From Digital Factories to Systems of Agency: Architecting the Autonomous Enterprise

FEATURING RESEARCH FROM FORRESTER

How Agentic AI Reshapes Your Enterprise

EXECUTIVE SUMMARY

Digital factories have enabled enterprises to achieve scale, speed, and standardization in application development. However, their reliance on predefined workflows and human-led execution limits adaptability and real-time decision-making.

Systems of agency are the next evolution in enterprise architecture, where AI agents act as goal-driven orchestrators. These systems interpret intent, dynamically plan execution, and coordinate services through APIs within governed guardrails, shifting enterprises from automation to autonomous operations. Therefore, there is need for redefinition of how enterprises approach IT systems engineering itself.

At Tech Mahindra, this thinking is shaping our evolution from traditional ADMS (Application Development Maintenance Services) to Agentic Development & Modernization Services (ADMS), anchored in the belief that enterprises must now be engineered for autonomy, not just efficiency. This paper focuses on how enterprises can build a strong foundation for engineering the autonomous enterprise of today and tomorrow by re-architecting platforms into composable, agent-ready systems and evolving software engineering from human-driven workflows to agent-orchestrated execution.

IN THIS DOCUMENT

From Digital Factories to Systems of Agency: Architecting the Autonomous Enterprise

Research From Forrester: How Agentic AI Reshapes Your Enterprise

About Tech Mahindra

TABLE OF CONTENTS

REIMAGINING ENTERPRISE ARCHITECTURE IN THE AGE OF AGENCY

Over almost a decade, enterprises have invested increasingly in digital factories to modernize legacy systems, accelerate product innovation, and embed agile and DevSecOps practices across engineering organizations. These factories to a good extent succeeded in breaking monolithic systems into modular services, migrating workloads to hybrid, multi-cloud environments, and embedding automation across development, testing, and deployment pipelines. These factories delivered speed, scale, and operational efficiency.

However, the digital factory model was built around a fundamental assumption: humans remain the primary owners and drivers of enterprise workflows. Applications are designed and built for user interaction, automation pipelines are deterministic, and typical business processes are structured around form-based engagement. We are now at a structural inflection point.

The rapid evolution of large language models, reasoning engines, and knowledge graphs has enabled a refreshed architectural paradigm: the System of Agency. In this model, AI agents become active participants in enterprise workflows. They can interpret intent, reason based on contextual knowledge, orchestrate service- and API-based workflows dynamically, and execute goal-driven processes within certain guardrails.

Enterprises are evolving across three interdependent systems:

- Systems of Foundation (core cloud, infra, network, and security)
- Systems of Intelligence (Data, Analytics)
- Systems of Agency

This vision is expanded into a detailed blueprint for transformation, centered on three defining layers of the system of agency:

- **X-perience** – Experience as intent and conversation
- **Faceless Applications** – API-native, headless business logic
- **AgentifAled Business Processes** – Goal-driven orchestration

Together, these layers move enterprises beyond just digital acceleration and toward architected and governed autonomy.

