust foundation for governed, explainable, and scalable AI-native enterprise operations.*

Keywords

AI-Native Operations · ITIL 5 · DVSG · Dynamic Value Stream Governance · Enterprise Architecture · Runtime Governance · Context Engineering · Model Context Protocol (MCP) · Agent Orchestration · Explainability · Human-in-the-Loop · EU AI Act · ISO 42001 · DORA · NIS2

Document Purpose

This document is intended for enterprise architects, CIOs, CTOs, AI governance leads, service management practitioners, and policy makers who are responsible for shaping the next generation of operating models. It is positioned as an academic and practitioner-oriented reference that complements — rather than replaces — established frameworks such as ITIL, TOGAF, SIAM, and COBIT. Its purpose is to provide a coherent vocabulary, an actionable architectural model, and a governance philosophy fit for the AI-native era.

Table of Contents

1. Introduction

Enterprise IT is entering a transition that is fundamentally larger in scope and impact than the move to cloud computing, DevOps, virtualization, or even the internet itself. For decades, enterprise operations were strictly deterministic. Systems behaved according to predefined logic, workflows were orchestrated through fixed processes, and governance was largely static. Documentation described environments that evolved slowly enough for human operators to maintain absolute control.

Artificial Intelligence disrupts this paradigm entirely. Modern AI-native enterprises no longer operate through purely deterministic execution. Instead, they increasingly rely on autonomous reasoning, dynamic orchestration, contextual retrieval, multi-agent collaboration, probabilistic decision-making, runtime optimization, and continuously adapting operational ecosystems. This evolution fundamentally alters the nature of operations itself.

The critical challenge facing organizations today is the attempt to implement AI using operational models that were never designed for autonomous systems. While traditional IT governance frameworks remain extremely valuable — ITIL, TOGAF, SIAM, COBIT, and modern cloud operating models all provide essential foundations — they were architected around **service-centric and human-centric** operations. AI-native enterprises require a paradigm that incorporates **operational cognition, runtime governance, explainability pipelines, context engineering, orchestration fabrics**, and **distributed decision systems**.

This white paper introduces **Dynamic Value Stream Governance (DVSG)** and the **AI-Native Operational Architecture (7-Layer Model)** as a practical operational implementation architecture for the AI-native direction emerging in ITIL 5.

2. The Evolution of Enterprise Operations

The history of enterprise IT can be accurately described as a sequence of progressive abstraction layers. Mainframes abstracted hardware complexity; virtualization abstracted infrastructure; cloud computing abstracted platforms; containers abstracted deployment; and Kubernetes abstracted orchestration. AI now abstracts operational reasoning itself.

ERA	PRIMARY ABSTRACTION	OPERATIONAL IMPLICATION
Mainframe	Hardware	Centralized compute, deterministic batch processing

ERA	PRIMARY ABSTRACTION	OPERATIONAL IMPLICATION
Virtualization	Infrastructure	Resource pooling, elastic capacity
Cloud	Platforms	On-demand services, global distribution
Containers	Deployment	Portable, immutable workload units
Kubernetes	Orchestration	Declarative scheduling and scaling
AI-Native	**Operational Reasoning**	**Autonomous cognition, probabilistic decisions, adaptive governance**

This transition is fundamentally distinct from previous technological shifts. Earlier advancements automated *execution*, whereas AI automates *cognition*. This distinction is critical: once systems begin reasoning dynamically, organizations can no longer rely solely on static governance, predefined integrations, fixed workflows, or centralized human coordination. The operational model itself must evolve to accommodate this shift.

The emergence of ITIL 5 is strategically important in this context. It reflects a growing realization within enterprise IT that future operations will be AI-native. However, while ITIL 5 effectively outlines the governance philosophy, service principles, value orientation, and desired operational outcomes, a significant gap remains between **operational intent** and **operational execution**. This gap becomes acutely visible when organizations attempt to operationalize AI agents, Retrieval-Augmented Generation (RAG) pipelines, orchestration systems, copilots, automation ecosystems, or comprehensive AI governance programs.

3. The Limitations of Existing Frameworks

The core issue is not that existing frameworks are flawed; rather, they were designed for a different operational reality. ITIL successfully transformed service management and operational governance by introducing lifecycle thinking, service value systems, continual improvement, and value co-creation. ITIL 5 evolves this further toward AI-native service ecosystems. However, ITIL primarily defines operational philosophy, governance principles, organizational practices, and service management outcomes. It does not yet fully articulate *how* autonomous AI ecosystems should execute operationally, how runtime orchestration should function, how context should be governed, or how explainability becomes operationalized.

