Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

From AI Assistants to AI Workforces: A Multi-Agent Enterprise Operating Model for Saudi Vision 2030 Organizations

Choudhry Bilal Mazhar Hussain[1]

[1]Arabian Agricultural Services Company (ARASCO), Riyadh, Kingdom of Saudi Arabia

Corresponding author: cbmr.bilal@gmail.com | ORCID: 0009-0005-7625-7296

Abstract

This paper develops a multi-agent enterprise artificial intelligence (AI) operating model for organizations seeking to move beyond isolated AI assistants toward coordinated AI workforces, positioned within Saudi Vision 2030 and the Kingdom's national data-and-AI strategy. Using design science research, the study specifies the Saudi Enterprise Multi-Agent AI Operating Model (SEMAI) as a conceptual artifact synthesized from enterprise AI, multi-agent systems, digital transformation, responsible AI, and information systems design science literature. The artifact is demonstrated and formatively evaluated through a structured scenario walkthrough based on a generalized Saudi enterprise context; the study does not claim empirical validation or measured deployment outcomes. SEMAI comprises five layers integrating human roles, AI assistants, collaborative agents, enterprise systems, and governance controls. It is accompanied by design requirements, design principles, an artifact specification, an evaluation rubric, a maturity model with progression criteria, governance and accountability controls, and a vendor-neutral model with a Microsoft-oriented reference implementation. The contribution is a reusable, governance-aware operating model for transitioning from task-level AI assistants to accountable, human-supervised, multi-agent AI workforces suitable for digitally mature Vision 2030 organizations.

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

Keywords: Multi-agent AI; Enterprise AI; AI workforce; AI governance; Saudi Vision 2030

1. Introduction

Organizations are entering a new phase of enterprise AI adoption. The first phase emphasized analytics dashboards, automation scripts, chatbots, and individual AI assistants. The emerging phase is more complex: enterprises are beginning to design coordinated AI workforces in which multiple specialized agents interact with human employees, enterprise systems, workflows, and governance controls. Microsoft documentation already reflects this direction: Copilot Studio supports custom AI agents that can understand natural language, answer questions, automate tasks, and connect to Fabric data agents as connected agents for agent-to-agent collaboration. Microsoft also documents the ability to connect Microsoft Foundry agents to Copilot Studio agents, while emphasizing testing, permissions, observability, identity, traceability, and human oversight when external agents are used [18,19].

Saudi Arabia is an appropriate context for examining this transition. Saudi Vision 2030 frames data, AI, digital transformation, innovation, and future-economy competitiveness as national priorities, while SDAIA's National Strategy for Data & AI describes a national ambition to position the Kingdom as a global hub for data and AI, build local data-and-AI talent, attract investment, strengthen research and innovation, and stimulate adoption through a collaborative ecosystem. The strategy's 2030 orientation is explicitly linked to becoming an "industry leader" utilizing and exporting data and AI. Therefore, enterprise AI workforces are not merely a technology trend; they are relevant to the Kingdom's economies-of-the-future agenda, advanced computing capability, workforce transformation, and knowledge-based economic diversification [28,35,36].

However, current enterprise AI adoption remains fragmented. Many organizations deploy copilots, chatbots, workflow automations, analytics assistants, and AI-enabled knowledge tools as separate capabilities. These tools may improve individual productivity, but they do not automatically become an operating model. A production-

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

grade AI workforce requires orchestration, shared context, secure access to enterprise data, integration with operational systems, accountability, model monitoring, human decision rights, and measurable business outcomes. The gap is not the absence of AI assistants; it is the lack of practical operating architecture for coordinating multiple agents across business functions.

This study addresses the following research questions:

RQ: How can Saudi Vision 2030 organizations design and govern multi-agent enterprise AI workforces that coordinate human employees, AI assistants, enterprise systems, and decision workflows while maintaining accountability, security, and business value?

The paper contributes four artifacts: (1) SEMAI, the Saudi Enterprise Multi-Agent AI Operating Model; (2) a vendor-neutral operating architecture with a Microsoft ecosystem reference implementation; (3) a five-level AI workforce maturity model with measurable progression criteria; and (4) a structured demonstration and formative evaluation scenario based on a representative Saudi enterprise. The paper intentionally avoids overlapping with prior work on dashboards, cybersecurity governance, or Fabric governance frameworks by focusing specifically on **multi-agent AI workforce design**, not analytics dashboards or cyber-control reporting.

2. Literature Review

2.1 Enterprise AI and Digital Transformation

Enterprise AI research has shifted from experimental machine-learning pilots toward enterprise-wide AI transformation. Earlier work framed AI as a tool for automation, prediction, and decision support; more recent scholarship positions AI as a general-purpose organizational capability requiring changes to data architecture, governance, skills, operating models, and decision rights [2,3,6,8,20,25,40]. This shift builds on a longer tradition of business intelligence and decision-support research linking data and analytics to organizational decision quality [5,24,31]. Digital transformation research similarly emphasizes that technology value emerges not from tool adoption alone but

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

from organizational redesign, capability building, and business-model adaptation [4,34,35].

For Saudi Vision 2030 organizations, this distinction is important. SDAIA's strategy does not describe AI only as a software capability; it links AI to workforce skills, policy and regulation, investment, research and innovation, and ecosystem development. This broader framing suggests that enterprise AI adoption should be studied as an operating model rather than as isolated technical deployment [28].

2.2 Construct Definitions for Enterprise Agentic AI

Table 1 defines the core constructs used throughout this study.

Table 1. *Core construct definitions for enterprise agentic AI.*

Construct	**Definition**	**Autonomy and workflow role**	**Accountability and governance implication**
AI assistant	User-facing AI capability that responds to prompts, summarizes content, retrieves knowledge, drafts text, or triggers simple tasks.	Low to moderate autonomy; supports individual productivity and task assistance.	Requires usage policy, access control, source grounding, and user review.
AI agent	Role-specific AI component with defined objectives, tool access, context, and workflow responsibilities.	Moderate autonomy; performs bounded actions within a defined business process.	Requires owner assignment, tool permissions, testing, monitoring, and audit logs.

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

Collaborative agent	Specialized agents are designed to exchange information, coordinate tasks, or hand off work with other agents.	Moderate autonomy; contributes to multi-step workflows across functions or systems.	Requires orchestration rules, escalation paths, conflict handling, and traceability.
Multi-agent system	Coordinated set of agents that share context, allocate tasks, interact with tools, and support distributed decision workflows.	Bounded collective autonomy; executes process-level orchestration under controls.	Requires governance spine, observability, incident response, model monitoring, and human approval gates.
AI workforce	Enterprise operating arrangement in which multiple agents and human roles jointly perform repeatable work across business functions.	Process-level autonomy; coordinates task execution, recommendations, and exception routing.	Requires business ownership, RACI, decision rights, performance metrics, and risk controls.
Enterprise AI operating model	Management system that defines how AI capabilities are governed, deployed, monitored, measured, and improved across the enterprise.	Organizational-level coordination; aligns AI workforces with strategy, architecture, compliance, and value realization.	Requires executive sponsorship, governance committees, lifecycle controls, and continuous assurance.

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

2.3 Agentic AI and AI Assistants

AI assistants are typically user-facing systems that respond to prompts, retrieve knowledge, summarize content, generate drafts, or trigger workflows. In Microsoft's ecosystem, Copilot Studio is described as a graphical low-code platform for building custom AI agents that can answer inquiries, automate tasks, and deploy across channels such as Microsoft Teams, websites, and Microsoft 365 Copilot. These capabilities support productivity, but an assistant remains limited when it operates as a single generalist [18].

