Digital Enterprise Architecture Advisors

Operating Models & Enterprise Architecture in the Age of AI

How the choice of business operating model shapes the design, governance, and ambition of an AI Target Operating Model — a strategic guide for C-level executives and senior enterprise architects.

Digital Enterprise Architecture Advisors LLC

Document type: Strategic White Paper on AI TOM Design

Domain: Enterprise Architecture & Digital Strategy

Sub-domain: AI Target Operating Model Design

Primary audience: C-suite (strategic framing) and Chief Enterprise Architects (architectural implications)

Version: 1.0

Date: July 2026

Author: Christian Kobsa; Digital Enterprise Architecture Advisors LLC

Classification: Internal Use / Non-Confidential

Table of Contents

Executive Summary

Artificial intelligence is no longer a peripheral capability that organizations can graft onto existing structures. It is rapidly becoming the primary mechanism through which enterprises sense market signals, make decisions, automate work, and create competitive differentiation. This transformation demands a commensurate evolution in how organizations design their operating models and enterprise architecture.

This white paper addresses a question that sits at the intersection of strategy and architecture: how does the fundamental choice of business operating model shape the design of an AI Target Operating Model (AI TOM)? The answer has profound implications for where to invest, how to govern, what to centralize, and how fast an organization can realistically move.

The paper is structured around the Ross, Weill, and Robertson framework from MIT's Center for Information Systems Research (CISR), which defines four archetypal operating models along two dimensions: business process standardization and business process integration. These two dimensions — and the critical diagnostic questions that accompany them — determine not only the appropriate enterprise architecture, but the entire logic of how AI capability should be built, governed, and scaled.

> *The central argument of this paper: there is no universally correct AI Target Operating Model. The right AI TOM is the one that is coherent with — and, in some cases, deliberately designed to evolve — the organization's fundamental operating model choice. Getting this alignment wrong is one of the most common and costly failures in enterprise AI programs.*

The paper covers: the foundational operating model framework and its architectural implications; a structured analysis of how each of the four operating model types shapes the AI TOM across data, governance, workforce, and technology dimensions; the relationship to leading AI architecture frameworks including MIT CISR's emerging research, McKinsey's agentic organization model, and the AI-First Enterprise Architecture stack; and a set of strategic recommendations for the Chief Enterprise Architect and C-suite.

1. Foundations: Operating Models and Enterprise Architecture

1.1 The MIT CISR Operating Model Framework

Enterprise architecture, as defined by MIT CISR, is 'the organizing logic for business processes and IT capabilities reflecting the integration and standardization requirements of the firm's operating model.' This definition is deceptively simple. It establishes that architecture is not a technology exercise — it is a strategic expression of how the organization has chosen to operate.

The operating model itself is determined by two fundamental questions, which every organization must answer honestly before designing its architecture or its AI strategy:

1. To what extent is the successful completion of one business unit's transactions dependent on the availability, accuracy, and timeliness of other business units' data?
2. To what extent does the company benefit by having business units run their operations in the same way?

The first question measures the required degree of business process integration. The second measures the required degree of business process standardization. Together, they produce a 2×2 matrix of four operating model archetypes:

	← Low Integration	High Integration →
Low Standardization	**DIVERSIFICATION** *Low Std / Low Int* Independent BUs, own processes and data. Holding company archetype.	**COORDINATION** *Low Std / High Int* Shared data, flexible processes. Common customer view across divisions.
High Standardization	**REPLICATION** *High Std / Low Int* Identical processes, independent data. Franchise or multi-site archetype.	**UNIFICATION** *High Std / High Int* Single integrated enterprise. Airline, logistics, platform archetype.

These four archetypes are not merely descriptive labels — they carry direct, concrete implications for every major architectural decision: shared services strategy, data platform design, application rationalization, governance structures, and now, AI capability design.

