Business Reference Architecture for the AI-Enabled and Agentic Enterprise

A TOGAF 10-Aligned Framework for Capabilities, Value Streams, Operating Models, and Governance

A White Paper – April 2026

Author:

Christian Kobsa
Strategic Enterprise Architect, Business Architect, IT Consultant
Digital Enterprise Architecture & Advisory (DEAA)

Abstract

Enterprises are shifting from application-centric architectures to **agentic AI ecosystems** that operate across hybrid-cloud environments. This white paper presents a complete reference architecture for securely deploying AI agents at enterprise scale, integrating IBM's agentic-AI vision, the Model Context Protocol (MCP), modern data engineering practices, and the Agent Development Lifecycle (ADLC). It outlines the architectural principles required for enterprise-grade AI—acceptable agency, secure-by-design patterns, interoperability, hybrid deployment, evaluation-first development, and continuous governance—and provides a practical roadmap for implementing AI systems across public cloud, private cloud, on-premises, and edge environments.

Table of Contents

Executive summary

Artificial intelligence has evolved from a discrete technology investment into a pervasive business capability. The next frontier—**agentic AI**—introduces autonomous, goal-driven systems that can reason, plan, and act with varying degrees of independence. These systems can orchestrate complex workflows, collaborate with other systems and humans, and continuously improve based on feedback.

To harness this power responsibly, enterprises need a **TOGAF 10-aligned Business Reference Architecture** (**BRA**) that:

- Treats AI as a **cross-cutting enabler**, not an isolated technology stack.
- Integrates **agentic AI** into business capabilities, value streams, and operating models.
- Embeds **governance**, **ethics**, and **risk management** into the AI lifecycle.
- Leverages **process-orchestration patterns** (e.g., Camunda-style hybrid BPMN + agents) for traceable, resilient automation.

This white paper provides:

- A **TOGAF-aligned Business Reference Architecture** for AI-enabled and agentic enterprises.
- A **Business Capability Map**, **Value Stream Model**, **Operating Model**, and **Governance Model**—each as a standalone diagram.
- A **deep integration of agentic-AI design patterns** and **process-orchestration insights** into the **TOGAF 10 ADM**.
- A **transformation roadmap** that balances innovation, control, and continuous learning.

1. Introduction

1.1 Purpose

This document defines a **Business Reference Architecture** (**BRA**) that integrates:

- Traditional business capabilities and value streams.
- AI-enabled capabilities (analytics, ML, automation).
- **Agentic AI**—autonomous, goal-driven agents that reason, plan, and act.

It is aligned with **TOGAF 10** and is intended for:

- Enterprise and Business Architects
- CIOs, CDOs, CAIOs
- Strategy, Transformation, and Operations Leaders
- Risk, Compliance, and Governance Executives

1.2 Why a Business Reference Architecture for AI and agentic AI?

AI introduces new capabilities, roles, and governance needs. Agentic AI amplifies this by:

- **Autonomous decision-making** within defined constraints.
- **Dynamic orchestration** of processes and services.
- **Continuous learning** from outcomes and feedback.

However, AI and agentic AI do **not** replace:

- Core business capabilities and human judgment.
- Regulatory and ethical obligations.
- Enterprise-wide governance and risk management.

A balanced BRA ensures AI and agents are **embedded**, not bolted on, and that autonomy is **bound by governance**.

2. TOGAF foundations for the AI-enabled and agentic enterprise

2.1 Alignment with TOGAF 10 business architecture

The BRA follows TOGAF's Business Architecture domains:

- **Business Strategy**
- **Business Capabilities**
- **Value Streams**
- **Information Concepts**
- **Organization & Roles**
- **Governance & Controls**

AI and agentic AI are integrated as **cross-cutting concerns** across all domains.

2.2 Architecture principles

- **AI as an enabler, not a goal:** AI and agents serve business outcomes.
- **Human-centered decision authority:** Humans retain ultimate accountability.
- **Transparency and explainability:** Decisions by agents must be traceable and explainable.
- **Modularity and reusability:** Agents, models, and workflows are composable building blocks.
- **Data as a strategic asset:** Data quality, lineage, and semantics are first-class concerns.
- **Governance by design:** Ethics, risk, and compliance are embedded from the outset.
- **Interoperability:** Agents interact with systems, APIs, and humans via well-defined contracts.
- **Security and ethical compliance:** Privacy, security, and fairness are non-negotiable.

