



ORCHESTRATION ALIGNS AI AGENTS WITH BUSINESS MAP AGENTIC AI TO BUSINESS NEEDS WITH SCALABLE ORCHESTRATION PRACTICES

Abstract

As enterprises rapidly adopt AI agents, they risk inadvertently creating fragmented AI ecosystems, escalating costs and limiting interoperability. Robust integration, automation, and orchestration practices are critical if enterprises are to govern AI agents, ensure accurate outputs, and realize sustained value from their AI investments.

This paper highlights the need for a strategic approach to agentic AI orchestration. It explores a three-dimensional framework built on modular integration, enterprise knowledge, and packaged data. It also examines how effective agentic AI orchestration can scale across the enterprise—safely, reliably, and predictably—unlocking greater business value for organizations.

Abstract	01
Introduction	02
Need for a Strategic Approach to Agentic AI Orchestration	02
Three Dimensions of Agentic AI Integration and Orchestration	02
Dimension 1: Create modular capabilities	03
How it works	03
Impact	03
Takeaway	03
Dimension 2: Build enterprise knowledge and context	03
How it works	04
Impact	04
Takeaway	05
Dimension 3: Package enterprise data as products	05
How it works	05
Impact	05
Takeaway	05
Building Unified, Hybrid Workflows	06
Recommendations: From Strategy to Practice	06
Conclusion	07
References	07
About the Authors	07

Introduction

Organizations across industries are rapidly embedding artificial intelligence (AI) agents into business workflows, from automated customer service to predictive maintenance. Just as organizations typically run their operations on heterogeneous environments, most of them rely on AI agents procured from different vendors and ecosystems, such as Microsoft Copilot, SAP Joule, and Salesforce Agentforce. AI agents are often built on different architectures and distributed across multiple cloud environments. The result is a fragmented AI ecosystem where agents operate in silos, making coordination, interoperability, and governance increasingly challenging.

Need for a Strategic Approach to Agentic AI Orchestration

Industry experts warn that uncoordinated AI agents increase operational risk rather than deliver business value. This makes agent orchestration a critical capability. Orchestration acts as a control plane that oversees how work is executed. It ensures that AI agents collaborate effectively, route tasks to the right capabilities, and operate within predefined enterprise policies. From a business perspective, orchestration transforms loosely coupled AI bots into dependable collaborators by aligning them with enterprise processes, data, and governance frameworks.

In an orchestrated model, agents operate within a unified framework that enforces consistency, auditability, and control. Decision paths are traceable and reversible. High-impact workflows are protected with guardrails, such as automated approvals for high-confidence actions and human review for critical decisions. This hybrid approach keeps humans in the loop for strategy, risk, and ethics, while allowing agents to execute routine tasks at speed and scale.

With robust governance of AI agents, business outcomes improve significantly. Research shows that organizations with strong data and process foundations extract greater value from AI. By combining intelligence, automation, and orchestration, enterprises can move beyond isolated AI experiments to reimagine end-to-end processes as integrated human-AI workflows. These could include order-to-cash, claims handling, or supply chain operations. The result is faster cycle times, fewer errors, continuous responsiveness, and stronger enterprise-grade controls. All of these are essential to realize the full promise of agentic AI.

