Digital Enterprise Architecture Advisors

WHITE PAPER

AI in Enterprise Architecture

Integrating Modern AI into Capabilities, Governance, and Target Architectures

Prepared by:

Christian Kobsa

Strategic Enterprise Architect

Digital Enterprise Architecture Advisors LLC

Publication Date:

February 2026

EXECUTIVE SUMMARY

Artificial Intelligence is reshaping how enterprises operate, make decisions, and deliver value. Yet most organizations struggle to scale AI beyond isolated pilots. The issue is not the technology — it is the absence of a coherent enterprise- level approach.

Enterprise Architecture (EA) is uniquely positioned to turn AI into a governed, scalable, enterprise-wide capability. This white paper provides:

- A capability-driven approach to AI adoption
- Guidance for integrating AI across all EA domains
- A comprehensive AI governance model
- A full EA-grade AI Reference Architecture
- Practical steps for embedding AI into the enterprise
- operating model

The goal is simple: help organizations move beyond experimentation and build AI into the core of how they operate.

Table of Contents

1. Introduction: Why AI Belongs at the Center of Enterprise Architecture

AI is no longer a future trend — it is a structural shift in how enterprises create value. From decision automation to knowledge retrieval to customer engagement, AI is reshaping business models and operating models alike.

However, most enterprises face the same challenges:

- AI initiatives are fragmented
- Data readiness is inconsistent
- Governance is unclear
- Technology choices are tactical
- Business alignment is weak

Enterprise Architecture provides the frameworks, governance, and strategic alignment needed to address these challenges. EA ensures AI is not a collection of tools, but a **coherent enterprise capability**.

2. Anchor AI in Business Capabilities

AI should begin with **capabilities**, not technologies. Capabilities are stable, business-aligned constructs that transcend organizational structures and systems. They provide the ideal foundation for AI adoption.

Why capabilities matter

Mapping AI to capabilities ensures:

- Clear alignment with strategic priorities
- Identification of high-value use cases
- Visibility into dependencies (data, skills, governance)
- A structured AI investment portfolio
- A stable foundation for long-term planning

Where AI enhances capabilities

AI strengthens capabilities in several predictable ways:

- **Decision support** — augmenting human judgment with predictive or generative insights
- **Forecasting & planning** — improving accuracy and speed of planning cycles
- **Customer interaction** — enabling personalized, conversational, context-aware engagement
- **Automation & optimization** — reducing manual effort and improving throughput
- **Knowledge management** — extracting, organizing, and retrieving institutional knowledge

This capability-driven approach prevents "AI for AI's sake" and ensures investments are tied to measurable outcomes.

3. Integrate AI Across All EA Domains

AI is not a standalone component. It reshapes every architectural layer — business, information, application, and technology. EA ensures these changes are coherent, governed, and strategically aligned.

3.1 Business Architecture

AI changes how work gets done. It introduces new decision points, new automation opportunities, and new human-machine interactions.

Key impacts

- AI-augmented processes (e.g., automated triage, predictive maintenance)
- New roles:
    - AI Product Owner
    - Model Steward
    - Prompt Engineer
    - AI Risk Officer

- Updated decision governance and RACI models

AI becomes part of the operating model, not an add-on.

3.2 Information Architecture

AI is only as good as the data it consumes. This layer determines whether AI initiatives scale or stall.

Key elements

- Data lineage, quality, metadata, and access controls
- Vector databases and embeddings for semantic search and RAG
- Policies for training data, inference data, and model outputs
- AI knowledge assets cataloged as enterprise information

Information Architecture becomes the backbone of AI governance and reliability.

3.3 Application Architecture

AI should be exposed as **standardized, composable services** that applications can consume.

Key elements

- Modular AI services (LLM APIs, inference endpoints, RAG orchestrators)
- Integration patterns: synchronous scoring, event-driven inference, batch pipelines
- Replaceable model components and providers
- AI governance integrated into SDLC and architecture review processes

This ensures AI is not tightly coupled to any single application or vendor.

3.4 Technology Architecture

Modern AI requires specialized infrastructure and operational capabilities.

Key elements

- GPU/accelerator strategy and model hosting patterns

- MLOps/LLMOps pipelines for training, deployment, and monitoring
- Observability for drift, bias, hallucination, and cost
- Security controls for prompts, data, and model access

Technology Architecture ensures AI is scalable, reliable, and cost-controlled.

4. Extend EA Governance to Cover AI

AI introduces new risks — ethical, operational, regulatory, and reputational. EA must extend governance frameworks to address these risks consistently across the enterprise.

AI governance elements

- Model lifecycle governance (approvals, versioning, retirement)
- Responsible AI controls (fairness, transparency, explainability)
- Risk classification for AI use cases
- Human-in-the-loop requirements for critical decisions
- Security guardrails (prompt injection, data leakage prevention)

EA becomes the steward of enterprise-wide AI governance.