3. Business strategy layer

3.1 Strategic drivers

- **Operational efficiency** through intelligent automation and agentic orchestration.
- **Enhanced customer experience** via personalization and proactive engagement.
- **Risk reduction** through predictive analytics and autonomous monitoring.
- **Revenue growth** via new AI-enabled products and services.
- **Workforce augmentation** with AI copilots and agents.
- **Faster decision cycles** through real-time insights and autonomous actions.

3.2 Strategic outcomes

- **Adaptive operating models** that respond dynamically to context.
- **AI-augmented workforce productivity** across roles and functions.
- **Data-driven decisioning at scale**, with agents acting as decision co-workers.
- **Continuous optimization** of processes and value streams.
- **Responsible and compliant AI deployment** with clear accountability.

3.3 Strategic risks

- **Model bias** and **unfair outcomes**.
- **Data quality degradation** and **drift**.
- **Regulatory non-compliance** (AI acts outside policy).
- **Model and agent drift** (behavior diverges from intent).
- **Over-automation** and **loss of human oversight**.
- **Security vulnerabilities** in agentic orchestration (e.g., prompt injection, API abuse).

4. Business capability reference model

This section defines the **capability map** that integrates AI-specific and traditional capabilities, and highlights where **agentic AI** extends them.

4.1 Standalone diagram: AI-enabled business capability map

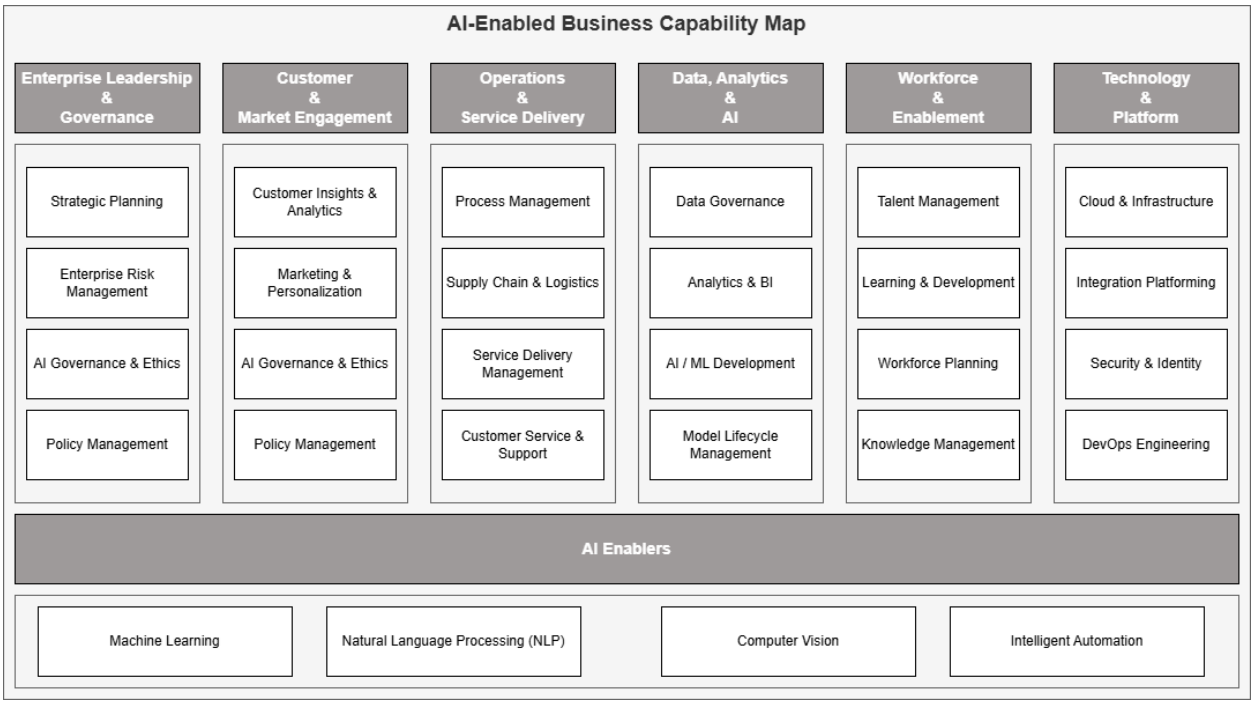

Columns and capabilities:

1. **Enterprise Leadership & Governance**
   - Strategic Planning
   - Enterprise Risk Management

- AI Governance & Ethics
- Policy Management
- Regulatory Compliance

2. **Customer & Market Engagement**
   - Customer Insight & Analytics
   - Marketing & Personalization
   - Sales Enablement
   - Customer Service & Support
3. **Operations & Service Delivery**
   - Process Management
   - Supply Chain & Logistics
   - Service Delivery Management
   - Quality Assurance
4. **Data, Analytics & AI**
   - Data Governance
   - Analytics & BI
   - AI / ML Development
   - Model Lifecycle Management
5. **Workforce & Enablement**
   - Talent Management
   - Learning & Development
   - Workforce Planning
   - Knowledge Management
6. **Technology & Platform**
   - Cloud & Infrastructure
   - Integration Platforms
   - Security & Identity
   - DevOps & Engineering

Bottom band: AI Enablers

- Machine Learning
- Natural Language Processing
- Computer Vision
- Intelligent Automation

Agentic AI extension (conceptual overlay):

On top of this map, agentic AI introduces cross-cutting capabilities such as:

- **Agent Orchestration & Coordination** (within Data, Analytics & AI + Technology & Platform).
- **Agent Policy & Guardrails Management** (within Enterprise Leadership & Governance).
- **Agent-Enabled Customer Interaction** (within Customer & Market Engagement).
- **Agent-Driven Process Optimization** (within Operations & Service Delivery).

5. Value stream reference model

AI and agentic AI enhance value streams but do not replace them.

5.1 Standalone diagram: AI-enhanced value stream model

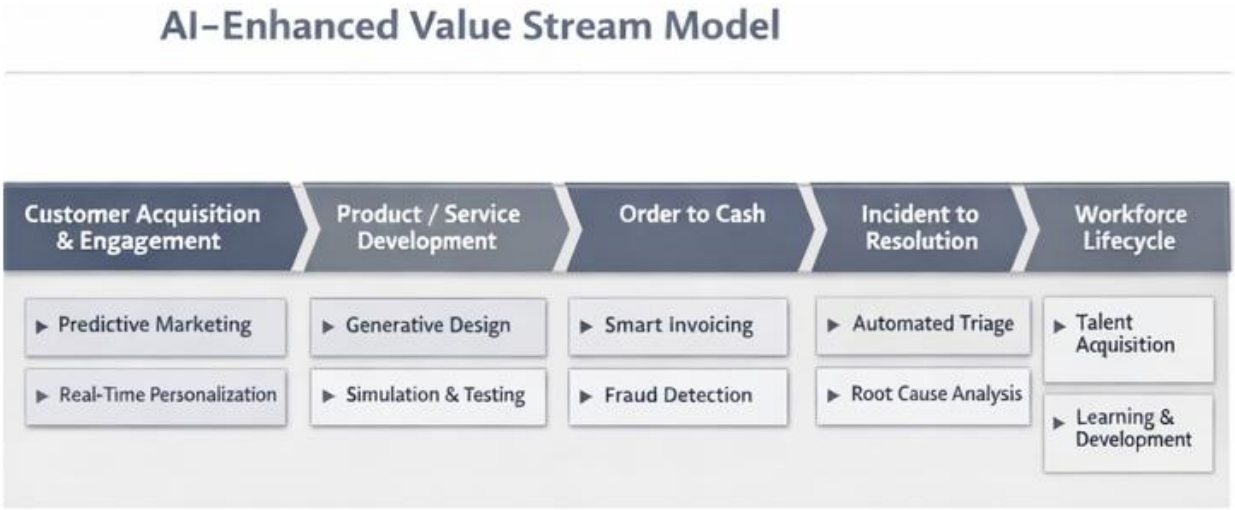

- Title: AI-Enhanced Value Stream Model
- Layout: Five horizontal columns (value streams) in a row; beneath each, two AI accelerators.