Figure 1

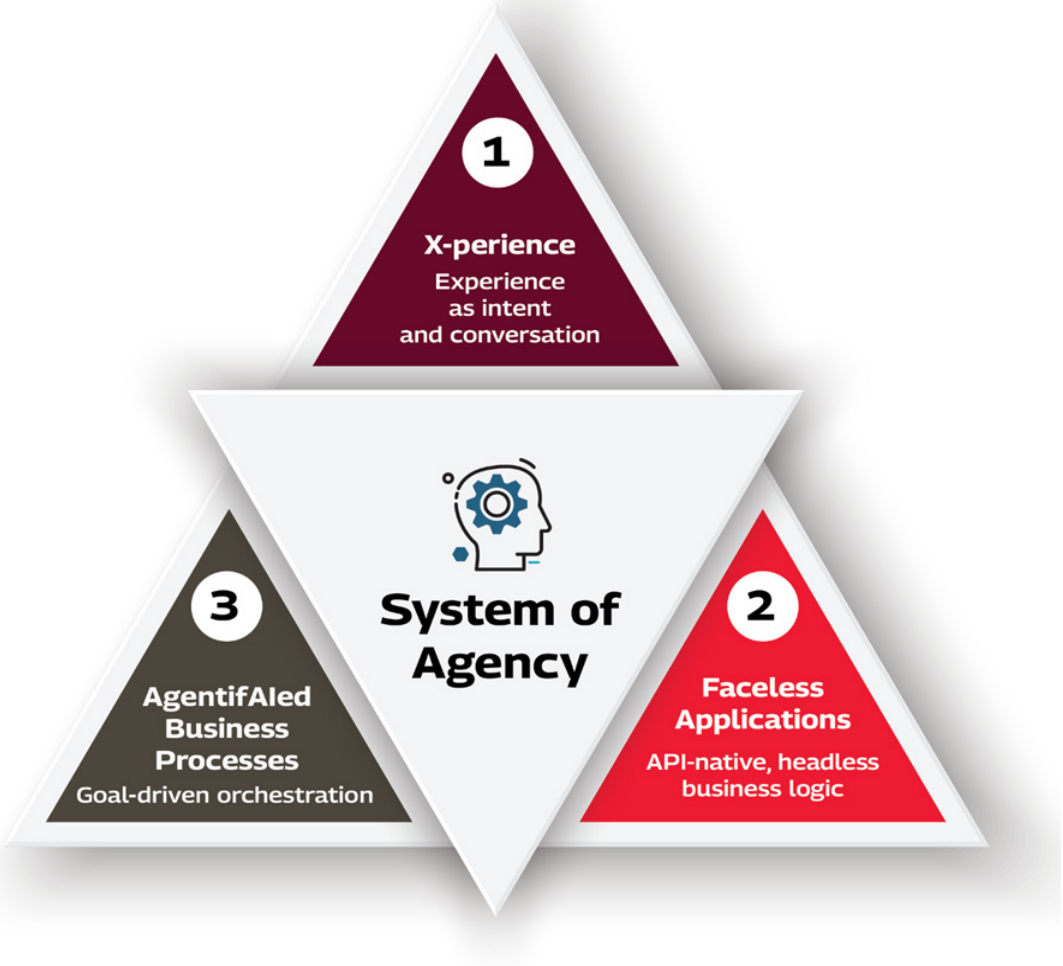

THE STRUCTURAL LIMITATIONS OF DIGITAL FACTORIES

Digital factories transformed the software engineering discipline. They introduced cloud-native architectures and foundations, agile practices, microservices-based design, DevSecOps automation, and platform engineering best practices. Release cycles came down from months to days. Technical debt was reduced. CI/CD pipelines got institutionalized along with automated quality gates. However, digital factories optimized delivery mechanics, but did not provide the execution intelligence.

Applications remained user-triggered. Workflows were pre-orchestrated. Automation scripts still remained deterministic. Artificial intelligence was embedded as a feature rather than designed as an architectural layer, and with an AI-first mindset.

Notably, advanced hyper-automation claims relied on rules-based engines or robotic process automation. These technologies improved efficiency but failed to introduce adaptive reasoning.

The next transformation wave differs in nature. It is not concerned with how quickly code is deployed but how intelligently systems operate.

The question is no longer “How fast can we build and release?” It shifts to: How intelligently and autonomously can we execute?

THE EMERGING ENTERPRISE ARCHITECTURE AND STACK

The emerging foundational architecture articulates a layered evolution of the enterprise asset stack

Systems of Foundation

This layer provides the required stability, security, scalability, and resilience. It includes hybrid and multi-cloud infrastructure, containerized platforms, API gateways, microservices architecture, and Zero Trust security modeling. Digital factories have significantly strengthened this layer, but certain aspects still need further evolution, such as security models and agent access management.

Systems of Intelligence

This layer is crucial for reasoning; here is where machine learning platforms, LLM farms, knowledge graphs, ontology models, and governance frameworks reside. This layer enables context-driven reasoning, semantic mapping, and enterprise-wide inference and insights capabilities.

Systems of Agency

This is the most dynamic layer, introducing AI agents capable of autonomous orchestration. It leverages conversational interfaces over legacy form-based UIs, multimodal engagement, and goal-driven, intent-aware planning engines to move from automation to agency.