5. EA-Grade AI Reference Architecture

This reference architecture provides a vendor-neutral, composable blueprint for designing, governing, and operating AI as an enterprise capability.

5.1 Core Principles

- **Business-anchored** — AI enhances capabilities and decisions
- **Composable** — models, vector stores, and orchestrators are replaceable
- **Governed** — risk, ethics, and compliance are built-in
- **Observable** — AI systems monitored like mission-critical services
- **Secure** — prompts, data, and outputs treated as sensitive surfaces

6. Layered Architecture

Each layer builds on the previous one, creating a coherent end-to-end AI ecosystem.

6.1 Business & Use-Case Layer

Components

- AI use-case portfolio and capability mapping
- Value and risk classification
- AI product ownership

Narrative

This layer ensures AI is tied to strategy and value. Without it, AI becomes a scattershot of disconnected experiments.

6.2 Data & Knowledge Layer

Components

- Source systems and data pipelines
- Feature store for ML
- Document/content store
- Vector store for embeddings and semantic retrieval
- Metadata, lineage, and sensitivity classification

Narrative

This layer transforms raw data into AI-ready knowledge. Vector stores and embeddings enable RAG, enterprise search, and knowledge copilots — the most impactful AI use cases today.

6.3 Model Lifecycle Layer

Components

- Model development environments
- Model registry with versioning and lineage

- Training pipelines
- Evaluation and testing (performance, fairness, robustness)
- Promotion workflows

Narrative

This is the AI factory. It ensures models are built, validated, approved, and deployed in a controlled manner.

6.4 Inference & Orchestration Layer

Components

- Real-time and batch inference services
- RAG orchestration (retrieval → context → LLM → post-processing)
- Tool/function calling
- Guardrails and policy enforcement
- Caching, cost optimization, and model routing

Narrative

This layer operationalizes AI. It standardizes how applications consume AI and ensures outputs are safe, compliant, and auditable.

6.5 Integration & Experience Layer

Components

- Enterprise applications consuming AI services
- Chat interfaces, copilots, and embedded AI
- Workflow/BPM integration
- Decisioning integration (rules engines, human review)

Narrative

AI must be embedded into workflows, not bolted on. This layer ensures AI enhances user experience and operational efficiency.

6.6 Platform & Infrastructure Layer

Components

- Compute and storage (GPU/accelerator pools)
- Cloud/on-prem/hybrid hosting strategy
- MLOps/LLMOps platform
- Observability stack
- Security services (IAM, secrets, KMS)

Narrative

This layer ensures AI workloads are scalable, performant, and secure. It also manages cost — a critical concern with LLMs.

6.7 Governance, Risk & Compliance Layer

Components

- AI policy framework
- Model risk management
- Responsible AI practices
- Security and privacy controls
- Audit and traceability

Narrative

AI governance is not optional. It protects the organization from regulatory, ethical, and operational risks.

7. End-to-End Flows

7.1 Model Lifecycle Flow

- Use-case definition
- Data preparation

- Model development
- Governance review
- Deployment
- Monitoring and retraining

7.2 RAG-Based Inference Flow

- User request
- Pre-processing
- Retrieval
- Context assembly
- LLM invocation
- Post-processing
- Logging and audit

These flows illustrate how the architecture operates in practice.

8. Mapping to EA Viewpoints

Business Architecture

- AI-augmented processes
- Capability uplift
- New roles and decision governance

Information Architecture

- Data domains
- Vector stores
- Metadata and classification

Application Architecture

- AI services

- Orchestrators
- Integration patterns

Technology Architecture

- Platforms
- Infrastructure
- Observability and security

Security & Risk Architecture

- AI-specific threats
- Controls
- Monitoring

9. AI Capability Map

- AI strategy & portfolio management
- Data & knowledge management
- Model engineering & lifecycle
- AI runtime & orchestration
- AI integration & experience
- AI governance & risk
- AI operations & observability

This capability map becomes the basis for maturity assessments and roadmaps.

10. Embedding AI into the Enterprise Operating Model

AI becomes sustainable only when embedded into the enterprise operating model.

Key integration points

- Portfolio management and prioritization
- Investment planning and sequencing

- Risk management and compliance
- Architecture review boards
- Change management and adoption
- Vendor and platform strategy

This ensures AI becomes a **first-class citizen** in enterprise decision-making.

Conclusion

AI is not just another technology wave — it is a structural shift in how enterprises operate. Enterprise Architecture is the discipline best positioned to guide this transformation.

By combining:

- A capability-driven approach
- Cross-domain architectural integration
- Strong governance
- A robust AI reference architecture

...organizations can move beyond experimentation and build AI into the core of their operating model.