1. **Customer Acquisition & Engagement**
   a. Predictive Marketing
   b. Real-Time Personalization
2. **Product / Service Development**
   c. Generative Design
   d. Simulation & Testing
3. **Order to Cash**

e. Smart Invoicing
f. Fraud Detection

4. **Incident to Resolution**
   g. Automated Triage
   h. Root Cause Analysis
5. **Workforce Lifecycle**
   i. Talent Acquisition
   j. Learning & Development

Agentic AI extension:

- **Customer Acquisition & Engagement:**

  Agentic marketing agents autonomously test campaigns, adjust targeting, and orchestrate multi-channel engagement based on performance signals.

- **Product / Service Development:**

  Design agents generate options, run simulations, and propose trade-offs, collaborating with human product owners.

- **Order to Cash:**

  Finance agents monitor payment behavior, trigger running workflows, and coordinate with fraud agents.

- **Incident to Resolution:**

  Support agents triage tickets, gather context from multiple systems, and orchestrate resolution steps across teams and tools.

- **Workforce Lifecycle:**

  Talent agents match candidates to roles, propose learning paths, and monitor skill gaps.

6. Information concepts & knowledge architecture

Agentic AI requires richer information structures than traditional analytics.

6.1 Core information domains

- Customer
- Product / Service
- Operations
- Finance
- Workforce
- Risk & Compliance
- AI / Model Metadata

6.2 AI-specific information concepts

- **Training data lineage** (source, transformations, usage).
- **Model versioning and metadata** (hyperparameters, training context).
- **Model performance metrics** (accuracy, fairness, drift indicators).
- **Model risk classification** (criticality, impact, regulatory exposure).
- **Explainability artifacts** (feature importance, decision traces).
- **Prompt libraries and prompt governance** (for LLM-based agents).

6.3 Agentic AI-specific information concepts

- **Agent goals and policies** (what agents are allowed and not allowed to do).
- **Agent plans and execution traces** (steps taken, decisions made).
- **Agent-to-agent and agent-to-system interaction logs.**
- **Context graphs / knowledge graphs** used by agents for reasoning.
- **Vector embeddings** for semantic search and retrieval-augmented generation.