Agentic AI extends this model by giving AI components role-specific objectives, tools, memory/context, and workflow responsibilities. In enterprise settings, agents may support HR onboarding, procurement triage, IT service management, finance variance explanation, or analytics interpretation. The literature on agentic systems suggests that specialization and coordination are central: one agent may classify intent, another may retrieve data, another may execute a workflow, and another may present a decision recommendation [15,27,30,32,38,39,41,42].

Recent literature further indicates that generative AI can enhance productivity when embedded into organizational work systems, but these benefits depend on governance, adoption, and workflow integration rather than model capability alone [40]. Emerging surveys of LLM-based multi-agent collaboration and agentic AI also emphasize the need for definitions, orchestration mechanisms, evaluation metrics, and risk controls, which support SEMAI's emphasis on structured artifact specification and formative evaluation [41,42].

2.4 Multi-Agent Systems

Multi-agent systems have a long research tradition in distributed artificial intelligence, coordination, negotiation, task allocation, autonomous systems, communication protocols, shared beliefs, distributed knowledge, and emergent behavior [7,10,30,32,38,39]. Their relevance to enterprises is increasing because modern organizations rarely operate through one process, one dataset, or one decision path. Enterprise work is distributed across departments, roles, systems, and policies.

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

SEMAI translates classic MAS concepts into enterprise operating principles: coordination becomes workflow orchestration; task allocation becomes agent specialization; communication becomes governed agent-to-agent protocols; autonomy becomes decision-rights constraints; emergent behavior becomes a requirement for observability and control; and distributed knowledge becomes governed enterprise data access.

Microsoft's current documentation provides a practical enterprise example. Fabric data agents can be consumed inside Copilot Studio as connected agents, enabling a custom Copilot Studio agent to securely access enterprise data through a Fabric data agent and ground responses in organizational knowledge. Microsoft Foundry agents can also be connected to Copilot Studio agents, but Microsoft explicitly states that organizations remain responsible for reviewing and testing connected agents, ensuring appropriate data flows, meeting requirements and laws, enforcing permissions, and establishing observability, identity, traceability, and human oversight. These requirements align closely with responsible multi-agent enterprise design [18,19].

2.5 AI Governance and Responsible AI

AI governance addresses the policies, controls, roles, and assurance mechanisms needed to ensure AI systems are reliable, secure, transparent, accountable, and aligned with organizational objectives [9,11,14,21,22,26,43]. In multi-agent systems, governance becomes more complex because decisions may emerge from agent collaboration rather than from a single model output. A procurement recommendation, for example, may combine outputs from a demand agent, supplier-risk agent, budget agent, and human approver.

Saudi organizations also operate in a national data-and-AI context shaped by SDAIA, NDMO, PDPL, and sector-specific cybersecurity and compliance expectations. The national AI strategy emphasizes practical and ethical data-and-AI applications, governance, infrastructure, and ecosystem development. Therefore, a Saudi enterprise AI workforce model should embed governance by design: role-based access, data classification, human decision rights, auditability, agent metadata, model

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

evaluation, and escalation rules [28]. These requirements build on established data-governance research [1,16,37] and emerging AI-management standards such as ISO/IEC 42001 [14].

Recent Saudi public-sector guidance also reinforces the need for ethical and responsible use of AI technologies, which strengthens the rationale for embedding responsible AI controls directly into SEMAI's governance layer [43].

2.6 Research Gap

Existing research provides foundations for enterprise AI, digital transformation, multi-agent systems, and AI governance. However, there are three gaps remaining. First, much of the literature treats AI assistants, AI governance, and enterprise architecture separately. Second, multi-agent systems research often emphasizes technical coordination but less often translates into business operating models for large enterprises. Third, Saudi Vision 2030 research has emphasized national AI strategy and digital transformation, but less attention has been given to how enterprises can operationalize AI agents as governed workforces.

This paper fills that gap by proposing a practical operating model that integrates enterprise AI assistants, multi-agent collaboration, Microsoft ecosystem architecture, governance controls, human decision rights, and Vision 2030 alignment.

3. Research Methodology

This study uses design science research methodology. Design science is suitable because the objective is not to test a statistical hypothesis but to create and communicate an artifact that addresses a practical organizational problem. Design science research is well established in information systems [12,13,17,23,33]. Peffers *et al.* describe design science research as a methodology for information systems research that includes problem identification, objectives of a solution, design and development, demonstration, evaluation, and communication [23]. This study follows that logic.

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

The problem identification stage defines the transition challenge from AI assistants to AI workforces. The objectives stage specifies the need for a model that is enterprise-grade, governance-aware, vendor-neutral at the principal level, Microsoft-compatible at the implementation level, and relevant to Saudi Vision 2030. The design stage develops SEMAI, design requirements, design principles, artifact specification, maturity model, and reference architecture. The demonstration stage applies the model to a representative enterprise onboarding scenario. The evaluation is formative and conceptual rather than empirical; it assesses fit against predefined criteria including completeness, governance coverage, human accountability, enterprise integration, security and privacy alignment, workflow orchestration, scalability, auditability, business-value measurability, and Vision 2030 alignment.

The demonstration scenario uses generalized, non-confidential observations representative of a large Saudi enterprise operating a combined Microsoft and SAP environment. Such an environment typically involves Microsoft 365, SAP, Power BI, Fabric modernization planning, AI and Copilot enablement, governing analytics, and data-foundation development, together with user-facing agents. These observations are used only to build a generalized demonstration scenario and structured formative evaluation; no confidential operational measures, personal data, credentials, system identifiers, or commercial commitments are disclosed.

4. Design Science Operationalization

Table 2 operationalizes the design science research methodology across its six phases.

Table 2. *Design science methodology operationalization.*

DSR phase	**Activity performed in this study**	**Output/artifact**	**Evidence used**	**Evaluation criterion**

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

Problem identification	Defined fragmentation between AI assistants, enterprise workflows, data platforms, and governance.	Problem statement and research question.	Enterprise AI, MAS, digital transformation, and governance literature.	Relevance to enterprise AI adoption and Vision 2030 priorities.
Objectives of solution	Specified requirements for governed multi-agent AI workforces.	Design requirements and design principles.	Design science, responsible AI, AI RMF, ISO/IEC 42001, SDAIA strategy, and Microsoft agent documentation.	Completeness, accountability, security, auditability, and business value.
Design and development	Constructed SEMAI layers, decision rights, maturity model, and reference architecture.	SEMAI artifact, architecture, maturity model, and governance tables.	Kernel theories from MAS, enterprise AI, data governance, and digital transformation.	Internal coherence, reusability, and operational specificity.
Demonstration	Applied SEMAI to employee onboarding	Scenario walkthrough	Generalized Saudi enterprise	Fit to realistic multi-system, multi-role

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

		and operational readiness.	and agent workflow.	context with Microsoft/SAP ecosystem assumptions.	enterprise constraints.
Formative evaluation	Assessed the scenario against explicit evaluation criteria and risk controls.	Evaluation of rubric, value metrics, and risk mitigation table.	Responsible AI controls, governance requirements, and operating-model criteria.	Coverage of agent roles, governance, integration, scalability, and measurable value.	
Communication		Presented the artifact, implications, limitations, and future research agenda.	Revised manuscript.	Design science reporting guidance.	Transparency, boundary conditions, and reproducibility of the artifact logic.

5. Multi-Agent Enterprise AI Operating Model

5.1 Design Requirements and Principles

Table 3 presents the SEMAI design requirements and their corresponding design principles.

Table 3. *SEMAI design requirements and principles.*

Requirement	**Design principle**	**Operational implication**
Coordinate multiple specialized agents	Agent specialization with governed orchestration.	Define agent roles, handoffs, task ownership, and conflict-resolution rules.