1.2 Why Operating Model Choice Precedes Architecture

A recurring failure pattern in enterprise transformation programs is the inversion of this sequence: organizations deploy technology platforms — ERP systems, data warehouses, AI tools — and then discover that the platform's inherent logic conflicts with the actual operating model. A highly federated conglomerate that implements a unified ERP creates governance chaos. A tightly unified airline that fragments its AI investment by business function undermines the very integration that gives it operational advantage.

The same principle applies, with even greater force, to AI. AI systems are uniquely sensitive to data architecture decisions. A large language model hallucinating on poorly governed, inconsistently defined data is not an AI problem — it is an architecture problem rooted in a mismatch between operating model requirements and data design. An agentic AI system that cannot traverse organisational data boundaries because those boundaries were never meant to be crossed is not a technology limitation — it is an operating model constraint.

> *The Chief Enterprise Architect's role is to ensure that AI investment decisions are made downstream of — and coherent with — a clear, explicit, and board-endorsed operating model choice. Where that choice is ambiguous or contested, clarifying it is the first architectural task.*

1.3 The AI-First Enterprise Architecture Stack

For context throughout this paper, we reference the AI-First Enterprise Architecture framework, which defines a four-layer stack for organizations embedding AI into core operations:

- Layer 1 — AI Layer: Intelligence and Automation. Houses AI copilots, large language models, autonomous agents, ML models, decision engines, predictive analytics, and automation workflows. The quality of output from this layer is entirely dependent on what flows upward from the layers below.
- Layer 2 — Semantic Layer: Meaning and Context. The architecture's most distinctive and often under-invested layer. Provides business ontologies, knowledge graphs, a business glossary, operational rules and constraints, and explainability components. Without this layer, AI systems become black boxes that different business units trust differently.
- Layer 3 — Data Products: Domain-Owned, Reusable Assets. Versioned, discoverable, quality-certified data assets owned by domain teams. Each carries a service level agreement, a schema contract, and access policies. AI model teams consume products, not raw sources.

- Layer 4 — AI-Ready Data Management: Foundation and Governance. The substrate on which everything rests. Encompasses governance, data quality, catalogue and metadata, lineage, the Feature Store, and the Vector Database for retrieval-augmented generation.

How much of this stack an organization should build centrally, how much should be federated to business units, and which layers are even technically feasible — all of these questions are answered primarily by the operating model choice. The operating model is the lens through which the architecture stack must be interpreted.

2. The Four Operating Models and Their AI TOM Implications

The following analysis examines each operating model archetype in turn, addressing its AI governance model, data strategy, model ownership, semantic layer requirements, agentic orchestration design, workforce implications, and primary risk profile.

2.1 Diversification — Low Standardization, Low Integration

Organisational archetype

The Diversification model characterizes holding companies and conglomerates whose business units operate in entirely different markets, serve different customers, and employ fundamentally different business processes. A group owning a logistics subsidiary, a media business, and a healthcare services company is the canonical example. Each unit is, in operational terms, a standalone enterprise. The group holding structure exists for financial, capital allocation, or portfolio management reasons — not to create operational synergies between units.

Implications for the AI Target Operating Model

The Diversification model is structurally the most challenging for enterprise-wide AI because the two inputs AI most depends on — shared data and consistent processes to learn from — are precisely what this model withholds by design. There is no meaningful enterprise data layer to feed a shared AI platform, and no process consistency against which to train shared models.

This does not render AI ineffective — it renders enterprise-wide AI architecturally unjustifiable. The AI TOM must therefore be radically decentralized:

- Each business unit builds and owns its own AI capabilities independently, matched to its domain, data, and competitive context.
- The corporate center's role is minimal: it may negotiate shared cloud infrastructure agreements, set a thin governance framework covering ethics, risk, and regulatory compliance, and establish capital allocation criteria for AI investment. It does not mandate platforms, models, or data architectures.
- There is no enterprise-wide semantic layer, knowledge graph, or feature store. Each unit that requires these capabilities builds its own, sized to its own needs.
- The AI Centre of Excellence, if one exists at all, functions as an advisory body or internal consultancy — sharing knowledge and best practices — rather than as a platform owner or governance authority.