7. Organizational structure & roles

7.1 Standalone diagram: AI operating model

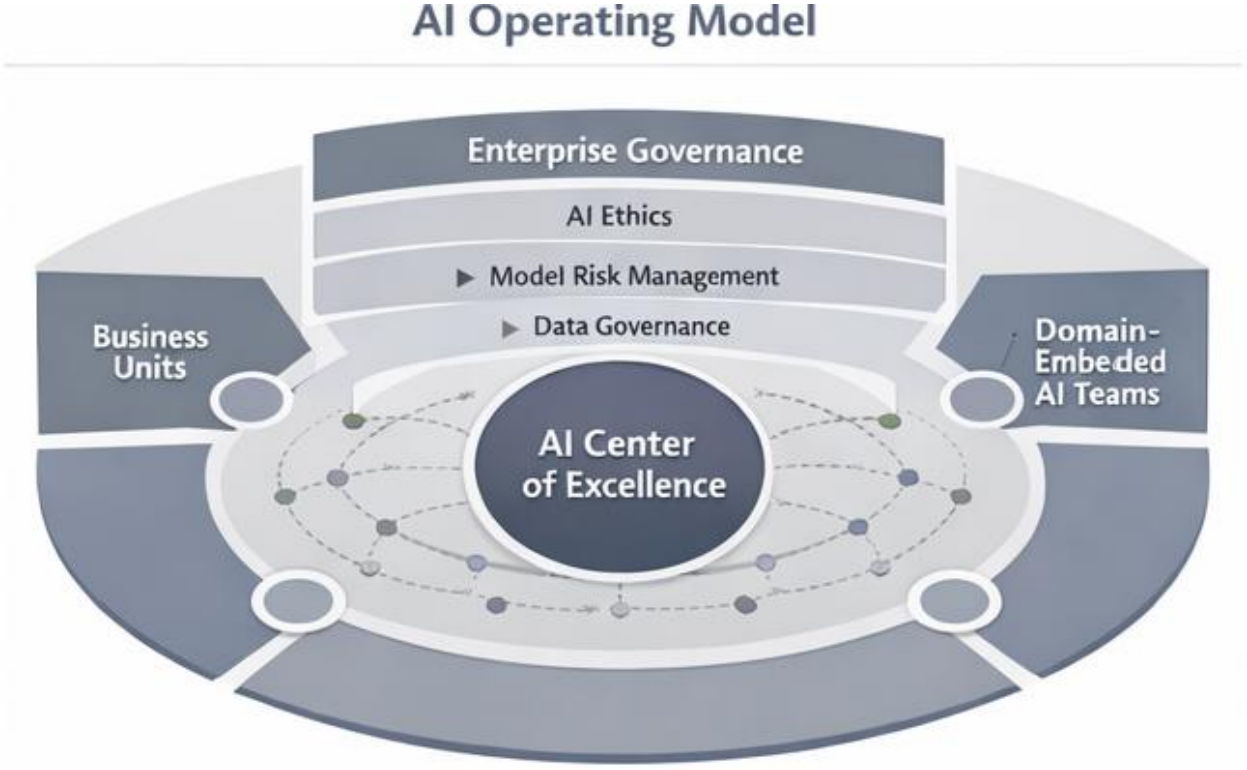

Diagram description:

Layout:

- Top horizontal bar: Enterprise Governance with three labels underneath:
  - AI Ethics
  - Model Risk Management
  - Data Governance
- Center circle: AI Center of Excellence
- Left segment: Business Units (connected to the center circle with arrows).
- Right segment: Domain-Embedded AI Teams (also connected to the center circle).

This is a **hub-and-spoke model:**

- **Hub:** AI Center of Excellence (CoE)
  - Defines standards, patterns, and reusable components.
  - Provides shared platforms (MLOps, orchestration, monitoring).

        - Supports complex or cross-domain AI initiatives.
- **Spokes:** Domain-Embedded AI Teams
    - Build and operate domain-specific models and agents.
    - Collaborate with business units on use cases.
    - Apply CoE standards and patterns.
- **Enterprise Governance:**
    - AI Ethics Board
    - Model Risk Management
    - Data Governance Council

7.2 AI-specific and agentic roles

- Chief AI Officer / Head of AI
- AI Product Owner
- AI Architect / Agentic AI Architect
- Data Scientist / ML Engineer
- MLOps Engineer
- Model Risk Officer
- AI Ethics Lead
- Prompt Engineer / Conversation Designer
- Process Orchestration Architect (BPMN + agents)

8. Governance & controls

8.1 Standalone diagram: AI governance & controls

Diagram description:

Layout:

Four horizontal layers stacked vertically.

1. **Top layer:** Enterprise Oversight
   a. Label: Executive Accountability
2. **Second layer:** Three parallel bars
   b. Model Risk Management
   c. AI Ethics Board
   d. Security & Compliance
3. **Third layer:** AI Lifecycle Management
   e. Sub-labels beneath:
      i. Model Development
      ii. Validation & Monitoring
      iii. Drift & Decommissioning
4. **Bottom layer:** Data Governance & Quality

8.2 Governance domains

- **AI Ethics:** fairness, transparency, human oversight.
- **Model Risk Management:** classification, validation, stress testing, approvals.
- **Security & Compliance:** privacy, access control, regulatory adherence.
- **AI Lifecycle Management:** from ideation to retirement.
- **Data Governance & Quality:** lineage, stewardship, quality controls.

8.3 Agentic AI governance specifics

- **Agent policy frameworks:** define allowed actions, escalation rules, and boundaries.
- **Autonomy levels:** classify agents by degree of autonomy (assistive, semi-autonomous, fully autonomous with human override).
- **Kill switches and override mechanisms:** ensure humans can intervene.
- **Behavior monitoring:** detect anomalous or policy-violating agent behavior.
- **Auditability:** maintain detailed logs of agent decisions and actions.

9. Agentic-AI design patterns and Camunda-style process orchestration

This section elaborates the **agentic-AI design patterns** and **process-orchestration insights** and sets up their integration into TOGAF ADM.

9.1 Core agentic-AI design patterns

1. **Task-Oriented Agent Pattern**

- Agents are assigned specific tasks (e.g., “triage incident,” “generate proposal”).
- They use tools (APIs, databases, BPMN processes) to complete tasks.
- They report back with results and explanations.