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

across business functions.		
Preserve human accountability for material decisions.	Human-in-the-loop and human-on-the-loop control.	Use approval gates, escalation paths, decision registers, and override procedures.
Ground agent outputs in trusted enterprise data.	Governed data access by design.	Use certified data sources, sensitivity labels, lineage, role-based access, and audit logs.
Prevent uncontrolled automation.	Autonomy proportional to decision risk.	Match autonomy level to decision type, confidence, impact, and policy requirements.
Enable measurable enterprise value.	Value measurement embedded in operations.	Track productivity, quality, risk, adoption, financial impact, and capability development.
Support transferability beyond one vendor stack.	Vendor-neutral model, platform-specific implementation.	Separate SEMAI principles from Microsoft reference components and document substitutes.

5.2 SEMAI: Saudi Enterprise Multi-Agent AI Operating Model

The proposed SEMAI framework consists of five layers, as illustrated in Figure 1.

Layer 1: Human workforce. This includes executives, managers, analysts, operational staff, process owners, data stewards, and compliance teams. Humans define goals, approve high-impact decisions, resolve exceptions, and remain accountable for business outcomes.

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

Layer 2: AI assistants. These are user-facing assistants such as Microsoft 365 Copilot, HR assistants, procurement assistants, IT service assistants, finance assistants, and analytics assistants. They support individual work but do not independently govern cross-functional processes.

Layer 3: Collaborative agents. This layer contains specialized agents that coordinate tasks. Examples include a policy agent, SAP data agent, analytics agent, supplier-risk agent, workflow agent, and approval-routing agent. Microsoft's connected-agent model is relevant because Copilot Studio can connect to Fabric data agents and Microsoft Foundry agents, enabling agent-to-agent collaboration and specialized task handling [18,19].

Layer 4: Enterprise systems and data platforms. This includes SAP S/4HANA, Microsoft Fabric, OneLake, Power Platform, SharePoint, Teams, Power BI, Dataverse, and APIs. A representative enterprise target flow moves data from SAP S/4HANA through Fabric Data Factory into OneLake and on to Power BI using Direct Lake.

Layer 5: Governance and trust. This layer covers identity, permissions, data classification, audit trails, responsible AI controls, human oversight, model evaluation, and legal/regulatory alignment. Microsoft's guidance for Foundry-connected agents explicitly highlights the need for appropriate data flows, permissions, observability, identity, traceability, and human oversight [19].

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

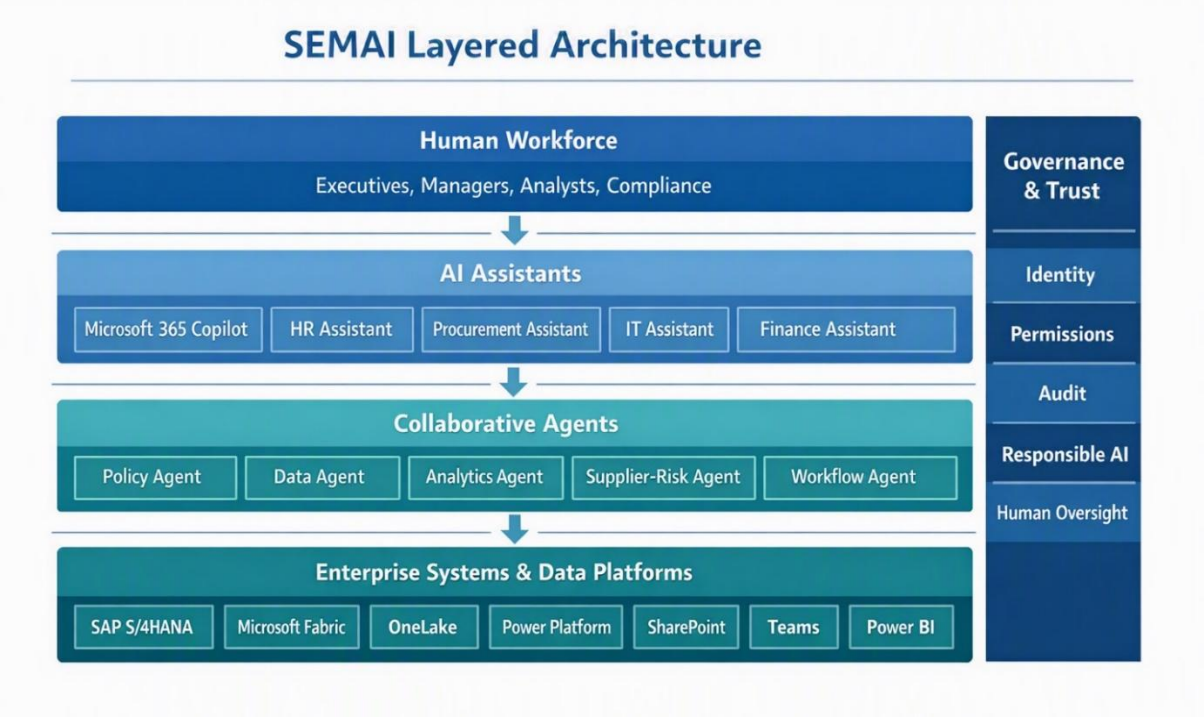

Figure 1. *SEMAI layered architecture, showing the human workforce, AI assistants, collaborative agents, and enterprise systems layers with a governance and trust spine spanning all layers.*

5.3 Formal Artifact Specification

Table 4 provides the formal artifact specification for each SEMAI layer.

Table 4. *Formal SEMAI artifact specification.*

Layer	**Primary actors**	**Core function**	**Dependencies**	**Control and risks**	**Success measures**
Human workforce	Executives, process owners, managers, analysts, data stewards,	Define goals, approve decisions, resolve exceptions,	Role definitions, governance committees, approval policies.	Risk of unclear ownership; mitigated by RACI and	Decision latency, override quality, accountable outcome ownership.

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

	compliance teams.	own outcomes.		decision registers.	
AI assistants	Copilots, departmental assistants, knowledge assistants.	Support individual productivity and guided task completion.	Knowledge bases, user permissions, collaboration platforms.	Risk of hallucination and misuse; mitigated by grounding and user review.	Adoption rate, user satisfaction, task completion time.
Collaborative agents	Policy, analytics, SAP/data, workflow, risk, and routing agents.	Coordinate specialized work across processes and systems.	Agent registry, APIs, orchestration logic, monitoring.	Risk of conflicting outputs; mitigated by arbitration rules and escalation.	Successful handoffs, exception rate, workflow completion quality.
Enterprise systems and data	SAP, data platform, BI, workflow, document, and collaboration systems.	Provide transactional, analytical, and operational context.	Data pipelines, APIs, metadata, lineage, master data, integration security.	Risk of data leakage or stale data; mitigated by classification and certified sources.	Data freshness, access compliance, certified source usage.
Governance and trust	AI governance board,	Set policies, assurance	AI policy, risk framework, model	Risk of unmanaged	Audit completeness, incident

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

	legal, cybersecurity, data office, internal audit.	controls, monitoring, auditability, and human oversight.	inventory, logs, incident process.	autonomy; mitigated by lifecycle controls and monitoring.	closure time, policy compliance, risk reduction.

5.4 Vendor-Neutral SEMAI and Microsoft Reference Implementation

SEMAI is vendor-neutral at the operating-model level. Its core principles—agent specialization, governed orchestration, trusted enterprise data, human accountability, auditability, and value measurement—can be implemented using different cloud, data, workflow, and AI platforms. The Microsoft-oriented architecture in this paper is therefore presented as one reference implementation, not as the only possible realization of SEMAI. In a Microsoft ecosystem, Copilot Studio can provide the orchestration and user-interaction layer; Fabric data agents can provide governed access to analytical data; Microsoft Foundry agents can support specialized reasoning, evaluation, and custom model orchestration; Power Platform can enable workflow automation and approvals; SAP can provide transactional business context; Fabric and OneLake can provide a governed data foundation; and SharePoint and Teams can provide collaboration surfaces. Equivalent implementations could substitute other orchestration frameworks, enterprise data platforms, identity providers, workflow engines, collaboration platforms, or ERP systems, if they preserve the same governance, security, traceability, and human-approval controls.