> *The temptation in Diversification organizations is to impose a unified AI platform for cost efficiency. This temptation should generally be resisted. The absence of shared processes means a common platform generates more governance overhead than value. The correct enterprise-level AI investment is in shared infrastructure — cloud, compute, security tooling — not in shared intelligence.*

Data strategy

Data lakes and AI platforms are built and maintained independently by each business unit. Enterprise data governance extends only as far as PII classification standards, security baselines, and regulatory compliance requirements that apply across all units regardless of their market context.

Workforce and governance

Workforce AI transformation is a business unit initiative. The corporate center may establish a minimum AI literacy standard and an ethics policy, but role redesign, upskilling investment, and AI adoption strategy are entirely the responsibility of individual business unit leadership. This is not a failure of enterprise ambition — it is an appropriate expression of the operating model's logic.

Primary risk

The dominant risk in Diversification organizations is not over-centralization — it is the inverse: duplication of effort, inconsistent standards, and missed opportunities for cross-unit learning. The AI CoE's most valuable contribution in this model is building a knowledge sharing network that allows business units to learn from each other's AI investments without requiring architectural integration.

2.2 Coordination — Low Standardization, High Integration

Organisational archetype

The Coordination model characterizes organizations whose business units serve overlapping customers or share critical operational data, but execute their core processes in different ways. A financial services group whose retail banking, wealth management, and insurance divisions all serve the same client population — and therefore need a consistent, accurate view of that client — but who each operate with their own processes, risk frameworks, and product structures is the canonical example. Healthcare systems with multiple clinical specialties sharing a patient record, and professional services firms serving shared accounts across different practice areas, also exhibit this pattern.

Implications for the AI Target Operating Model

Coordination organizations have the shared data that AI needs but lack the process standardization that would allow shared AI models to be trained and deployed uniformly across the enterprise. The AI TOM is therefore data-unified but model-federated:

- The primary enterprise-level AI investment is a shared data foundation: a unified customer or entity view (a 'Client 360' or 'Member 360'), a shared data catalogue, lineage tracking, and a governance layer. This is the natural and highest-value home for an enterprise knowledge graph and semantic layer.
- AI models themselves are built and owned by business units, tailored to their specific processes — but they consume from a common, governed data platform. The enterprise data products layer is the critical shared asset; the AI layer above it is federated.
- The shared semantic layer is the most critical enterprise-level AI investment in the Coordination model. It ensures that when different business units' AI systems reason about 'the same customer,' they are working from the same underlying facts, definitions, and relationship structures. Without this, AI proliferation creates a new category of cross-unit inconsistency that is harder to detect and more dangerous than inconsistency in manual processes.
- The AI CoE functions primarily as a data and platform owner, not as a model owner. It governs the shared data products, maintains the business ontology and glossary, sets integration standards, and runs the knowledge graph infrastructure.

> *The Coordination model is where knowledge graphs and retrieval-augmented generation deliver the most distinctive business value. The organization's core architectural commitment is already to shared entity integration — the AI investment amplifies this commitment rather than creating new infrastructure.*

Data strategy

Enterprise investment in a unified semantic layer — ontologies, knowledge graphs, a business glossary, and an entity resolution capability — is the defining data architecture commitment in the Coordination model. Individual business unit data products feed into this shared layer. AI teams across the enterprise consume from it.

Governance imperatives

The Coordination model creates a distinctive AI governance risk: two business units' models may reach different — and potentially contradictory — conclusions about the same customer or counterparty, because they have trained on the same underlying data but with different objectives and feature engineering choices. In regulated industries, where inconsistent decisions create legal and regulatory exposure, this risk must be addressed explicitly in the AI governance framework. Explainability, audit trail requirements, and cross-unit decision consistency standards are non-negotiable governance components in this model.