The transition from systems of intelligence to systems of agency defines the boundary between assistance and autonomy.

DEFINING THE SYSTEM OF AGENCY

A system of agency is an architectural design and construct in which AI agents operate as intent and goal-driven orchestrators within the enterprise boundaries.

Unlike traditional workflow engines that follow predefined sequences, agents:

- Interpret objectives and intent expressed in natural language formats
- Plan for execution paths and workflows dynamically based on reasoning
- Invoke faceless services through APIs
- Monitor outcomes continuously and adapt paths and workflows based on signals
- Escalate exceptions under governed ‘policy as code’ driven guardrails

This model shifts enterprise design from deterministic workflows to goal-centric paths. Business processes are no longer handled through static pipelines; they are dynamically adapted strategies executed within controlled boundaries and guardrails to meet business goals.

The system of agency rests on three core layers: X-perience, Faceless Applications, and AgentifAIed Business Processes.

X-perience: Experience as Intent not Interface

Generally, traditional experience design prioritizes UI driven navigation. Users interact with forms, screens, and dashboards, explicitly triggering backend business logic.

X-perience redefines user interaction as the expression of intent.

Conversational interfaces are becoming mainstream, and chat-to-app patterns enable users to articulate goals and intent rather than manually trigger and complete workflows.

In this model:

- The interpretation takes precedence, while the interface becomes secondary
- Context persists across user sessions and also adapts based on feedback signals
- Agents dynamically generate workflows and interaction pathways.

Thus, experience becomes continuous, adaptive, and contextual.

Faceless Applications: Logic Without Interface Coupling

Faceless Applications expose business logic through secure, composable APIs independent of UI layers. This decoupling is essential for agent orchestration.

Agents require modular service endpoints to act autonomously. Legacy systems must be refactored to:

- Support API-first service design
- Enable event-driven communication
- Implement identity and access management frameworks for both humans and AI agents
- Operate under Zero Trust security models

Faceless architecture evolves core applications from user interaction-centric systems into orchestration-ready services.

AgentifAIed Business Processes: From Automation to Goal Execution

AgentifAIed processes differ from traditional deterministic automation. These processes begin with goals and objectives rather than predefined workflows.

The core vision for this is to have use cases, for example: AI-driven compliance mapping, agent-led onboarding, and intelligent financial planning

In an agentified environment:

- A business goal is identified and defined in natural language.
- The agent reasons using contextual knowledge graphs and world maps
- The agent creates a plan and strategy for execution to accomplish the goal.
- APIs are invoked dynamically based on agent-led orchestration.
- Outcomes are validated and refined.

This architecture enables adaptive execution across multiple industries, from BFSI regulatory and compliance monitoring to manufacturing for predictive optimization and telecom billing lifecycle orchestration; the possibilities are limitless.

THE TRANSFORMATION JOURNEY: FROM DIGITAL FACTORY TO SYSTEM OF AGENCY

5A Framework — Architecting the Autonomous Enterprise

Transitioning to an agency-based system is not a linear modernization program. It is a journey into a multi-year architectural and organizational shift that requires thought-through and phased execution. This transformation is a journey that reshapes the enterprise's operating logic.

The transformation can be divided across five distinct phases:

Phase 01
Architect the Foundation
Creating an Agent-Ready ecosystem

The first phase ensures that the Systems of Foundation are stable, resilient, composable and scalable enough to support dynamic orchestration.

Enterprises must:

- Rationalize the hybrid cloud environments to eliminate any shadow infrastructure
- Standardize API gateways and microservices frameworks
- Implement Zero Trust identity and access management architecture for both human and agent actors
- Consolidate data into a data lake or federated data platforms with governance overlays.

To support AI agents as first-class citizens, identity and access management frameworks should fundamentally change. This covers credential lifecycle management, scoped access policies, and an auditable execution trail for agents.

Phase 1 core outcomes includes:

- Lower API and service fragmentation
- Secure and traceable (audit-ready) machine-to-machine authentication for agent-driven action
- Data accessibility for reasoning engines and knowledge graphs
- Elastic and scalable Infrastructure for inference workloads

Without this foundation, agentification efforts may remain brittle and siloed.