The Microsoft reference implementation provides operational specificity but also introduces technology-stack dependence. Potential risks include licensing uncertainty, tenant configuration complexity, connector limitations, data residency and privacy constraints, platform roadmap changes, agent interoperability gaps, and vendor lock-in. Organizations adopting SEMAI should therefore document which capabilities are vendor-neutral, which are platform-specific, which components are substitutable, and

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

which legal, security, procurement, and compliance approvals are required before production deployment. Figure 2 depicts this Microsoft-oriented reference architecture.

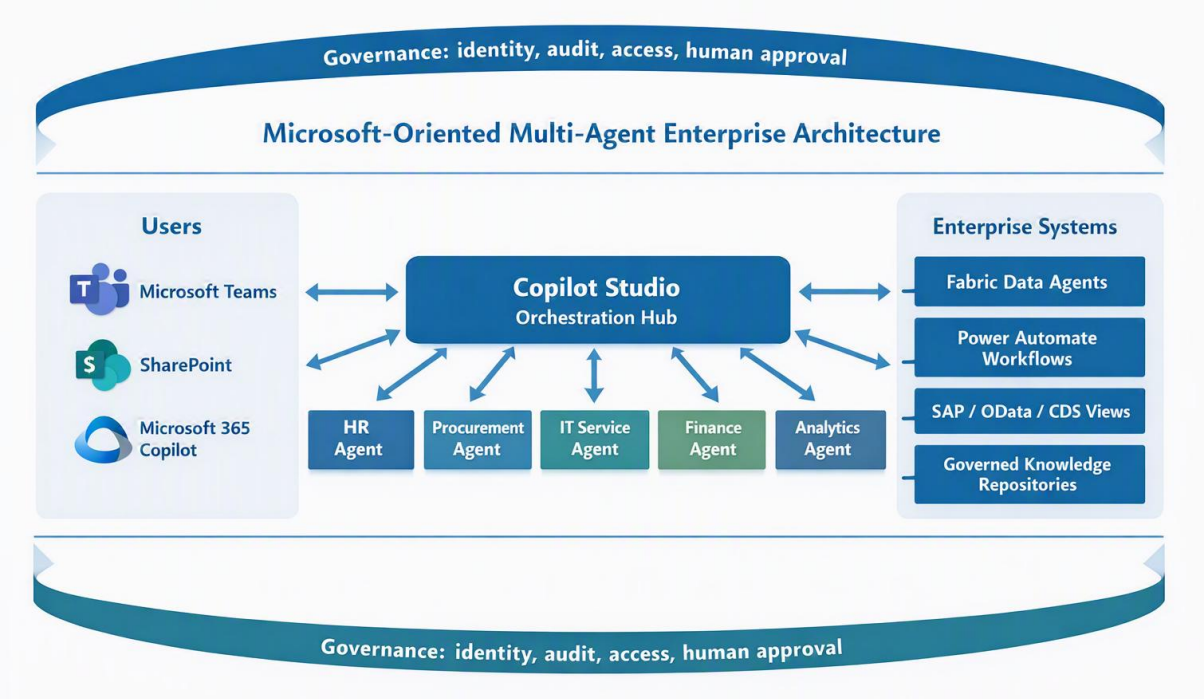

Figure 2. *Microsoft-oriented multi-agent enterprise architecture, with Copilot Studio orchestrating specialized agents across Fabric, Power Platform, and SAP under enterprise governance controls.*

5.5 Decision-Rights Framework

SEMAI distinguishes between **information**, **recommendation**, and **action**. Information responses may be automated if based on approved knowledge or governed data. Recommendations require explanation and confidence indicators. High-impact actions require human approval. This prevents the AI workforce from becoming an uncontrolled automation layer and aligns agent autonomy with organizational accountability. Table 5 details these decision types, the corresponding agent autonomy, and the required governance controls.

Table 5. *SEMAI decision-rights framework: decision types, agent autonomy, and required governance controls.*

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

Decision type	Example	Agent autonomy	Required control
Information retrieval	"Show current purchase order status"	High	Access control and audit log
Analytical explanation	"Why did inventory variance increase?"	Medium	Certified data source and explanation
Recommendation	"Which supplier should be prioritized?"	Medium-low	Human review and rationale
Transactional action	"Create purchase request"	Low	Approval workflow and traceability
High-impact decision	"Change procurement allocation"	Very low	Executive/business owner approval

6. AI Workforce Maturity Model

The maturity model defines five levels. Figure 3 depicts these levels, and Table 6 summarizes their labels, core capabilities, and main risks.

Level 1: Individual AI assistants. Users adopt copilots for writing, summarization, search, and simple productivity tasks. Benefits are local and difficult to measure.

Level 2: Functional AI assistants. Departments deploy HR, finance, procurement, IT, or analytics assistants. Agents are useful but remain siloed.

Level 3: Connected agents. Agents begin to collaborate. For example, a Copilot Studio agent invokes a Fabric data agent for governed analytics or a Foundry agent for specialized reasoning. Microsoft's connected-agent documentation supports this direction [18,19].

Level 4: Process-level AI workforce. Agents coordinate end-to-end workflows such as onboarding, procurement, incident triage, budget variance explanation, or compliance reporting. Human approval is embedded at decision gates.

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

Level 5: Enterprise AI workforce. Multi-agent systems operate across functions with shared governance, observability, model monitoring, role-based access, and executive value tracking.

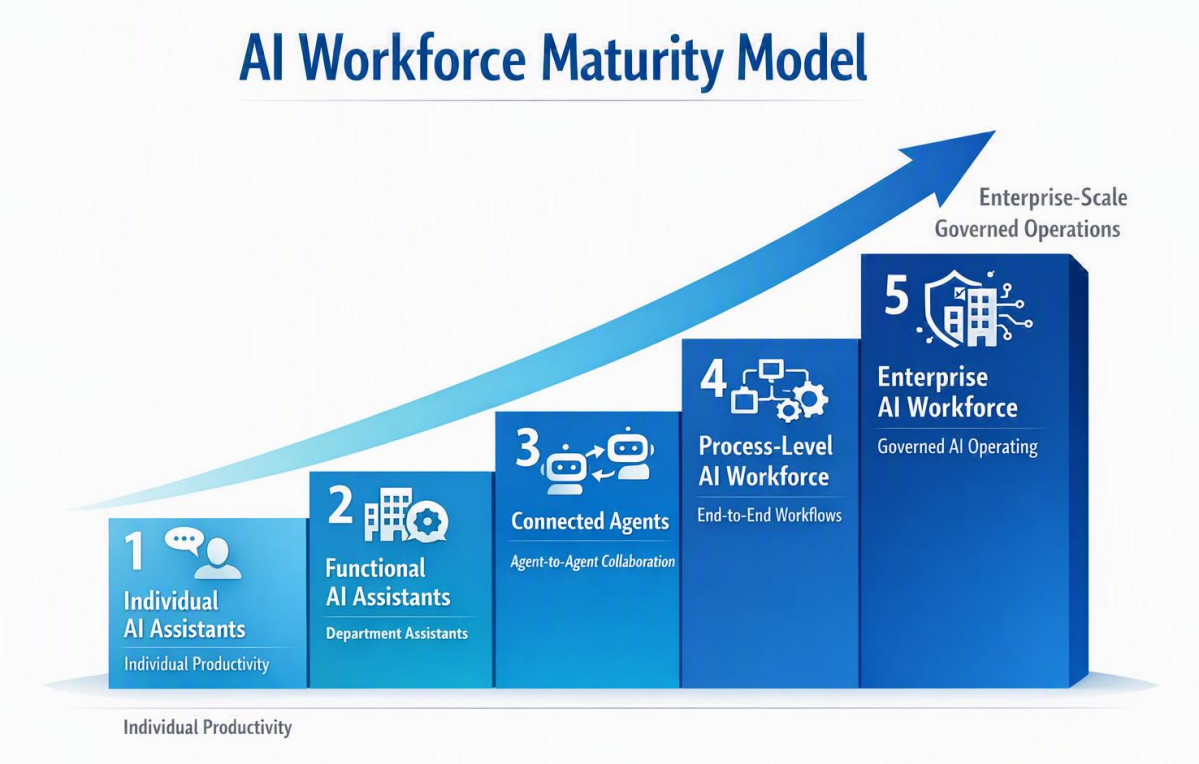

Figure 3. *Five-level AI workforce maturity model, from individual AI assistants to a governed enterprise AI workforce.*