Three Dimensions of Agentic AI Integration and Orchestration

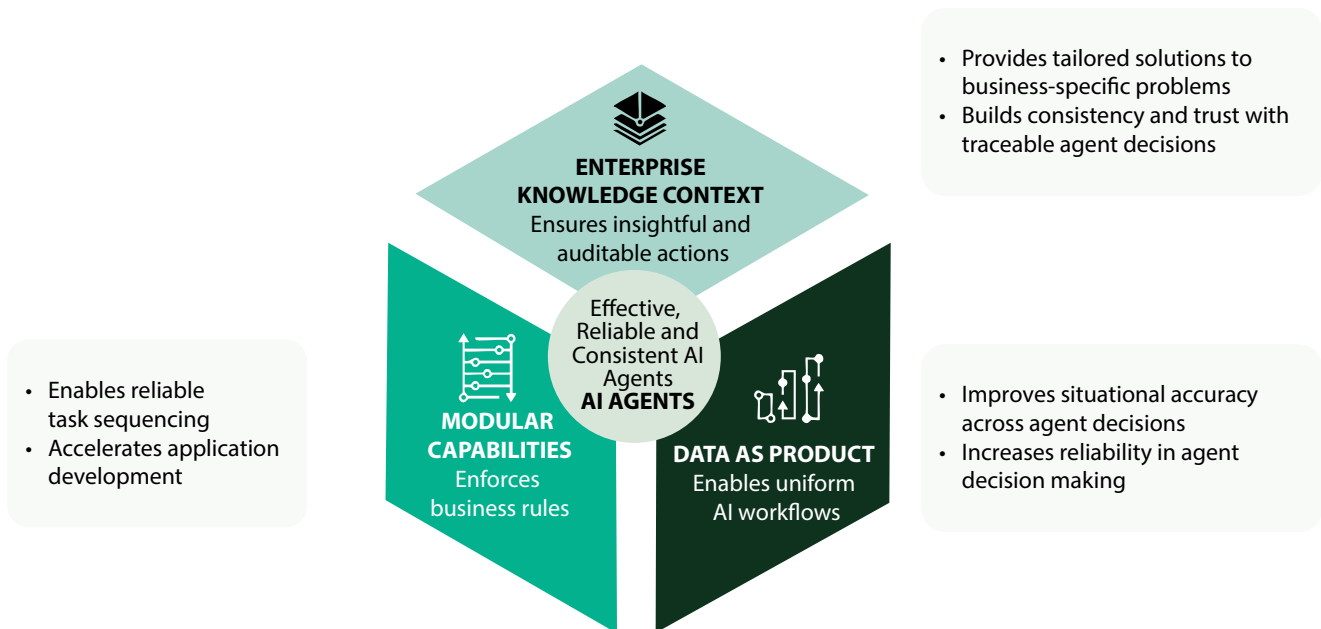


Fig 1: Key aspects of the three dimensions of agentic AI orchestration

Dimension 1: Create modular capabilities

Integration is crucial to agentic workflows. Agents are only as powerful as the systems they can sense and act on. A robust integration strategy using application programming interfaces (APIs), connectors, and microservices ensures that data flows freely and consistently between agents and enterprise applications. Without strong integration, AI agents will struggle with the same silos and blind spots that humans face, resulting in partial visibility and poor decisions.

How it works

Effective orchestration treats every existing enterprise capability as a building block for AI. Legacy systems, microservices, APIs, robotic process automation (RPA) bots, and workflows are exposed as tools that agents can use. For example, a loan-processing bot might call a credit-check microservice API or update records in a customer relationship management (CRM) system.

Industry experts call this 'exposing capabilities via standardized interfaces' using the emerging Model Context Protocol (MCP), which provides a standard protocol for AI agents to connect to enterprise systems and data. Forrester predicts that nearly 30% of enterprise apps will embed their own MCP servers by 2026, creating hubs where AI agents can securely plug into each business application. This means that the agent orchestrator can dynamically route an agent to the right API across finance, human resources, inventory, and other business functions. Just as a smartphone uses USB-C ports, the MCP provides a uniform connection point.

Impact

When enterprises offer every business function as a standardized service or 'skill', agents can sequence tasks reliably. They can break complex processes into modular steps, calling services for data fetch, validation, and execution. This modular approach also accelerates development. Rather than coding each agent from scratch, operations teams can convert existing workflows into AI agents on a governed platform. The result is Lego®-like, reusable functionality, such as payment processing and policy or customer lookup, that can be orchestrated that can be orchestrated to support new and emerging use cases.

In practice, the orchestrator itself may be a workflow engine or dedicated agent controller that can:

- Manage task delegation by determining which agent or service handles each step
- Run parallel execution where multiple agents tackle a problem together
- Perform error recovery such as retries and human fallback

Takeaway

Strong integration enforces business rules. For instance, an agent should call a payments API only after required approvals. Enterprises should expose core business capabilities through stable, well-defined interfaces. They must coordinate all AI agents through a shared execution layer that enforces rules, guardrails, and policies. By wrapping business capabilities as callable APIs, MCPs, or agents, organizations can ensure that the agentic AI ecosystem remains disciplined and governed. Such a system succeeds by anchoring to existing processes, compliance requirements, and operational controls, rather than operating in an ad-hoc manner.

Dimension 2: Build enterprise knowledge and context

Agents thrive on contextual knowledge within the enterprise. In addition to live data, they need access to corporate information, such as policies, customer history, product catalogs, standard operating procedures (SOPs), and regulatory guidelines. Without this institutional knowledge, agents may hallucinate or violate policy. For example, an insurance claims bot must know company coverage rules to avoid issuing unauthorized payments. Further, without clear context of organizational rules, agents may act in ways that violate company policies. The solution lies in equipping AI agents with enterprise knowledge and process context, embedding corporate constraints into agent workflows.