Workforce and primary risk

Workforce AI transformation requires particular attention to cross-unit collaboration. AI models that reason across shared customer data must be designed and validated by teams that include representation from all business units that share that data — not just the unit that builds the model. The primary architectural risk is the proliferation of business unit AI models that each optimize for local performance metrics while collectively degrading the consistency and trustworthiness of enterprise-wide intelligence.

2.3 Replication — High Standardization, Low Integration

Organisational archetype

The Replication model characterizes organizations that operate the same business model across multiple geographies, sites, or franchises — where each unit is self-contained and does not depend on other units' data, but every unit operates to the same process standard. Global franchise networks, multinational retailers operating an identical store model, and professional services firms with standardized methodologies deployed across independent country practices are the canonical examples.

Implications for the AI Target Operating Model

Replication organizations possess what AI requires most for high-quality model development: large volumes of process-consistent data generated across many independent instances of the same business activity. The same transaction type, the same customer interaction pattern, the same operational decision occurs thousands or millions of times across independent units — creating ideal conditions for training generalizable, high-quality AI models. The AI TOM is therefore model-centralized but data-locally-deployed:

- The enterprise builds AI models centrally, trains them on aggregated data from all units (subject to privacy, data residency, and jurisdictional constraints), and deploys them as a shared model library that each unit instantiates and runs locally.
- The Feature Store — which standardizes the engineering of ML features to ensure consistency between training and inference — is the defining shared infrastructure in this model. Its value is directly proportional to the standardization of the underlying processes.
- Local units consume central models but do not need to share their operational data with each other in real time. The integration requirement remains low, consistent with the operating model logic.
- The AI CoE functions as a model factory and platform owner: it defines the model catalogue, manages training pipelines, governs retraining cycles, monitors model drift across all deployments, and publishes model updates to all units simultaneously.
- Automation Workflows, Decision Engines, and AI agents are strong candidates for centralized development and distributed deployment — because the decisions they make are structurally identical across units.

> *The Replication model offers the highest return on centralized AI investment of any operating model type. A single model improvement, once validated, delivers value simultaneously across every unit in the network. Conversely, the cost of a flawed central model is also amplified across every unit — making robust model validation and staged rollout disciplines critical governance requirements.*

Localization imperative

Replication organizations typically operate across multiple jurisdictions with different regulatory environments, languages, and local market conditions. A centrally trained AI model may perform differently — or be legally impermissible — in certain geographies. This is not a theoretical concern: the EU AI Act, for example, classifies many workplace and customer-facing AI applications as high-risk, imposing transparency, human oversight, and notification requirements that may not apply in other jurisdictions. The AI TOM must incorporate a localization layer: structured mechanisms to fine-tune, constrain, or adapt central models for local regulatory and market contexts without compromising the standardization that makes them cost-effective to maintain.

Workforce and primary risk

Workforce redesign in the Replication model has a unique advantage: because processes are standardized, a role redesign validated in one unit can be rolled out across all units. This compresses the time and cost of AI-driven workforce transformation significantly. The primary risk is brittleness: a flawed central model, a biased training dataset, or a governance failure affects every unit simultaneously. Model monitoring, staged deployment, and clear rollback procedures are essential risk management disciplines.

2.4 Unification — High Standardization, High Integration

Organisational archetype

The Unification model characterizes enterprises that operate as a single integrated entity: shared processes, shared data, shared systems, and a single consistent view of the business at every level. Global airlines, logistics operators, platform businesses, and single-product manufacturers with unified global supply chains are the canonical examples. The enterprise competes as one — local variation is an exception to be managed, not a design principle.