Table 6. *Five-level AI workforce maturity model: levels, labels, core capabilities, and main risks.*

Level	Label	Core capability	Main risk
1	Assistant	Individual productivity	Unmeasured value
2	Functional	Department assistant	Siloed adoption
3	Connected	Agent-to-agent collaboration	Weak orchestration
4	Process workforce	End-to-end workflow support	Accountability gaps

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

5	Enterprise workforce	Governed AI operating model	Scale and assurance complexity

Table 7 specifies the capability, governance, and data-integration criteria for progressing between maturity levels.

Table 7. *Maturity progression criteria for SEMAI.*

Level	**Required capabilities**	**Governance maturity**	**Data/integration maturity**	**Metrics and exit criteria**
1. Individual assistant	Personal productivity use cases, prompt guidance, basic user awareness.	Acceptable-use policy and user review.	Limited enterprise data grounding.	Adoption tracked; move to Level 2 when recurring departmental use cases and owners are defined.
2. Functional assistant	Department-specific assistants for HR, IT, procurement, finance, or analytics.	Department owner, access controls, and support process.	Approved knowledge sources and limited system integration.	Usage, satisfaction, response quality; move to Level 3 when agents can securely call governed data or tools.
3. Connected agents	Agent-to-agent collaboration, tool use, and specialized reasoning.	Agent registry, testing, traceability, and escalation rules.	API, data platform, identity, and metadata integration.	Successful handoffs and exception rate; move to Level 4 when end-to-end workflow

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

				orchestration is stable.
4. Process workforce	Multi-agent support for full business workflows with human approval gates.	Decision rights, RACI, monitoring, incident response, and audit logs.	Certified datasets, ERP/workflow integration, and lineage.	Cycle-time reduction, error reduction, override frequency; move to Level 5 when cross-functional governance and value tracking are institutionalized.
5. Enterprise workforce	Reusable AI workforce patterns across functions and processes.	Enterprise AI governance board, lifecycle assurance, model monitoring, and continuous improvement.	Enterprise data foundation, integrated workflow platforms, and observability.	Portfolio value, audit completeness, risk reduction, and sustained adoption.

7. Demonstration Scenario and Formative Evaluation

The demonstration scenario is based on a generalized setting representative of a large Saudi enterprise undertaking analytics and AI modernization. It is used to demonstrate artifact fitness and support formative evaluation, not to claim empirical validation. In such settings, a common strategic objective is to position IT shared services as a data-driven intelligence function that empowers business units with analytics and AI in support of Saudi Vision 2030. Typical strategic objectives include trusted data foundations, high-impact AI use cases, self-service analytics, responsible AI governance, and internal capability development.

7.1 Scenario: AI Workforce for Employee Onboarding and Operational Readiness

A new employee joins a business function requiring system access, policy orientation, equipment provisioning, and reporting access. In a traditional workflow, HR, IT,

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

procurement, and business managers coordinate manually. In SEMAI, the process is handled by an AI workforce:

1. **HR Agent** verifies onboarding requirements and policy documents.
2. **IT Service Agent** checks application access needs.
3. **Procurement Agent** determines whether equipment or service requests are required.
4. **Analytics Agent** checks role-based reporting access.
5. **SAP/Fabric Data Agent** retrieves governed master-data and role context.
6. **Foundry Specialist Agent** evaluates complex policy exceptions.
7. **Human Manager** approves high-impact or exception-based actions.
8. **Power Platform Workflow Agent** initiates approved workflows.

This scenario maps directly to the Microsoft ecosystem. Copilot Studio provides the interaction layer; Foundry handles advanced specialized reasoning; Fabric data agents provide governed data access; Power Platform executes workflows; SAP provides transactional context; SharePoint and Teams provide employee-facing channels. Common deployment patterns, such as embedding a Copilot Studio agent into SharePoint, show that user-facing agent delivery is already practical in representative Saudi enterprise environments. Figure 4 illustrates this onboarding workflow, and Table 8 maps the demonstration components to represent enterprise evidence and implications.

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

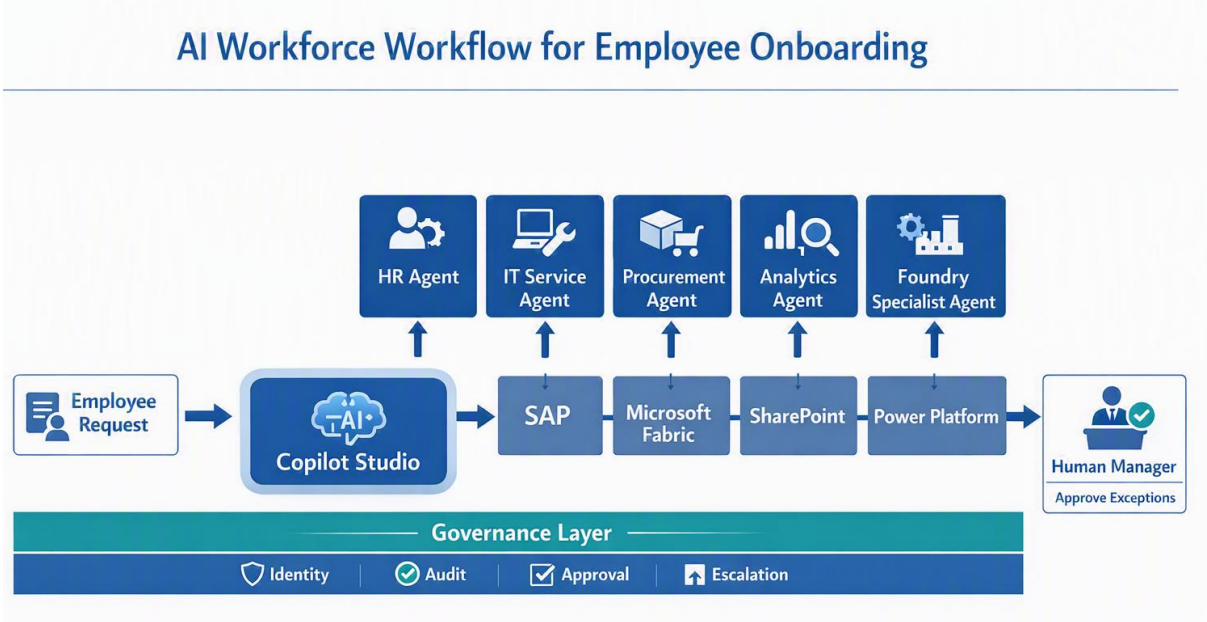

Figure 4. *AI workforce workflow for employee onboarding, showing agent orchestration through Copilot Studio with a governance layer and human approval of exceptions.*

Table 8. *SEMAI demonstration mapping: components, representative enterprise evidence, and demonstration implications.*

SEMAI component	Representative enterprise evidence	Demonstration implication
Enterprise systems	SAP, Microsoft 365, Power BI, Fabric planning	Multi-system integration is required
Data foundation	OneLake, Fabric Data Factory, Purview, and Direct Lake target flow	AI agents require governed data
User-facing agent	Copilot Studio agent embedded into SharePoint	Agent access should meet users where they work
Governance	Purview lineage, sensitivity, and access controls with SDAIA/NDMO alignment	Governance must be built into architecture

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

Operating model	AI governance, center of excellence, and capability development	AI workforce requires people and process design

The scenario demonstrates the framework by application, not by statistical testing or live enterprise measurement. It shows that SEMAI captures realistic enterprise constraints: fragmented systems, role-based access, governance requirements, SAP integration, Microsoft 365 surfaces, data-platform modernization, workflow approvals, and human accountability. The scenario should therefore be interpreted as structured demonstration and formative evaluation rather than empirical proof of performance improvement.