Similarly, TOGAF provides robust enterprise architecture governance, capability mapping, and strategic alignment, yet it largely concludes at design-time governance rather than addressing runtime operational cognition. SIAM excels at coordinating multiple service providers, but AI-native enterprises increasingly require coordination between autonomous digital agents rather than solely external organizations. AIOps platforms provide observability, event correlation, and operational automation, but they rarely address enterprise governance, explainability, Human-in-the-Loop (HITL) control, architectural traceability, or operational value governance.

FRAMEWORK	STRENGTH	GAP FOR AI-NATIVE OPERATIONS
ITIL / ITIL 5	Service value systems, continual improvement	Does not fully define runtime execution of autonomous systems
TOGAF	Enterprise architecture governance	Stops at design-time; lacks runtime cognition
SIAM	Multi-vendor service integration	Not designed for autonomous digital agent coordination
COBIT	Governance and risk control	Static control catalog, limited runtime applicability
AIOps Platforms	Observability and event correlation	Lack governance, explainability, HITL, and value traceability

Consequently, organizations are creating fragmented AI ecosystems characterized by disconnected copilots, isolated RAG engines, unmanaged prompts, shadow AI, duplicated knowledge, and inconsistent governance models. This fragmentation generates **operational entropy**, which becomes increasingly dangerous as the AI ecosystem expands.

4. The Fundamental Shift Introduced by DVSG

Dynamic Value Stream Governance (DVSG) was developed to address a critical governance limitation in modern enterprises. Traditional governance largely focuses on process adherence, compliance, service management, and operational control. DVSG shifts the governance paradigm toward **continuous operational value creation**.

The central question shifts from "Are processes being followed correctly?" to

"Is the operational ecosystem continuously creating measurable value — safely, explainably, and adaptively?"

In AI-native operations, workflows continuously adapt, agents dynamically reason, retrieval pipelines evolve, context changes in real time, and orchestration decisions become probabilistic. Static governance cannot sufficiently control such ecosystems. Governance itself must become dynamic, operational, observable, measurable, explainable, and continuously adaptive.

DVSG integrates governance directly into runtime execution. Governance manifests as operational telemetry, confidence scoring, explainability pipelines, HITL validation, policy enforcement, and adaptive orchestration controls. This facilitates a transition from *governance as documentation* to **governance as runtime architecture**.

5. The AI-Native Operational Architecture — 7-Layer Model

The 7-Layer Operational Model was created to operationalize the transition to AI-native enterprises. It does not replace ITIL, TOGAF, SIAM, or other enterprise architecture frameworks; instead, it acts as the **execution architecture beneath them**. The relationship is clear: ITIL 5 defines *what* AI-native operations should achieve, while the 7-Layer Model defines *how* those operations are implemented technically and operationally.

This distinction is paramount. Most contemporary AI discussions remain tool-centric, whereas the 7-Layer Architecture is operational-centric. It treats AI ecosystems as governed operational systems rather than isolated AI applications.

LAYER	NAME	OPERATIONAL FUNCTION
1	Architecture as Code Foundation	Executable, versioned architectural baseline
2	Knowledge & Memory	Unified, governed operational memory fabric
3	Skills & Intelligence	Cognitive execution within governed boundaries
4	Governance & Guardrails	Runtime, executable governance and compliance
5	Agent Orchestration	Distributed, specialized cognitive collaboration
6	Runtime Automation	Translation of cognition into operational action

LAYER	NAME	OPERATIONAL FUNCTION
7	Human & Interaction	Governed human-AI collaboration and accountability

5.1 Layer 1 — Architecture as Code Foundation

The first layer establishes architectural traceability and operational baseline governance. AI-native operations require significantly more dynamic change than traditional environments. Consequently, architecture itself must become executable, versioned, and operationally integrated. This layer incorporates tools and methodologies such as *Structurizr, DSL-driven architecture*, *BPMN process models, Git-based version governance, infrastructure as code*, and operational model traceability. Architecture evolves from mere documentation into a living operational contract, aligning governance, operations, orchestration, and deployment.

5.2 Layer 2 — Knowledge & Memory Layer

The second layer addresses a major weakness in enterprise AI: **fragmented context**. Most AI systems fail not due to weak models, but because of poor retrieval quality, fragmented operational memory, disconnected documentation, and unmanaged contextual relevance. This layer introduces RAG, scatter-gather retrieval, hybrid search fusion, semantic ranking, parallel context pre-processing, metadata enrichment, and governed knowledge pipelines. Systems such as Obsidian, BookStack, Gitea, Jupyter, PostgreSQL, documentation repositories, and operational telemetry sources are integrated into a unified operational memory fabric, transforming documentation into *operational intelligence*.