Implications for the AI Target Operating Model

Unification organizations are architecturally the most AI-ready, because they already possess both conditions AI requires: shared, integrated data and consistent, standardized processes to reason about. The full AI-First Enterprise Architecture stack — all four layers — can be built once and deployed enterprise-wide. The AI TOM is therefore fully centralized and enterprise-wide:

- The complete AI-First Architecture stack is justified and valuable at the enterprise level. A single Feature Store, Vector Database, Knowledge Graph, model catalogue, and AI platform serves all business functions.
- Agentic orchestration is most powerful in the Unification model. AI agents can traverse the full enterprise data and process graph — because it is unified. A customer-facing agent handling a complex enquiry can access supply chain, financial, and operational data in a single task graph without crossing organisational or data boundaries.
- The Decision Engine and Automation Workflows can enforce enterprise-wide rules consistently, because there are no legitimate process variations to accommodate.
- The AI CoE owns the full stack: platforms, models, governance, standards, and the agentic workforce. It is a strategic enterprise function, not an advisory or shared-services body.

- AI governance is most straightforwardly centralized in this model — a single governance framework, a single ethics policy, and a single model lifecycle management process apply to all AI deployments.

> *The Unification model offers the most coherent foundation for an enterprise agentic network — the emerging model in which specialist AI agents collaborate across organisational functions, orchestrated by a central coordination layer, to execute complex multi-step business processes autonomously. This is the direction in which the most advanced AI architectures are evolving.*

Resilience and innovation imperatives

Unification organizations face a distinctive set of AI risks. A single AI model or platform failure can affect the entire enterprise simultaneously — there is no unit-level isolation. The AI TOM must invest proportionally in resilience, model redundancy, fallback procedures, and incident response capabilities. There is also an innovation risk: highly unified architectures can become slow to change, and AI capabilities are evolving at a pace that rewards rapid iteration. The governance model must include structured sandboxing and experimentation mechanisms that allow new AI capabilities to be tested and validated without destabilizing the core platform.

Workforce and primary risk

Workforce transformation in the Unification model is a single, coordinated enterprise program — the most efficient form of AI-driven role redesign. The primary risk is over-centralization: when the AI platform becomes a bottleneck for every AI initiative in the enterprise, innovation slows and business unit frustration mounts. The governance model must balance central control with structured pathways for business-unit-initiated AI experimentation within the enterprise platform's guardrails.

3. Cross-Framework Synthesis

3.1 The AI TOM Comparison Matrix

The following matrix summarizes the AI TOM implications across all four operating model archetypes:

Dimension	Diversification	Coordination	Replication	Unification
AI Governance	Fully federated — BU autonomy	Hybrid: central data, federated AI	Central model factory, local deploy	Fully centralized CoE
Data Strategy	Per-BU data lakes	Enterprise shared entity layer	Std data products, local instances	Single unified data platform
AI Model Ownership	BU-owned, no sharing	BU models, enterprise data	Central models, BU consumption	Enterprise models, central ops
Semantic Layer	None / per-BU	Enterprise knowledge graph critical	Standardized ontologies, replicated	Single enterprise ontology
Agent Orchestration	BU-level agents only	Shared customer agents, BU ops	Central agent templates deployed	Full enterprise agentic network
Primary AI Risk	Fragmentation & duplication	Inconsistent cross-BU decisions	Localization & regulatory variance	Brittleness & over-centralization
CoE Role	Advisory / policy only	Data platform owner	Model factory & lifecycle owner	Full stack platform & governance

3.2 Relationship to the McKinsey Agentic Organization Framework

McKinsey's emerging framework for the agentic organization identifies five pillars of transformation: business model, operating model, governance, workforce and culture, and technology and data. It argues that AI agents — systems capable of planning, reasoning, and executing multi-step tasks autonomously — are moving from experimental capability to core operational infrastructure, with the range of tasks AI can complete without human supervision expanding rapidly.

The McKinsey framework's internal decomposition of the AI layer adds important structure to the four-layer AI-First stack: the AI layer is not a flat collection of tools, but a managed execution environment with a dedicated coordination plane — the agentic orchestration layer — that decides which agents are invoked, in what sequence, with what context, and under what constraints. Agents act only through a standardized toolbox abstraction layer, and backend systems are accessed exclusively through this toolbox. This architectural constraint is what ensures agents never bypass governance.