7.2 SEMAI Evaluation Rubric

Table 9 presents the formative evaluation rubric used to assess SEMAI against the study's predefined criteria.

Table 9. *Formative evaluation rubric for SEMAI.*

Criterion	**Evaluation question**	**Expected evidence**
Completeness of agent roles	Are all relevant human, assistant, collaborative-agent, system, and governance roles defined?	Role catalogue, agent registry, RACI, process map.
Governance coverage	Are identity, permissions, policy, monitoring, audit, and escalation controls embedded?	Governance control matrix and audit trail design.
Human accountability	Are approval thresholds, overrides, decision owners, and exception paths specified?	Decision-rights table, approval workflow, decision register.
Enterprise integration	Can agents access certified data and approved systems without bypassing controls?	API design, data lineage, access model, integration testing.

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

Security and privacy	Are data minimization, classification, access control, and incident response addressed?	Security review, privacy impact assessment, data classification.
Workflow orchestration	Can the AI workforce complete cross-functional workflows with traceable handoffs?	Scenario walkthrough, workflow logs, handoff success rate.
Scalability	Can patterns be reused across functions without redesigning governance each time?	Reusable agent templates, platform standards, operating playbook.
Auditability	Can decisions, data sources, tools, prompts, outputs, approvals, and overrides be reconstructed?	Trace logs, decision register, monitoring dashboard.
Business-value measurability	Are productivity, quality, risk, adoption, and financial outcomes measurable?	Value scorecard and baseline/target metrics.
Vision 2030 alignment	Does the AI workforce support national priorities in data, AI, skills, innovation, and productivity?	Vision 2030/SDAIA mapping and capability-development plan.

7.3 Accountability and Responsible AI Controls

Table 10 defines the SEMAI accountability matrix across key AI workforce activities, and Table 11 maps responsible AI principles to enterprise controls and their SEMAI implementation mechanisms.

Table 10. *SEMAI accountability matrix.*

Activity	Responsible	Accountable	Consulted	Informed

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

Agent design and use-case definition	AI product owner and solution architect	Business process owner	Data office, cybersecurity, legal, internal audit	End users and governance board
Data access approval	Data Steward and identity administrator	Data owner	Cybersecurity, privacy/legal, business owner	AI governance board
Prompt and agent-behavior review	AI engineer and domain expert	AI product owner	Risk, compliance, business users	Governance board
Workflow execution	Workflow owner and automation team	Business process owner	IT operations, support, compliance	Affected departments
Human override and exception handling	Assigned approver or manager	Business process owner	Legal, risk, cybersecurity when required	Governance board and audit
Incident response	Security operations and AI operations	CISO or delegated risk owner	Legal, privacy, business owner, vendor support	Executive sponsor and internal audit
Model monitoring and performance review	AI operations team	AI governance board	Business owner, data office, internal audit	Executive sponsor
Business outcome ownership	Process owner and business analytics lead	Executive sponsor	Finance, HR, IT, risk, data office	Governance board and stakeholders

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

Table 11. *Responsible AI control mapping for SEMAI.*

Responsible AI principle	Enterprise control	SEMAI implementation mechanism
Accountability	Decision owner, RACI, decision register, approval trail.	Human approval gates and accountable business owner for high-impact actions.
Transparency	Explanation, documentation, source citation, model card or agent card.	Agent output includes source, confidence, rationale, and limitations where relevant.
Security	Identity, role-based access, threat modeling, monitoring, incident response.	Agents use least-privilege access and monitored tool execution.
Privacy	Data minimization, classification, retention, privacy impact assessment.	Agents access only approved datasets and avoid unnecessary personal data exposure.
Reliability and safety	Testing, red teaming, monitoring, fallback, rollback, human review.	Agents are tested before deployment and monitored for drift, failure, and unsafe outputs.
Fairness	Bias checks, representative testing, impact review.	High-impact recommendations require review for discriminatory or inequitable effects.
Human oversight	Approval thresholds, override mechanism, escalation path.	Autonomy is reduced as decision impact increases.
Contestability	Appeal or challenge route for affected users.	Users can request human review of material AI-assisted decisions.

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

7.4 Value Measurement and Vision 2030 Alignment

Table 12 presents the SEMAI business-value measurement layer, and Table 13 maps Vision 2030 and SDAIA priorities to the corresponding SEMAI components and expected organizational outcomes.

Table 12. *SEMAI business-value measurement layer.*

Metric group	Example measures	Purpose
Productivity	Cycle-time reduction, task automation ratio, time saved per workflow, backlog reduction.	Measure operational efficiency and process acceleration.
Quality	Error rate, rework rate, first-time-right completion, output acceptance rate.	Assess whether AI-assisted workflows improve work quality.
Governance	Audit completeness, source-traceability rate, approval compliance, policy exceptions.	Verify that AI workforce activity remains controlled and reviewable.
Risk	Incident count, data-leakage events, unsafe-output rate, override frequency.	Monitor exposure from agent behavior and automation.
User adoption	Active users, repeat usage, user trust, satisfaction, training completion.	Assess acceptance and practical usability.
Financial/business impact	Cost avoidance, productivity value, service-level improvement, working-capital impact.	Link AI workforce deployment to measurable enterprise outcomes.
Workforce capability	AI literacy, agent-design skills, governance training, citizen-developer participation.	Measure organizational learning and capability development.

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

Table 13. *Vision 2030 and SDAIA alignment mapping.*

Vision 2030/SDAIA priority	Corresponding SEMAI component	Enterprise capability enabled	Expected organizational outcome
Data and AI adoption across the Kingdom	AI assistants, collaborative agents, and enterprise AI operating model.	Reusable AI-enabled workflows across business functions.	Broader practical adoption of AI in enterprise operations.
AI skills and workforce transformation	Human workforce layer and capability-development metrics.	AI literacy, agent stewardship, and citizen-developer participation.	Higher workforce readiness for future-economy roles.
Policies, regulations, and responsible AI	Governance and trust spine, responsible AI controls, RACI.	Accountable, auditable, and policy-aligned AI deployment.	Reduced governance risk and stronger institutional trust.
Research and innovation	Design science artifacts, maturity model, and evaluation rubric.	Reusable reference model for enterprise Agentic AI research and practice.	Contribution to applied computing and enterprise AI knowledge.
Collaborative AI ecosystem	Vendor-neutral principles and platform-specific implementation patterns.	Interoperable AI architecture and cross-functional orchestration.	Improved scalability across departments, sectors, and technology stacks.
Economic diversification and productivity	Business-value measurement layer.	Measurement of productivity, quality, cost avoidance, and	Evidence-based AI investment and improved enterprise competitiveness.

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

		service improvement.	

7.5 Implementation Risks and Mitigation Strategies

Table 14 summarizes the primary implementation risks and their corresponding mitigation strategies.