2. **Goal-Oriented Agent Pattern**

- Agents receive high-level goals (e.g., “reduce average resolution time by 20%”).
- They decompose goals into plans, execute steps, and adapt based on feedback.
- They coordinate with other agents and systems.

3. **Multi-Agent Collaboration Pattern**

- Specialized agents (e.g., data-retrieval agent, reasoning agent, execution agent) collaborate.
- A coordinator agent orchestrates their interactions.
- Useful for complex workflows (e.g., fraud investigations, claims handling).

4. **Human-in-the-Loop Pattern**

- Agents propose actions; humans approve, modify, or reject.
- Escalation rules define when human intervention is mandatory.
- Ideal for high-risk or regulated decisions.

5. **Guardrailed Agent Pattern**

- Agents operate within strict policy constraints.
- Policies are enforced via rule engines, BPMN gateways, or policy-as-code.
- Violations trigger alerts or automatic shutdown.

6. **Learning Loop Pattern**

- Agents log outcomes and feedback.
- Models and policies are updated periodically or continuously.
- Supports continuous improvement and adaptation.

9.2 Camunda-style hybrid orchestration patterns

Camunda's perspective on **agentic process orchestration** can be summarized as:

- **Hybrid BPMN + Agents:**
  - Use BPMN for **deterministic**, **auditable process structure**.
  - Use agents for **adaptive**, **context-sensitive decisions** within that structure.
  - Example: a BPMN process defines the steps of a claims process; an agent decides which documents to request, how to classify the claim, and when to escalate.
- **Agents as "smart tasks" in BPMN:**
  - Each agent is invoked as a service task or external task.
  - The process engine orchestrates sequence, parallelism, and compensation.
  - Agents handle unstructured work (interpretation, reasoning, content generation).
- **Event-Driven Orchestration:**
  - Agents subscribe to events (e.g., "payment failed," "ticket escalated").
  - BPMN processes are triggered or updated based on agent actions and events.
- **Case Management with Agents:**

  - For non-linear, variable workflows (e.g., investigations), agents help decide which tasks to open next.
  - The process engine maintains state and traceability.
- Governance and Observability:
  - The process engine provides a single pane of glass for monitoring agent-driven workflows.
  - Metrics: throughput, SLA adherence, error rates, agent decision patterns.

These patterns are crucial when embedding agentic AI into **mission-critical, regulated processes.**

10. TOGAF 10 ADM integration with agentic AI and orchestration

Now we integrate all of the above into a **phase-by-phase TOGAF 10 ADM view**, explicitly incorporating agentic-AI patterns and Camunda-style orchestration.

10.1 Phase: Preliminary

Objectives:

- Establish AI and agentic-AI **vision**, **principles**, and **policies**.
- Define **autonomy levels** and **governance structures**.
- Identify **reference patterns** (BPMN + agents, human-in-the-loop, guardrails).

Key activities:

- Define **AI and agentic-AI principles** (ethics, transparency, human oversight).
- Set up **AI Governance & Ethics Board** and **Model Risk Management** function.
- Select **process-orchestration platform** (e.g., BPMN engine) and **agent framework**.
- Define **reference architectures** for:
    - Agentic AI components (agents, tools, memory, policies).
    - Orchestration (BPMN, events, APIs).

Design patterns introduced:

- Guardrailed Agent Pattern
- Human-in-the-Loop Pattern

10.2 Phase A: Architecture vision

Objectives:

- Identify high-value AI and agentic-AI use cases.
- Define the **target vision** for AI-enabled and agentic value streams.

Key activities:

- Map **strategic drivers** to candidate use cases (e.g., agentic customer support, agentic claims handling).
- Use the **AI-Enhanced Value Stream Model** to identify where agents can create leverage.
- Define **success metrics** (e.g., reduced cycle time, improved NPS, reduced risk).

Design patterns applied:

- Goal-Oriented Agent Pattern (for strategic outcomes).
- Task-Oriented Agent Pattern (for specific use cases).

Orchestration considerations:

- Identify which processes will remain **fully deterministic** vs **hybrid** (BPMN + agents).
- Define **initial scope** for agentic orchestration pilots.