5.3 Layer 3 — Skills & Intelligence Layer

The third layer introduces cognitive execution. It encompasses Large Language Models (LLMs), reasoning engines, notebooks, execution runtimes, speech systems, AI models, experimentation environments, and analytical capabilities. Crucially, this intelligence layer operates within a governed architecture. The models consume governed context from Layer 2 and remain constrained by Layer 4 governance controls, preventing uncontrolled autonomous drift.

5.4 Layer 4 — Governance & Guardrails Layer

This layer is arguably the most strategically important within the architecture. While traditional governance frameworks remain static — policies, procedures, controls, and compliance documents — AI-native operations require **runtime governance**. This layer operationalizes HITL approvals, explainability, auditability, confidence scoring, traceability, policy validation, compliance controls, and operational oversight. Governance becomes executable and directly aligns with principles from the *EU AI Act*, *ISO 42001*, *DORA*, *NIS2*, and emerging AI operational regulations. Governance is embedded directly into runtime execution rather than imposed externally.

5.5 Layer 5 — Agent Orchestration Layer

This layer introduces distributed operational cognition. Rather than relying on a monolithic AI system, operational problems are decomposed into specialized expert agents — for example, Enterprise Architect Agents, Security Agents, Governance Agents, Financial Analysis Agents, Technical Architecture Agents, Compliance Agents, and Domain Architecture Agents. Queries are processed in parallel, and results are aggregated through consensus and synthesis mechanisms. This dramatically improves quality, explainability, bias reduction, and operational resilience, serving as the orchestration fabric of the AI-native enterprise.

5.6 Layer 6 — Runtime Automation Layer

The sixth layer transforms cognition into operational execution. It includes protocols and platforms such as MQTT, Node-RED, Camunda, smart home systems, process orchestration engines, industrial systems, automation runtimes, and operational control systems. This layer acts as the *execution nervous system* of the enterprise, translating decisions into actions.

5.7 Layer 7 — Human & Interaction Layer

The final layer preserves human accountability — a fundamental principle in enterprise AI. AI-native operations should augment human decision-making, not eliminate it. This layer introduces voice interfaces, dashboards, holographic assistants, mobile interfaces, conversational systems, and operational visualization layers. The objective transcends mere usability; the goal is **governed human-AI collaboration**.

6. Context Engineering as a Core Discipline

A pivotal concept introduced by this architecture is **Context Engineering**. Traditional software engineering optimized logic, code, APIs, and infrastructure. AI-native systems increasingly optimize *context quality*. The quality of retrieval, ranking, enrichment, preprocessing, and orchestration directly determines the quality of AI outcomes.

This fundamentally alters enterprise architecture. Context becomes an operational asset, a governance concern, and a strategic capability. The architecture leverages scatter-gather retrieval, dual-mode hybrid search fusion, multi-angle parallel queries, and parallel context preprocessing to maximize factual accuracy, explainability, confidence, and operational trustworthiness.

7. MCP — The Interoperability Backbone

The architecture relies heavily on the **Model Context Protocol (MCP)**. MCP functions similarly to how TCP/IP standardized internet communication. Without MCP, AI ecosystems become fragmented; with MCP, agents, tools, memory systems, orchestration engines, and governance services become interoperable. MCP enables secure context exchange, permissions management, tool discovery, auditability, and operational consistency, which are critical for enterprise-scale AI-native ecosystems.

8. Strategic Implications

The current market focuses disproportionately on AI tools, copilots, agents, and automation. Very few organizations are constructing governed AI-native operational architectures. This gap will constitute one of the most significant operational risks of the next decade. Organizations that fail to operationalize governance inside AI runtime systems will inevitably encounter uncontrolled AI sprawl, explainability failures, governance gaps, operational fragmentation, and escalating operational risk. The organizations that succeed will not necessarily possess the smartest models; rather, they will have the most robust operational architecture.

The future of enterprise AI will not be determined by prompts.
It will be determined by architecture.

---

9. Conclusion

The transition toward AI-native enterprises is inevitable. However, realizing AI-native operations requires more than copilots, orchestration engines, or isolated AI tooling. It necessitates operational architecture, runtime governance, context engineering, orchestration fabrics, explainability systems, and adaptive operational ecosystems.

While ITIL 5 increasingly defines the operational philosophy of this future, DVSG extends governance into dynamic operational value management. The 7-Layer Architecture operationalizes the runtime execution model required to implement that vision safely and at scale. Together, they provide a practical blueprint for governed, explainable, scalable, AI-native enterprise operations.

10. Copyright & Intellectual Property Notice

✓