The operating model matrix maps directly onto the McKinsey framework as follows:

- Diversification: Agentic orchestration exists only at the business unit level. There is no enterprise-wide agent coordination layer. Each unit implements its own agent workforce appropriate to its domain.
- Coordination: A shared agentic capability exists for processes that traverse shared data — for example, a customer intelligence agent that queries the enterprise knowledge graph. Business unit operational agents remain federated.
- Replication: Central agent templates are developed and validated by the AI CoE, then deployed to all units as standardized configurations. The orchestration layer is centralized; execution is distributed.
- Unification: A full enterprise agentic network is both technically feasible and strategically justified. Specialist agents — customer, finance, operations, risk — are coordinated by a central orchestration layer that has access to the unified enterprise data and process graph.

3.3 Relationship to MIT CISR's Emerging AI Research

MIT CISR's most recent research, published in late 2025, identifies a new two-dimensional framework for enterprise IT operating models in the AI era, with two dimensions — the degree of shared enterprise platforms and the degree of business-unit-driven innovation — producing four new IT operating model archetypes. Critically, CISR's 2026 research program explicitly investigates how the enterprise IT operating model affects AI initiative intake, delivery, and governance.

CISR's research to date suggests that top-performing enterprises in the AI era are achieving both speed and scale by combining modularity (which enables rapid experimentation) with reuse architected by strong IT leadership (which preserves governance and efficiency). This finding is directly coherent with the operating model matrix: the appropriate balance between modularity and reuse depends on whether the operating model calls for standardization, integration, both, or neither.

CISR's four-stage Enterprise AI Maturity Model — which progresses from piloting AI, through building capabilities, to scaled AI ways of working, and finally to governed and AI-driven enterprise operations — provides a useful maturity dimension to overlay on the operating model matrix. The rate at which an organization can progress through these maturity stages is directly shaped by its operating model: Unification organizations can move fastest, because each maturity investment applies enterprise-wide; Diversification organizations progress most slowly at the enterprise level, but may have business units that are highly advanced in isolation.

3.4 The Operating Model as a Strategic Variable, Not a Fixed Constraint

The most strategically significant insight for the Chief Enterprise Architect is that the operating model is not a permanent constraint — it is a strategic variable that AI can influence and, in some cases, justify changing.

AI creates new integration value that previously did not exist. A Diversification organization that builds a shared AI-driven customer intelligence platform may find itself migrating toward Coordination — not by managerial decree, but because the economic value of shared AI infrastructure creates integration incentives that outweigh the historical rationale for full business unit independence.

A Replication organization that deploys centralized AI models and feeds operational performance data from all units back into a shared optimization platform may find that units which were previously independent begin to inform and improve each other's operations — a form of emergent integration that pushes the operating model toward Unification.

> *The Chief Enterprise Architect is uniquely positioned to model and communicate these operating model migration dynamics to the C-suite. The question is not only 'what AI TOM fits our current operating model?' but also 'what operating model should we be migrating toward, and how can our AI architecture accelerate that migration?'*

This is a strategic advisory role that goes beyond technology. It requires the CEA to participate in — and in some cases lead — the conversation about organisational design and long-term business strategy. The architecture follows the operating model; but the architecture can also make new operating model positions viable that were previously impractical.

4. Strategic Recommendations for the Chief Enterprise Architect

4.1 Establish Operating Model Clarity Before AI Architecture Decisions

The most consequential action a Chief Enterprise Architect can take before designing an AI TOM is to ensure that the organization's operating model choice is explicit, documented, and endorsed at board level. In many large international organizations, the operating model is implied rather than stated — and different parts of the leadership team hold different assumptions about it.

Use the two diagnostic questions from the MIT CISR framework as a structured executive conversation:

- To what extent is the successful completion of one business unit's transactions dependent on the availability, accuracy, and timeliness of other business units' data?
- To what extent does the company benefit by having business units run their operations in the same way?