10.3 Phase B: Business architecture

Objectives:

- Design the **business architecture** with AI and agents embedded.
- Refine the **Business Capability Map** and **Value Streams** with agentic roles.

Key activities:

- Use the **AI-Enabled Business Capability Map** to:
  - Identify capabilities that will host agents (e.g., Customer Service & Support, Process Management).
  - Define new capabilities (e.g., Agent Orchestration & Coordination, Agent Policy Management).
- Model **business processes** using BPMN, explicitly marking:
  - Where agents are invoked as tasks.

    - Where human approvals are required.
    - Where guardrails and policies are enforced.

Design patterns applied:

- Multi-Agent Collaboration Pattern (for cross-capability workflows).
- Human-in-the-Loop Pattern (for high-risk decisions).

Orchestration considerations:

- Design **hybrid BPMN models** where:
    - The process engine orchestrates structure.
    - Agents handle unstructured decisions and content.
- Define **case management patterns** for variable workflows (e.g., investigations, escalations).

10.4 Phase C: Information systems architecture

Objectives:

- Design **data architecture** and **application architecture** to support agents and orchestration.

Key activities:

- Define **data pipelines** for training and operating models (batch + streaming).
- Introduce **vector databases** and **knowledge graphs** for semantic reasoning.
- Define **APIs and integration patterns** for:
    - Agents to call systems (tools).
    - BPMN engine to invoke agents.
- Model **information concepts** for agent policies, goals, and execution traces.

Design patterns applied:

- Learning Loop Pattern (data + feedback for continuous improvement).
- Guardrailed Agent Pattern (policy enforcement via rules and APIs).

Orchestration considerations:

- Ensure **idempotent**, **observable APIs** for agent tasks.
- Implement **event-driven architecture** for agent triggers and process updates.

10.5 Phase D: Technology architecture

Objectives:

- Define the **technical platforms** for AI, agents, and orchestration.

Key activities:

- Select **cloud / on-prem / hybrid** infrastructure for AI workloads.
- Define **MLOps platform** for model lifecycle management.
- Define **agent runtime** (LLM, tools, memory, policies).
- Define **process-orchestration platform** (BPMN engine, event bus).
- Implement **security controls** (identity, access, encryption, network segmentation).

Design patterns applied:

- Multi-Agent Collaboration Pattern (runtime support).
- Guardrailed Agent Pattern (policy enforcement at runtime).

Orchestration considerations:

- Ensure **high availability** and **resilience** for the BPMN engine and agent services.
- Implement **centralized logging and monitoring** for both processes and agents.

10.6 Phase E: Opportunities & solutions

Objectives:

- Identify **solution building blocks** and **transition architectures**.
- Prioritize **agentic AI initiatives**.

Key activities:

- Group use cases into **solution clusters** (e.g., Customer Experience, Operations Optimization).
- Define **pilot solutions** combining:
    - Agents (task-oriented, goal-oriented).
    - BPMN processes.
    - Data and model services.
- Evaluate **build vs buy** for agent frameworks and orchestration components.

Design patterns applied:

- Task-Oriented Agent Pattern (for pilots).
- Human-in-the-Loop Pattern (for early deployments).

Orchestration considerations:

- Use **Camunda-style hybrid orchestration** to ensure traceability and control in pilots.
- Instrument pilots with **metrics and observability** from day one.

10.7 Phase F: Migration planning

Objectives:

- Plan the **phased rollout** of agentic AI capabilities.
- Manage **organizational change** and **risk**.

Key activities:

- Define **migration waves** by domain or value stream.
- Plan **decommissioning** of legacy automation where agents take over.
- Define **training and upskilling** for business and IT teams.

Design patterns applied:

- Learning Loop Pattern (use pilot feedback to refine rollout).

Orchestration considerations:

- Ensure **backward compatibility** and **fallback paths** (e.g., manual handling if agents fail).
- Use BPMN to orchestrate **coexistence** of old and new flows.

10.8 Phase G: Implementation governance

Objectives:

- Ensure implementations conform to architecture, policies, and risk controls.

Key activities:

- Apply **AI Governance & Controls** model:
    - Model approvals, ethics reviews, risk assessments.
    - Security and compliance checks.
- Monitor **agent behavior** and **process performance**.
- Enforce **guardrails** and **override mechanisms**.