How it works

Enterprises deploy knowledge repositories and retrieval systems. Agents use techniques like retrieval-augmented generation (RAG) to query internal knowledge on demand. Confluent implementations underscore the need for agents to be grounded in a real-time, contextualized, and trustworthy information. As agents use reasoning capabilities, they draw snippets from approved documents. For instance, an agent could consult the latest safety manual when making factory scheduling decisions or fetch the customer's contract when handling a support case.

These knowledge repositories become part of the orchestration fabric. The agent control layer manages which documents the agent can query and tags outputs with their source metadata. Enabling such capabilities calls for flexible architecture that protects proprietary organizational context and institutional knowledge, keeping it under enterprise control even as AI tools and vendors evolve.

Impact

Embedding knowledge into agentic workflows yields two benefits:

- Augments LLMs with deep domain understanding: On their own, large language models (LLMs) know only broad patterns. When augmented with local data, they can provide answers tailored to the business.
- Builds consistency and trust: Every agent's decision is traceable to a known source, significantly lowering risk. Orchestrated agents execute the same set of actions consistently and operate within a controlled framework, ensuring full visibility into all actions.

Takeaway

Embedding AI agents with enterprise knowledge ensures insightful and auditable actions.

Notably, the knowledge layer ensures that the AI agent uses only accurate and vetted information.

In practice, a user request initiates a governed flow: the agent queries the knowledge store for relevant policies, proposes an action, routes it for human review and approval, and executes the action upon approval.

Dimension 3: Package enterprise data as products

AI agents make decisions in real time. Hence, their insights must be based on current information. For agentic AI to scale successfully, it needs strong datasets. However, 80% of companies cite data limitations as a barrier to agentic AI. Disparate, delayed, or poor-quality data can cause agents to fail or produce incorrect outputs.

AI agents need real-time data streams as well as historical stores that act as first-class products for agents. Modern enterprises accumulate massive amounts of data across warehouses, data lakes, and event streams, such as Kafka topics as well as feeds from the Internet of Things (IoT). Orchestration demands that agents seamlessly tap into all relevant data sources.

How it works

Orchestration-ready architecture treats data ingestion like a product. Enterprises must package data into discoverable, governed data products. This means cleaned, unified, and documented datasets or feeds that AI agents can consume. Some examples of data products with an API or streaming endpoint include market prices, inventory levels, CRM data, and machine metrics. Real-time data streams, once captured, can be made widely accessible with clear metadata and access controls. The best practice is to layer a streaming platform, such as Kafka or Confluent, under the agentic AI ecosystem. The Confluent Data Streaming Platform acts as a coordination path, allowing agents to subscribe to event streams for real-time context. Platforms like Apache Flink or ksqldb for Confluent Platform also enrich and prepare data dynamically.

Impact

With this approach, AI agents can read events in real time. An order-processing agent can immediately identify a stockout event through a stream, while a support agent can access the latest usage logs from a database. Importantly, integration via APIs, extract, transform, load (ETL), or infrastructure platform as a service (iPaaS) must ensure that data flows remain governed and consistent. Only then can agents reason with reliability. In practice, this would look like unified customer data across CRM and billing systems—bundled as an integrated data product—allowing an agent to personalize support queries. Real-time data feeds, such as production sensor data, can help an agent to respond to anomalies immediately. By assembling static datasets (in data warehouses as well as big data platforms) with streaming sources in the orchestration layer, agents gain complete visibility. This real-time context enables agents to augment generative AI (GenAI) outputs with situational accuracy.

Takeaway

Unified data enables unified AI workflows. With well-governed data products and streaming pipelines, agents become a dynamic front end to an enterprise's real-time view. In effect, data orchestration and agent orchestration form a single continuum. Agents consume data products, while orchestration ensures that every agent action is logged and traceable to the data that informed it.



Building Unified, Hybrid Workflows

When enterprises assemble the capabilities, knowledge, and data described above, they can build hybrid workflows that combine deterministic process steps with GenAI tasks. This resembles the integration of traditional business process management (BPM) with modern LLMs. For example, in a loan approval flow, an AI agent can draft the initial assessment using customer data and policy documents. It then routes the case through the bank's existing workflow engine for compliance checks and finally summarizes the decision for a loan officer in natural language.

The orchestrator is in charge of sequencing these actions. It schedules the LLM-based steps where nuance is needed. It also calls on traditional automation for precise actions, such as writing to databases and triggering emails. AI agents handle the uncertain reasoning, while the orchestrator maintains strict control over approval, audit trails, and final execution.

Hybrid workflows give business leaders the best of both worlds. They can accelerate innovation by enabling agents to ideate, simulate and propose options quickly while preserving reliability. Some organizations are already piloting workflows where agents and humans operate in a loop. For example, an AI agent might suggest a product design based on customer feedback and research and development (R&D) data. A human engineer then reviews the proposal and adjusts parameters. Finally, the agent runs a cost analysis using supply chain data. This creates a continuous feedback loop, enabling a self-optimizing process bounded by business logic.