Phase 02
Augment Intelligence
Developing Enterprise Reasoning through Intelligence Enablement
Systems of Intelligence are constructed in the second phase.

Here, enterprises introduce:

- Governance-controlled LLM platforms
- Knowledge graphs that depict technical and business ontologies, layers of semantic metadata
- Policy engines for enforcing guardrails and constraints

Data warehousing is not the end of this phase. For reasoning, it is necessary to represent enterprise information in a machine-interpretable manner.

Phase 03
Abstract Applications
Faceless Refactoring – Decoupling Logic from Interface
Now, after Phase 2, when the foundation infrastructure and the intelligence ecosystem are built, enterprises need to refactor applications to become accessible to agents

This phase focuses on:

- Decoupling User Interface layers from backend business logic
- Transforming monolithic business services into composable APIs
- Enabling event-driven design and communication patterns
- Implementing service mesh architecture for observability

The objective of this phase is not to eliminate UI. It is UI independence.

Legacy workflows that depend heavily on human-triggered screen-based actions need to be abstracted into service endpoints. This transformation requires domain-driven design refactoring and elimination of hard-coded business logic embedded within user interface layers.

Phase 3 transforms the enterprise application landscape into an orchestration-ready ecosystem. This is where the journey of autonomy begins.

Phase 04
Agentification
Introducing Domain-Specific and Goal-Driven Agents
Agentification phase is where controlled AI agents are infused into defined business domains.

Instead of deploying enterprise-wide autonomy immediately, organizations need to begin with bounded contexts and domains. Some examples:

- Regulatory and compliance guardrails enforcement (Business and IT)
- New customer onboarding (Business)
- Predictive and preventive maintenance (IT Ops)
- Automated testing pipelines (SDLC)

Domain agents need to be introduced to explicit governance based on guardrails:

- Clearly defined objective and scope
- Constrained and entitlement-based API access
- Human-in-the-loop validation and verification checkpoints as needed
- Continuous monitoring dashboards, logging, and audit trails of agent actions

Agents are designed with memory layers, context-persistence mechanisms, and policy-enforcement engines.

Metrics introduced in this phase include:

- Human intervention frequency
- Decision accuracy rates
- Latency per goal execution
- Token consumption and efficiencies

Agentification must be iterative. Guardrail effectiveness and reasoning refinement happens through early controlled deployments.

Phase 05
Autonomous Enterprise
Scaling Agent Collaboration
The final phase extends beyond isolated agents toward coordinated multi-agent workflows.

In this stage:

- Agents collaborate across domains.
- Decision-making becomes event-triggered rather than user-initiated.
- Governance shifts toward behavior monitoring rather than access restriction.

For example, in telecom, onboarding agents may collaborate with churn-prediction and billing-optimization agents to continuously manage customer lifecycle value.

- Enterprises measure autonomy maturity through:
- Reduction in manual workflow orchestration
- Cross-agent coordination effectiveness
- Outcome-based KPIs rather than task completion metrics

The enterprise transitions from reactive execution to adaptive optimization.