Design patterns applied:

- Guardrailed Agent Pattern (runtime enforcement).
- Human-in-the-Loop Pattern (for critical decisions).

Orchestration considerations:

- Use the BPMN engine as a **control plane** for agentic workflows.
- Implement **alerting and incident management** for agent failures or anomalies.

10.9 Phase H: Architecture change management

Objectives:

- Enable **continuous improvement** of AI, agents, and processes.
- Manage **drift** and **evolving requirements**.

Key activities:

- Monitor **model and agent drift**; trigger retraining or policy updates.
- Use **Learning Loop Pattern** to incorporate feedback into models and processes.
- Update **reference architectures** and **patterns** as new capabilities emerge.

Orchestration considerations:

- Continuously refine BPMN models based on observed behavior and metrics.
- Introduce new agents or retire old ones through controlled change processes.

11. Transformation roadmap

11.1 Horizon 1: Foundation

Focus:

- Governance, data readiness, and initial orchestration.

Key actions:

- Establish **AI governance**, ethics, and risk management.
- Implement **data governance** and foundational data pipelines.
- Deploy **process-orchestration platform** and define **reference BPMN patterns**.
- Launch **low-risk agentic pilots** with strong human-in-the-loop controls.

11.2 Horizon 2: Expansion

Focus:

- Scaling agentic AI across value streams.

Key actions:

- **Extend agents** into Customer Acquisition, Order-to-Cash, and Incident-to-Resolution.
- Deploy **MLOps** and **agent lifecycle management** at scale.
- Introduce **multi-agent collaboration** for complex workflows.
- Integrate **Camunda-style hybrid orchestration** into core processes.

11.3 Horizon 3: Optimization

Focus:

- Continuous optimization and adaptive enterprise.

Key actions:

- Implement **autonomous optimization loops** (agents adjusting processes based on KPIs).
- Expand **goal-oriented agents** for strategic objectives.
- Institutionalize **continuous learning** across models, agents, and processes.
- Evolve the BRA as a **living architecture**.

12. Conclusion

This white paper defines a **TOGAF 10-aligned Business Reference Architecture** for an **AI-enabled and agentic enterprise**, preserving the full original structure and explicitly embedding:

- A **Business Capability Map** with AI enablers.
- An **AI-Enhanced Value Stream Model**.
- An **AI Operating Model** (hub-and-spoke with CoE and domain teams).
- An **AI Governance & Controls** layered model.

It extends that foundation with:

- **Agentic-AI design patterns** (task-oriented, goal-oriented, multi-agent, human-in-the-loop, guardrailed, learning loops).
- **Camunda-style hybrid orchestration** (BPMN + agents, event-driven workflows, case management).
- A **phase-by-phase TOGAF 10 ADM integration** that makes agentic AI and orchestration first-class citizens in enterprise architecture.

References

Foundational Enterprise Architecture Sources

The Open Group.
TOGAF® Standard, 10th Edition.
The Open Group, 2022.

The Open Group.
TOGAF® Series Guides (Business Architecture, ADM Techniques, Architecture Content Framework, etc.).
The Open Group Library.

AI, Agentic AI, and Governance Foundations

ISO/IEC JTC 1/SC 42.
Artificial Intelligence Standards Portfolio (AI Governance, Risk Management, Trustworthiness).

NIST.
AI Risk Management Framework (AI RMF 1.0).
National Institute of Standards and Technology, 2023.

Microsoft.
Responsible AI Standard (v2).
Microsoft, 2022.

Google DeepMind, OpenAI, Anthropic.
Publicly available research on agentic AI patterns, tool-use agents, and multi-agent systems.

Process Orchestration & BPMN Foundations

Object Management Group (OMG).
Business Process Model and Notation (BPMN) 2.0 Specification.

Camunda.
Documentation and best-practice guides on BPMN, orchestration, and automation patterns.

Academic

- Russell, S., & Norvig, P. *Artificial Intelligence: A Modern Approach*.

- Davenport, T., & Kirby, J. *Beyond Automation: Strategies for Augmenting Work with AI*.
- Gartner Research on AI Operating Models, MLOps, and Autonomous Enterprise Patterns.