

Enterprise AI Operating Model

A Scalable Operating Model for Strategy, Governance, Platforms, Data, Delivery, Risk, FinOps, Responsible AI, and Enterprise Adoption

Sanjeeve Kumar Gajadi*

*(HCL Technologies, Bengaluru, Karnataka, India)

Email: sanjeevekumar.g@hcltech.com

Abstract:

Artificial Intelligence is rapidly becoming an enterprise-wide capability rather than an isolated technology initiative. Organizations therefore require an operating model that defines how AI strategy, funding, governance, platforms, data, architecture, delivery, security, risk, operations, and business adoption work together at scale. This paper proposes an Enterprise AI Operating Model that connects executive strategy with federated business ownership, centralized guardrails, reusable AI platforms, governed business data, responsible AI, product-oriented delivery, AI FinOps, workforce enablement, and continuous performance management. The model defines organizational roles, decision rights, governance forums, product teams, AI platform services, lifecycle stages, intake and prioritization, architecture standards, risk controls, operational processes, and measurable value outcomes. It is designed to support predictive AI, Generative AI, AI copilots, enterprise agents, intelligent automation, and decision intelligence across hybrid and multi-cloud environments. The proposed operating model enables enterprises to industrialize AI adoption while reducing duplication, improving reuse, strengthening governance, accelerating delivery, and ensuring sustainable business value.

Keywords — Enterprise AI, AI Operating Model, Responsible AI, AI Governance, AI Platform, SAP BTP, Business Data Cloud, AI Product Management, AI FinOps, Enterprise Architecture, AI Center of Excellence, Federated Governance

I. INTRODUCTION

Enterprise AI adoption is moving from experimentation toward industrialization. Business units are deploying predictive models, Generative AI assistants, AI agents, intelligent automation, analytics, and embedded decision support across finance, supply chain, sales, service, manufacturing, human resources, software engineering, and corporate functions. Without a coherent operating model, these initiatives frequently become fragmented, expensive, difficult to govern, and hard to scale.

An Enterprise AI Operating Model defines how strategy, organization, processes, platforms, data, governance, funding, talent, risk, operations, and performance management work together to deliver AI capabilities. It clarifies who makes which decisions, which capabilities are centralized, which are federated, how reusable services are governed, and how business value is measured.

II. BUSINESS CHALLENGES

Organizations commonly face duplicated AI platforms, inconsistent model development practices, limited business ownership, weak data quality, uncertain accountability, uncontrolled experimentation, overlapping vendor tools, uncoordinated security

controls, and difficulty moving pilots into production. AI initiatives may be funded as isolated projects without lifecycle budgets for monitoring, retraining, cloud consumption, support, and governance.

These challenges are amplified by Generative AI because model choice, prompt engineering, Retrieval-Augmented Generation, enterprise knowledge, data residency, model risk, agentic automation, and usage-based cloud costs require coordinated decisions across many organizational functions.

III. OPERATING MODEL PRINCIPLES

The proposed model follows principles of business-value ownership, platform reuse, federated innovation, centralized guardrails, responsible AI by design, secure data access, product-oriented delivery, automation-first operations, transparent cost management, measurable outcomes, and continuous improvement.

A centralized team should not become a bottleneck for every AI use case. Conversely, fully decentralized AI adoption leads to technology sprawl and inconsistent controls. The model therefore combines enterprise standards and shared platforms with domain-level product ownership and delivery accountability.

IV. ORGANIZATIONAL STRUCTURE

Executive leadership defines AI ambition, risk appetite, investment priorities, and strategic outcomes. An Enterprise AI Steering Committee aligns funding, portfolio priorities, and cross-enterprise policy. A Responsible AI and Risk Council governs ethical, legal, privacy, security, and model-risk requirements. Enterprise Architecture defines reference architectures, standards, and technology roadmaps.

An AI Center of Excellence or AI Enablement Office provides reusable methods, patterns, training, evaluation standards, and advisory services. AI Platform Engineering operates shared model, data, integration, observability, and deployment capabilities. Federated business AI product teams own domain use cases and measurable business outcomes.

V. DECISION RIGHTS

Decision rights should be explicit to avoid governance ambiguity. Enterprise-level decisions include strategic model providers, approved AI platforms, identity architecture, data classification, security controls, regulatory policies, model evaluation standards, and risk acceptance thresholds. Domain teams decide backlog priorities, user experience, domain-specific data products, and operational adoption within enterprise guardrails.

High-risk use cases should require independent review before production deployment. Low-risk use cases may follow pre-approved patterns and automated controls. This tiered decision model accelerates routine adoption while preserving stronger governance for sensitive workloads.