These questions frequently surface productive strategic disagreements that, left unresolved, will generate expensive conflicts in AI program delivery. Surfacing and resolving them at the outset is an architectural investment with a high return.

4.2 Design the AI Governance Model to Match the Operating Model

AI governance is not a generic framework to be applied uniformly. The appropriate governance design depends directly on the operating model:

- Diversification: Light enterprise governance (ethics, risk, regulatory compliance baselines) with maximum business unit autonomy. A CoE that advises rather than controls.
- Coordination: Robust enterprise data governance and a shared semantic layer. Business unit AI model governance is federated, but shared data asset governance is centralized and non-negotiable.
- Replication: Centralized model governance with a disciplined model factory operating model. Localization governance for regulatory and market variation. A staged rollout discipline that tests model changes before enterprise-wide deployment.

- Unification: Full centralized AI governance — model lifecycle management, ethics review, performance monitoring, and incident response — as an enterprise function with board-level visibility.

4.3 Invest in the Right Layers of the AI-First Architecture Stack

The four-layer AI-First architecture stack should not be built in its entirety by every organization. The operating model determines which layers are genuinely justified and where the highest-value investments lie:

- The AI-Ready Data Foundation (Layer 4) is the correct first investment for all operating model types. Governance, data quality, catalogue, and lineage are prerequisites regardless of operating model position.
- The Data Products layer (Layer 3) should be built at the enterprise level in Coordination and Unification models, and at the business unit level in Diversification models. Replication models build standardized data products centrally and deploy them locally.
- The Semantic Layer (Layer 2) is the highest-priority enterprise investment in Coordination models, where shared entity definitions and knowledge graphs deliver the most distinctive value. In Unification models, it is equally critical but more straightforwardly achievable. In Diversification and Replication models, it may be built at the business unit level.
- The AI Layer (Layer 1) follows the data and semantic layers — in all operating model types. Investing in AI capability before the foundational layers are in place is the most common source of AI program failure.

4.4 Build the Workforce Strategy Into the Architecture Program

Architecture programs that treat workforce transformation as a change management workstream — a downstream consequence of technology deployment — consistently underperform those that treat it as a co-equal design constraint. Research consistently indicates that the human dimension of AI transformation accounts for 70% of the value created or destroyed.

For large international organizations, the workforce strategy must address three distinct layers:

- AI-fluent generalists: The broad workforce, trained in AI literacy, prompt-based interaction with AI tools, and critical evaluation of AI outputs. This is a minimum standard for all employees regardless of operating model type.
- AI domain specialists: Embedded within business units, combining deep domain expertise with applied AI skills. These individuals are the primary

translators between business need and AI capability. In Coordination models, they must also understand the shared data and semantic layer.

- AI platform engineers and architects: The central team responsible for infrastructure, security, model governance, and the AI-First architecture stack. In Unification and Replication models, this is a large and strategically critical function. In Diversification models, it may not exist at the enterprise level at all.

4.5 Plan for Operating Model Migration

As discussed in Section 3.4, AI investment can shift the economics of operating model choices in ways that justify deliberate migration from one quadrant to another. The CEA should build a regular operating model review into the enterprise architecture governance calendar — at minimum annually, and after any major AI capability deployment.

The five strategic decisions that define the AI TOM should be reviewed and, where necessary, updated:

Strategic Decision	Options	Recommended for Large International Orgs
Governance Model	Centralized / Federated / Hybrid	Hybrid: central standards, local execution
AI Platform Ownership	IT-led / Business-led / Shared	Shared CoE holding platform and policy
Data Architecture	Siloed / Unified / Federated Products	Federated data products with enterprise catalogue
Workforce Approach	Educate only / Redesign roles	Redesign roles and workflows; upskill broadly
Delivery Model	Big-bang / Phased / Continuous	Phased, outcome-driven, with fast early wins