Table 14. *Implementation risks and mitigation strategies.*

Risk	**Description**	**Mitigation**
Hallucination	Agent produces unsupported or inaccurate output.	Use grounded sources, citations, confidence indicators, testing, and human review.
Tool misuse	Agent calls an inappropriate tool or performs an unintended action.	Restrict tool permissions, define allowed actions, require approval for transactions.
Data leakage	Agent exposes sensitive or unauthorized information.	Apply least privilege, sensitivity labels, data minimization, and monitoring.
Conflicting agent outputs	Different agents produce inconsistent recommendations.	Use arbitration rules, source ranking, escalation paths, and human decision owner.
Automation bias	Users over-trust AI recommendations.	Require rationale, uncertainty disclosure, training, and contestability.
Vendor lock-in	Architecture becomes dependent on one platform roadmap or licensing model.	Separate vendor-neutral principles from reference implementation and document substitutes.

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

Regulatory misalignment	Deployment fails to meet privacy, security, or sector-specific requirements.	Perform legal, privacy, cybersecurity, and compliance reviews before production.
Weak observability	Agent decisions cannot be reconstructed or audited.	Maintain logs, decision registers, telemetry, and model/agent inventories.

8. Discussion

The proposed model reframes enterprise AI from a tool-adoption problem into an operating-model problem. This is important because many AI projects fail to scale when they remain disconnected from enterprise data, workflows, governance, and decision accountability. SEMAI addresses this by treating AI assistants as only one layer of a broader AI workforce.

The model also strengthens the link between AI and Vision 2030. SDAIA's strategy emphasizes data-and-AI adoption, skills, investment, research, innovation, and ecosystem development. SEMAI translates these macro-level ambitions into an enterprise-level implementation model: develop talent, govern AI, integrate enterprise data, connect agents, measure outcomes, and preserve human accountability [28,29].

The Microsoft architecture is practical because many Saudi enterprises already operate Microsoft 365, Teams, SharePoint, Power Platform, Power BI, and Azure-based environments. Microsoft documentation supports connected agents across Copilot Studio, Fabric data agents, and Foundry agents, while emphasizing requirements for permissions, testing, observability, traceability, and human oversight. This makes the architecture technically plausible, though production implementation requires security review, licensing validation, tenant configuration, and legal/privacy assessment [18,19].

The revised model should be interpreted as a design science artifact supported by conceptual demonstration and formative evaluation, not as a statistically validated or production-tested system. Its expected value lies in making the design logic reusable: organizations can specify agent roles, assign accountability, evaluate governance

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

coverage, track business value, and decide which technology stacks best implements the vendor-neutral principles. This positioning moderates the claims of scalability and enterprise readiness while still showing how SEMAI can guide implementation planning.

9. Limitations, Boundary Conditions, and Future Research

This study has several limitations. First, SEMAI was not tested in a live enterprise deployment. Second, the scenario is illustrative and generalized rather than based on measured operational outcomes. Third, no expert-panel, Delphi, practitioner survey, or longitudinal case-study validation was conducted in this revision. Fourth, Microsoft architecture is a reference implementation and may not generalize to organizations using different technology stacks. Fifth, the figures are conceptual illustrations, not empirical evidence. Sixth, governance controls require legal, regulatory, cybersecurity, privacy, and sector-specific review before deployment. Seventh, agent behavior and platform capabilities may change as vendor products evolve, which can affect interoperability, licensing, monitoring, and compliance assumptions.

SEMAI is most applicable to medium-to-large organizations with sufficient data maturity, defined process ownership, existing enterprise systems, AI governance capacity, workflow-automation opportunities, human approval structures, and executive sponsorship. It is especially relevant for digitally mature Saudi enterprises, government-related entities, and large private organizations operating integrated Microsoft, SAP, data-platform, or workflow environments. Smaller organizations or entities with weak data governance may need to begin with foundational data quality, identity, policy, and process-mapping work before adopting multi-agent AI workforces.

Future research should evaluate SEMAI using expert panels, Delphi validation, practitioner walkthroughs, controlled pilots, comparative analysis against existing enterprise architecture and AI governance frameworks, or longitudinal case studies. Future studies should measure outcomes such as cycle-time reduction, error rate, service backlog reduction, user trust, override frequency, exception rate, audit completeness, data-access compliance, cost avoidance, and workforce capability

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

development. Additional work should also examine sector-specific adaptations for government, manufacturing, healthcare, finance, logistics, education, and critical infrastructure.

10. Conclusion

This paper proposed SEMAI, a Saudi Enterprise Multi-Agent AI Operating Model for organizations seeking to move from isolated AI assistants toward governed AI workforces. The framework integrates human roles, AI assistants, collaborative agents, enterprise systems, and governance controls. It also provides Microsoft ecosystem architecture using Copilot, Copilot Studio, Microsoft Fabric, Microsoft Foundry, Power Platform, SAP integration, and governed enterprise data.

The study contributes to applied computing and informatics by translating multi-agent AI concepts into a practical enterprise operating model. It contributes to Saudi Vision 2030 research by linking enterprise AI workforce design to national priorities in advanced computing, data-and-AI adoption, workforce development, digital transformation, and future-economy competitiveness. The central conclusion is that enterprise AI value will not be achieved by deploying more assistants alone. It will be achieved when AI agents are organized as accountable, governed, human-supervised workforces embedded into enterprise processes and decision systems.

Declarations

Funding: This research received no external funding.
Conflicts of interest: The author declares no conflict of interest.
Data availability: No public dataset was generated or analyzed. The demonstration scenario uses generalized, non-confidential observations representative of a large Saudi enterprise and does not disclose confidential data.
Ethical approval: Not applicable; this study did not involve human or animal subjects, personal data, or samples.
Generative AI disclosure: Generative AI tools were used to assist with language refinement, structure, and formatting, and to create the illustrative conceptual figures. The figures are illustrative rather than factual or evidential. The author reviewed all

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

outputs and remains fully responsible for the content, interpretation, references, figures, and conclusions.

References

[1] Abraham, R., Schneider, J. and vom Brocke, J. Data governance: a conceptual framework, structured review, and research agenda. *International Journal of Information Management*, Vol. 49, pp. 424-438, 2019. https://doi.org/10.1016/j.ijinfomgt.2019.07.008

[2] Benbya, H., Davenport, T.H. and Pachidi, S. Artificial intelligence in organizations: current state and future opportunities. *MIS Quarterly Executive*, Vol. 19 No. 4, pp. 1-14, 2020.

[3] Berente, N., Gu, B., Recker, J. and Santhanam, R. Managing artificial intelligence. *MIS Quarterly*, Vol. 45 No. 3, pp. 1433-1450, 2021. https://doi.org/10.25300/MISQ/2021/16274

[4] Bharadwaj, A., El Sawy, O.A., Pavlou, P.A. and Venkatraman, N. Digital business strategy: toward a next generation of insights. *MIS Quarterly*, Vol. 37 No. 2, pp. 471-482, 2013. https://doi.org/10.25300/MISQ/2013/37:2.3

[5] Chen, H., Chiang, R.H.L. and Storey, V.C. Business intelligence and analytics: from big data to big impact. *MIS Quarterly*, Vol. 36 No. 4, pp. 1165-1188, 2012. https://doi.org/10.2307/41703503

[6] Davenport, T.H. and Ronanki, R. Artificial intelligence for the real world. *Harvard Business Review*, Vol. 96 No. 1, pp. 108-116, 2018.

[7] Dorri, A., Kanhere, S.S. and Jurdak, R. Multi-agent systems: a survey. *IEEE Access*, Vol. 6, pp. 28573-28593, 2018. https://doi.org/10.1109/ACCESS.2018.2831228

[8] Dwivedi, Y.K., Hughes, L., Ismagilova, E., Aarts, G., Coombs, C., Crick, T., Duan, Y., Dwivedi, R., Edwards, J., Eirug, A., Galanos, V., Ilavarasan, P.V., Janssen, M., Jones, P., Kar, A.K., Kizgin, H., Kronemann, B., Lal, B., Lucini, B., Medaglia, R., Le Meunier-FitzHugh, K., Le Meunier-FitzHugh, L.C., Misra, S.K., Mogaji, E.,

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

Sharma, S.K., Singh, J.B., Raghavan, V., Raman, R., Rana, N.P., Samothrakis, S., Spencer, J., Tamilmani, K., Tubadji, A., Walton, P. and Williams, M.D. Artificial intelligence (AI): multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy. *International Journal of Information Management*, Vol. 57, 101994, 2021. https://doi.org/10.1016/j.ijinfomgt.2019.08.002

[9] European Commission High-Level Expert Group on AI. *Ethics Guidelines for Trustworthy AI*. European Commission, Brussels, 2019.