VI. AI PRODUCT LIFECYCLE

The operating model uses a product lifecycle rather than a one-time project lifecycle. Stages include discovery, qualification, prioritization, architecture, data readiness, model development or configuration, evaluation, deployment, adoption, monitoring, optimization, and retirement.

Each AI product has a product owner, business sponsor, technical owner, data owner, risk owner, service-level objectives, value hypothesis, cost model, and operational support model. This ensures continued accountability after the initial launch.

VII. AI PLATFORM MODEL

A shared AI platform reduces duplicated infrastructure and provides standardized capabilities for model access, prompt management, RAG, vector

services, model orchestration, agent runtime, API integration, identity, secrets, observability, evaluation, deployment, and policy enforcement. SAP Business Technology Platform can provide enterprise services for integration, data, application development, identity, AI, and operations.

Platform teams should expose self-service capabilities with guardrails. Product teams consume approved platform services through templates, APIs, pipelines, and reusable reference architectures rather than building foundational capabilities independently.

VIII. DATA AND BUSINESS KNOWLEDGE

Enterprise AI requires trusted business context. The operating model therefore integrates data governance, metadata, semantic models, business data products, lineage, quality, master data, privacy, retention, and access controls. SAP Business Data Cloud and SAP Datasphere can support governed enterprise data and semantic business context.

Knowledge management should include structured data, unstructured documents, policies, process content, metadata, and enterprise taxonomies. RAG solutions should be treated as governed data products with ownership, quality metrics, access controls, freshness requirements, and lifecycle management.

IX. RESPONSIBLE AI AND RISK

Responsible AI is embedded into the operating model through policy, risk classification, evaluation, human oversight, auditability, incident management, and periodic recertification. Controls should address fairness, explainability, transparency, privacy, robustness, security, intellectual property, safety, and appropriate human involvement.

The risk process should be integrated with enterprise risk management rather than isolated within technical teams. Material residual risks should be visible to accountable business executives.

X. DELIVERY MODEL

AI delivery teams should combine product management, domain expertise, enterprise architecture, data engineering, AI engineering, integration, security, user experience, platform engineering, testing, and change management. Agile delivery supports iterative experimentation, but production promotion should be governed through defined quality and risk gates.

Reusable evaluation suites should measure model quality, groundedness, hallucination, toxicity, bias, privacy leakage, cybersecurity resilience, latency,

availability, and cost. Release pipelines should capture evaluation evidence automatically.

XI. AI FINOPS AND FUNDING

AI economics require active management because usage-based model inference, vector storage, data processing, GPU resources, and cloud services may produce rapidly changing costs. The operating model therefore integrates AI FinOps with product-level chargeback or showback, unit-cost metrics, budget thresholds, and model-routing optimization.

Funding should distinguish foundational platform investment from business product investment. Shared platform capabilities are financed as enterprise assets, while business units remain accountable for use-case value realization and domain operating costs.

XII. OPERATIONS AND OBSERVABILITY

AI operations combines model operations, platform operations, data operations, security operations, and business operations. Monitoring should include availability, latency, cost, model quality, drift, retrieval quality, security events, policy violations, data freshness, adoption, and business outcomes.

Incident processes should classify AI-specific failures such as harmful outputs, hallucinations, policy breaches, data exposure, unauthorized actions, model outages, and degraded recommendations. Operational teams require playbooks for rollback, model switching, prompt correction, data remediation, and temporary suspension.

XIII. WORKFORCE AND CHANGE MANAGEMENT

Successful AI adoption requires new workforce capabilities. Executives need AI governance literacy, product owners need value and risk literacy, architects need AI solution patterns, developers need secure AI engineering practices, and business users need training on appropriate use, verification, and escalation.

Communities of practice, role-based learning paths, reference implementations, internal certification, and reusable playbooks help scale responsible adoption. Change management should measure adoption behavior, user trust, process redesign, and realized productivity rather than only technical deployment.

XIV. PORTFOLIO MANAGEMENT

Enterprise portfolio governance should maintain an inventory of AI initiatives, products, models, agents, vendors, datasets, and dependencies. Portfolio

dashboards can identify duplication, unmanaged risk, low-value pilots, platform reuse opportunities, and strategic capability gaps.

Prioritization should consider business value, strategic alignment, data readiness, delivery feasibility, risk, regulatory impact, adoption potential, and total lifecycle cost. This prevents organizations from pursuing technically attractive but operationally weak use cases.