4.6 Adopt a Phased Delivery Approach Calibrated to Operating Model

The pace and scope of AI TOM delivery should be calibrated to the operating model. A single enterprise-wide transformation program is appropriate only in the Unification model, where a single architecture decision applies everywhere. In all other models, phased delivery with clear horizon boundaries is the appropriate approach:

- Horizon 1 (0–6 months): Establish AI governance structures aligned to the operating model. Complete a data infrastructure baseline assessment. Identify and prioritize three to five high-confidence AI use cases appropriate to the operating model type. Deliver measurable quick wins that build internal credibility and demonstrate the link between architecture quality and AI output quality.

- Horizon 2 (6–18 months): Expand AI adoption across priority domains. Build the shared infrastructure appropriate to the operating model (semantic layer for Coordination, Feature Store for Replication, full stack for Unification). Embed AI domain specialists in business units. Begin workflow redesign in high-value process areas.
- Horizon 3 (18 months+): Evolve toward agentic AI deployments in domains where the foundational layers are mature. Continuously review the operating model position and update the AI TOM accordingly. Mature governance from oversight to embedded real-time controls.

5. Conclusion

The emergence of AI as a core operational capability — not a peripheral analytical tool — demands a corresponding evolution in how enterprise architects think about operating model design. The Ross, Weill, and Robertson framework, developed in the context of ERP and enterprise systems strategy, turns out to be equally powerful as a lens for AI strategy. The two diagnostic questions at its heart — about integration and standardization requirements — map directly onto the two most consequential decisions in AI architecture: where data must be shared and governed, and where AI models can and should be standardized.

For the Chief Enterprise Architect of a large international organization, the practical implications are clear. Operating model clarity must precede AI architecture decisions. The AI governance model must be designed to match the operating model, not applied as a generic framework. The AI-First architecture stack should be invested in selectively, with the highest-priority layers determined by the operating model quadrant. The workforce strategy must be treated as a co-equal design constraint. And the operating model itself should be treated as a strategic variable — one that AI investment may, over time, justify revising.

MIT CISR's ongoing research program on enterprise AI maturity, now explicitly examining how the IT operating model shapes AI value capture, is expected to produce empirical findings over the next 12 to 18 months that will further substantiate the framework presented here. The Chief Enterprise Architect who engages with this emerging body of research — and who positions the operating model conversation at the heart of the organization’s AI strategy — will be among the architects best placed to deliver AI transformation that is coherent, governed, and durable.

> *Architecture is not what you build — it is the logic by which you decide what to build. In the age of AI, that logic begins with the operating model.*

References and Further Reading

Ross, J.W., Weill, P., and Robertson, D.C. (2006). Enterprise Architecture as Strategy: Creating a Foundation for Business Execution. Harvard Business Review Press.

MIT Center for Information Systems Research (2025). Enterprise IT Operating Models in the AI Era. MIT CISR Research Briefing XXV-12, December 2025. Thorogood, A. and Woerner, S.L.

MIT Center for Information Systems Research (2025). Business Models in the AI Era. MIT CISR Research Briefing XXV-10, October 2025. Weill, P., Sebastian, I., Woerner, S.L., and Benedict, T.

MIT Center for Information Systems Research (2025). Grow Enterprise AI Maturity for Bottom-Line Impact. MIT CISR Research Briefing XXV-8, August 2025. Woerner, S.L. et al.

MIT Center for Information Systems Research (2026). The Enterprise IT Operating Model's Role in Capturing AI Value. Active Research Program.

McKinsey & Company (2025). The Agentic Organization: Five Pillars of AI-Driven Transformation.

BCG (2025). Future-Built: AI Workforce Transformation in High-Performing Enterprises.

Deloitte (2026). State of AI in the Enterprise.

PwC (2025). Responsible AI in Practice: From Principles to Operating Model.

Gartner (2025). Enterprise AI Governance: From Periodic Oversight to Embedded Real-Time Controls.