[10] Ferber, J. *Multi-Agent Systems: An Introduction to Distributed Artificial Intelligence*. Addison-Wesley, Harlow, 1999.

[11] Floridi, L. and Cowls, J. A unified framework of five principles for AI in society. *Harvard Data Science Review*, Vol. 1 No. 1, 2019.

[12] Gregor, S. and Hevner, A.R. Positioning and presenting design science research for maximum impact. *MIS Quarterly*, Vol. 37 No. 2, pp. 337-355, 2013. https://doi.org/10.25300/MISQ/2013/37.2.01

[13] Hevner, A.R., March, S.T., Park, J. and Ram, S. Design science in information systems research. *MIS Quarterly*, Vol. 28 No. 1, pp. 75-105, 2004. https://doi.org/10.2307/25148625

[14] ISO/IEC. *ISO/IEC 42001: Information Technology – Artificial Intelligence – Management System*. International Organization for Standardization, Geneva, 2023.

[15] Jennings, N.R. An agent-based approach for building complex software systems. *Communications of the ACM*, Vol. 44 No. 4, pp. 35-41, 2001. https://doi.org/10.1145/367211.367250

[16] Khatri, V. and Brown, C.V. Designing data governance. *Communications of the ACM*, Vol. 53 No. 1, pp. 148-152, 2010. https://doi.org/10.1145/1629175.1629210

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

[17] March, S.T. and Smith, G.F. Design and natural science research on information technology. *Decision Support Systems*, Vol. 15 No. 4, pp. 251-266, 1995. https://doi.org/10.1016/0167-9236(94)00041-2

[18] Microsoft. Consume a Fabric data agent in Microsoft Copilot Studio. Microsoft Learn, 2026a. Available at: https://learn.microsoft.com/en-us/fabric/data-science/data-agent-microsoft-copilot-studio (accessed 19 July 2026).

[19] Microsoft. Connect to a Microsoft Foundry agent from Copilot Studio. Microsoft Learn, 2026b. Available at: https://learn.microsoft.com/en-us/microsoft-copilot-studio/add-agent-foundry-agent (accessed 19 July 2026).

[20] Mikalef, P. and Gupta, M. Artificial intelligence capability: conceptualization, measurement calibration, and empirical study on its impact on organizational creativity and firm performance. *Information & Management*, Vol. 58 No. 3, 103434, 2021. https://doi.org/10.1016/j.im.2021.103434

[21] NIST. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. National Institute of Standards and Technology, Gaithersburg, MD, 2023.

[22] OECD. *OECD Principles on Artificial Intelligence*. OECD, Paris, 2019.

[23] Peffers, K., Tuunanen, T., Rothenberger, M.A. and Chatterjee, S. A design science research methodology for information systems research. *Journal of Management Information Systems*, Vol. 24 No. 3, pp. 45-77, 2007. https://doi.org/10.2753/MIS0742-1222240302

[24] Power, D.J. *Decision Support Systems: Concepts and Resources for Managers*. Quorum Books, Westport, CT, 2002.

[25] Raisch, S. and Krakowski, S. Artificial intelligence and management: the automation-augmentation paradox. *Academy of Management Review*, Vol. 46 No. 1, pp. 192-210, 2021. https://doi.org/10.5465/amr.2018.0072

[26] Raji, I.D., Smart, A., White, R.N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D. and Barnes, P. Closing the AI accountability gap: defining

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

an end-to-end framework for internal algorithmic auditing. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pp. 33-44, 2020. https://doi.org/10.1145/3351095.3372873

[27] Russell, S. and Norvig, P. *Artificial Intelligence: A Modern Approach*, 4th ed. Pearson, Hoboken, NJ, 2021.

[28] Saudi Data & AI Authority (SDAIA). National Strategy for Data & AI (NSDAI). SDAIA, 2020/2025.

[29] Saudi Data & AI Authority (SDAIA). SDAIA and Vision 2030. SDAIA, Riyadh, 2026.

[30] Shoham, Y. and Leyton-Brown, K. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, Cambridge, 2008.

[31] Shollo, A. and Galliers, R.D. Towards an understanding of the role of business intelligence systems in organisational knowing. *Information Systems Journal*, Vol. 26 No. 4, pp. 339-367, 2016. https://doi.org/10.1111/isj.12071

[32] Stone, P. and Veloso, M. Multiagent systems: a survey from a machine learning perspective. *Autonomous Robots*, Vol. 8, pp. 345-383, 2000.

[33] Vaishnavi, V. and Kuechler, W. *Design Science Research Methods and Patterns: Innovating Information and Communication Technology*, 2nd ed. CRC Press, Boca Raton, FL, 2015.

[34] Verhoef, P.C., Broekhuizen, T., Bart, Y., Bhattacharya, A., Dong, J.Q., Fabian, N. and Haenlein, M. Digital transformation: a multidisciplinary reflection and research agenda. *Journal of Business Research*, Vol. 122, pp. 889-901, 2021. https://doi.org/10.1016/j.jbusres.2019.09.022

[35] Vial, G. Understanding digital transformation: a review and a research agenda. *Journal of Strategic Information Systems*, Vol. 28 No. 2, pp. 118-144, 2019. https://doi.org/10.1016/j.jsis.2019.01.003

Nexus Science Review
INTERNATIONAL JOURNAL OF INTERDISCIPLINARY RESEARCH

Received: 10/06/2026 - **Accepted:** 12/07/2026
Vol: 01.02
DOI: 10.66104/9zjxnd21
Pages: 1-38

[36] Vision 2030. Saudi Vision 2030 official portal, 2026. Available at: https://www.vision2030.gov.sa/en (accessed 19 July 2026).

[37] Weber, K., Otto, B. and Österle, H. One size does not fit all: a contingency approach to data governance. *Journal of Data and Information Quality*, Vol. 1 No. 1, pp. 1-27, 2009. https://doi.org/10.1145/1515693.1515696

[38] Weiss, G. (Ed.). *Multiagent Systems*, 2nd ed. MIT Press, Cambridge, MA, 2013.

[39] Wooldridge, M. *An Introduction to MultiAgent Systems*, 2nd ed. Wiley, Chichester, 2009.

[40] Al Naqbi, H., Bahroun, Z. and Ahmed, V. Enhancing work productivity through generative artificial intelligence: a comprehensive literature review. *Sustainability*, Vol. 16 No. 3, 1166, 2024. https://doi.org/10.3390/su16031166

[41] Tran, K.-T., Dao, D., Nguyen, M.-D., Pham, Q.-V., O'Sullivan, B. and Nguyen, H.D. Multi-agent collaboration mechanisms: a survey of LLMs. *arXiv*, 2025. https://doi.org/10.48550/arXiv.2501.06322

[42] Bandi, A., Kongari, B., Naguru, R., Pasnoor, S. and Vilipala, S.V. The rise of agentic AI: a review of definitions, frameworks, architectures, applications, evaluation metrics, and challenges. *Future Internet*, Vol. 17 No. 9, 404, 2025. https://doi.org/10.3390/fi17090404

[43] Saudi Press Agency. SDAIA launches AI framework to promote ethical, responsible use of advanced technologies in Saudi Arabia. SPA, 2024. Available at: https://www.spa.gov.sa/en/N2217019