The full value of LLMs emerges when they operate in a data-rich environment. Consider Confluent's model that describes streaming real-time business events for a knowledge base that AI agents can query even as they reason. By feeding enterprise documents and live data into LLM prompts, agents can generate insights that are rooted in the company's reality. For instance, an LLM with up-to-date inventory data can advise sales teams on substitutions in real time, while an LLM trained on past contracts and policies can draft compliant legal documents. With strong integration, context, and orchestration, GenAI assistants become context-aware collaborators, closing the gap between insight and action by coordinating across all enterprise systems.

Recommendations: From Strategy to Practice

For successful AI agent orchestration, enterprises need a structured approach comprising the following key steps:

- **Map high-impact processes to agentic scenarios:** Start by identifying a few end-to-end workflows, such as customer service triage or supply chain exception handling, where autonomous AI agents can deliver value. These become 'lighthouse' use cases that build momentum.
- **Invest in a modular tech stack:** Build or adopt an orchestration layer that exposes core enterprise systems via APIs or MCP. Several enterprise application vendors now provide built-in MCP or agent APIs. MCP is essential for new agents to plug into existing systems, such as enterprise resource planning (ERP), CRM, and data lakes.
- **Modernize data and knowledge foundations:** Treat data as a product. Create governed data pipelines for batch and real-time processing, supported by a single source of truth. Additionally, consolidate knowledge bases and enable RAG. This may involve deploying vector stores or enterprise search to make documents accessible to AI. Identify data and tool components that can be shared with multiple agent workflows. Gradually expand the agentic network within the organization while continuously refining the guardrails.
- **Govern and operate proactively:** Define clear policies, access controls, and audit trails for AI agents from the outset. Use human-in-the-loop thresholds for sensitive actions. Establish an oversight framework to assign agent identities, issue permissions, and monitor behavior continuously.

In the future, enterprises that excel at agent orchestration will gain a decisive advantage. By combining modular capabilities, real-time data, and contextual intelligence, organizations can deploy AI agents that not only reason effectively but also act with reliability and consistency. The opportunity is sizeable but capturing it requires intentional architectural choices.



Conclusion

Agentic AI delivers business value by closing the gap between insight and action. Successful agent orchestration is the key to amplify the value of agentic AI across the enterprise, safely and predictably. Isolated experimentation with standalone AI tools is no longer sufficient. To succeed in the agentic AI era, organizations must establish a centralized agent orchestration strategy, anchored on a clear set of foundational pillars.

References

- McKinsey & Company | [Building the foundations for agentic AI at scale](#)
- Pega | The missing ingredient: [Why orchestration is the key to success with AI agents](#)
- Confluent | [What is Agentic AI?](#)
- Elementum | [What Is Agentic AI Orchestration? An Enterprise Guide](#)
- Everest Group | [Rewiring Sales Services with Agentic AI: Orchestrating Outcomes through Systems of Execution](#)
- Forrester | [Predictions 2026: AI Agents, Changing Business Models, And Workplace Culture Impact Enterprise Software](#)
- Boomi | [Everything You Ever Wanted to Know About AI Agents \(But Were Afraid to Ask\)](#)
- McKinsey & Company | [The agentic organization: Contours of the next paradigm for the AI era](#)

About the Authors

Kapil Goyal

Vice President, Delivery Head – Infosys

Kapil Goyal is a Vice President and Global Delivery Head at Infosys, leading large transformation programs across Digital, Data, and Allied portfolios. He brings strong P&L ownership, enterprise delivery governance expertise, and experience scaling complex global engagements to drive businessaligned outcomes for strategic clients.

Ashok Kumar C. S.

AVP, Senior Industry Principal – Infosys

Ashok Kumar C. S. is an Associate Vice President and Senior Industry Principal at Infosys, heading consulting initiatives in agentic automation, data and platformled transformation. He advises CXO stakeholders on process modernization, integration strategy, and enterprise architecture, combining deep domain expertise with hands on transformation leadership.

For more information, contact askus@infosys.com



© 2026 Infosys Limited, Bengaluru, India. All Rights Reserved. Infosys believes the information in this document is accurate as of its publication date; such information is subject to change without notice. Infosys acknowledges the proprietary rights of other companies to the trademarks, product names and such other intellectual property rights mentioned in this document. Except as expressly permitted, neither this documentation nor any part of it may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, printing, photocopying, recording or otherwise, without the prior permission of Infosys Limited and/ or any named intellectual property rights holders under this document.