XV. MATURITY MODEL

AI operating maturity can be assessed across strategy, governance, organization, architecture, data, platforms, delivery, responsible AI, operations, talent, FinOps, and value management. Early-stage organizations operate through isolated pilots; developing organizations establish common standards; mature organizations run shared platforms and federated products; leading organizations use continuous controls, automated assurance, reusable business knowledge, and measurable AI portfolio optimization.

Maturity assessments should produce a practical roadmap rather than a score alone. Each gap should map to ownership, investment, target state, and measurable outcome.

XVI. KEY METRICS

Business metrics include revenue improvement, cost reduction, cycle-time reduction, productivity, customer satisfaction, decision quality, automation rate, and user adoption. Platform metrics include reuse, deployment lead time, availability, model latency, cost per transaction, and environment utilization. Governance metrics include review cycle time, policy compliance, risk exceptions, audit findings, and recertification completion.

Metrics should be balanced. High adoption without quality or controls is not success, and strict governance that prevents delivery is also not success. The operating model must optimize value, speed, trust, and sustainability together.

XVII. CONCLUSION

Enterprise AI cannot scale sustainably through disconnected experiments. It requires an operating model that connects executive strategy, business ownership, shared platforms, governed data, architecture, responsible AI, delivery, operations, talent, funding, and performance management.

The proposed Enterprise AI Operating Model provides a practical structure for industrializing AI across the enterprise. By combining centralized

guardrails with federated product ownership and reusable platform services, organizations can accelerate adoption while reducing duplication, controlling risk, improving transparency, and maximizing measurable business value.

TABLE I — ENTERPRISE AI OPERATING MODEL RESPONSIBILITY MATRIX

Capability	Enterprise Ownership	Federated Ownership
AI Strategy	Executive Steering Committee	Business sponsors
Architecture Standards	Enterprise Architecture	Domain architects
AI Platform	Platform Engineering	Product teams consume services
Data Governance	Enterprise Data Office	Domain data owners
Responsible AI	AI Risk Council	Product risk owners
Delivery	Methods and guardrails	Cross-functional product teams
Operations	Shared observability	Domain service ownership
FinOps	Enterprise cost model	Product budget accountability

REFERENCES

- [1] ISO/IEC 42001:2023, Information technology — Artificial intelligence — Management system.
- [2] National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0), 2023.
- [3] National Institute of Standards and Technology, AI RMF Generative Artificial Intelligence Profile, 2024.
- [4] ISO/IEC 27001:2022, Information security, cybersecurity and privacy protection — Information security management systems.
- [5] ISO 31000:2018, Risk management — Guidelines.
- [6] The Open Group, TOGAF Standard, 10th Edition, 2022.
- [7] ISACA, COBIT 2019 Framework: Governance and Management Objectives, 2019.
- [8] AXELOS, ITIL Foundation: ITIL 4 Edition, 2019.
- [9] Cloud Security Alliance, Security Guidance for Critical Areas of Focus in Cloud Computing.
- [10] CNCF, Cloud Native Landscape and Kubernetes Documentation.
- [11] SAP SE, SAP Business Technology Platform Documentation.
- [12] SAP SE, SAP AI Core and SAP AI Launchpad Documentation.
- [13] SAP SE, SAP Integration Suite Documentation.
- [14] SAP SE, SAP HANA Cloud Documentation.

- [15] SAP SE, SAP Cloud Identity Services Documentation.
- [16] SAP SE, SAP Cloud ALM Documentation.
- [17] SAP SE, SAP Business Data Cloud and SAP Datasphere Documentation.
- [18] SAP SE, SAP LeanIX Enterprise Architecture Documentation.
- [19] SAP SE, SAP Signavio Process Transformation Suite Documentation.
- [20] OECD, Recommendation of the Council on Artificial Intelligence.
- [21] European Union, Artificial Intelligence Act.
- [22] IEEE, Ethically Aligned Design: A Vision for Prioritizing Human Well-being with Autonomous and Intelligent Systems.
- [23] Project Management Institute, PMBOK Guide, Seventh Edition, 2021.
- [24] DAMA International, DAMA-DMBOK2: Data Management Body of Knowledge, 2017.

AUTHOR BIOGRAPHY

Sanjeeve Kumar Gajadi is an SAP Enterprise Architect and Lead Consultant with extensive experience in enterprise architecture, SAP Business Technology Platform, data and analytics, integration, cloud architecture, security, governance, and AI-enabled transformation. His professional interests include enterprise AI operating models, SAP BTP, data and AI platforms, governance, cloud architecture, and digital transformation.

ACKNOWLEDGMENT

The author acknowledges the contributions of standards bodies, enterprise architecture communities, SAP documentation teams, and responsible AI practitioners whose published guidance has informed the concepts synthesized in this paper.