Figure 2: Journey Roadmap

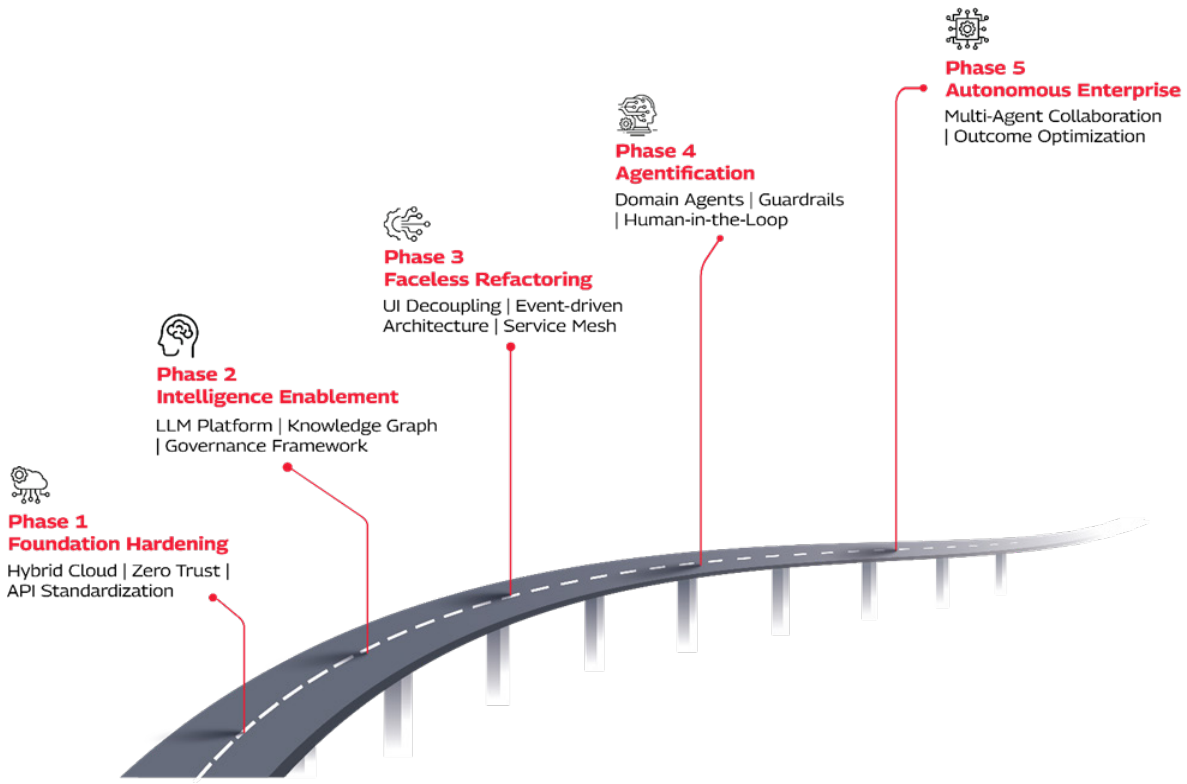

ORGANIZATIONAL EVOLUTION

Architecting autonomy requires parallel operating model transformation.

Roles evolve:

- Developers become Agent Engineers
- QA becomes continuous validation oversight
- Architects design autonomy frameworks
- Identity teams to manage both human and AI actors

Governance evolves from static compliance enforcement to dynamic behavioral supervision.

Leadership alignment and AI skilling—explicitly referenced in the foundational architecture. Digital factories to the system of agency will be strategically important to the enterprise.

ECONOMIC REFRAMING

Digital factories optimize release velocity and infrastructure cost.

Systems of agency optimize:

- Cost per autonomous decision
- Cost per orchestrated workflow
- Token economics per goal execution
- ROI per reduction in manual intervention

Financial models must account for inference elasticity and governance overhead.

ARCHITECTING AUTONOMY

The transition from digital factories to systems of agency represents a structural inflection in enterprise architecture. Where digital factories improved speed, systems of agency redefine intelligence.

Enterprises that successfully integrate:

- Intent-driven X-perience
- API-native faceless services
- Goal-driven agent orchestration

This shift merely does not modernize applications—it engineers adaptive, self-optimizing ecosystems capable of responding to change in real time.

The future enterprise is not faster. It is autonomous.

ABOUT THE AUTHOR

Kunal Purohit
President – Next Gen Verticals, Tech Mahindra
Kunal Purohit leads Next Gen Verticals at Tech Mahindra, driving AI-first, cloud, data, and cybersecurity-led transformation through industry-specific platforms and applications. He has been instrumental in advancing the company's AI-led positioning, scaling cognitive operating models, and expanding platforms such as TechM Orion (for Agentic AI) and Project Indus, a foundational large language model for Indic languages. Previously, as Chief Digital Services Officer, he led Tech Mahindra's digital solutions and capability units, strengthening global partnerships and accelerating next-generation service innovation.

Prior to Tech Mahindra, Kunal spent nearly two decades at HCL Technologies, where he led the Digital & Analytics business in Europe. Now, Kunal serves on the Mahindra Group Technology Council and has been recognized among India's top tech trailblazers.

Anjali Chhabra
Global Practice Head - Agentic Development & Modernization Services, Tech Mahindra
Anjali Chhabra is a technology leader with over 25 years of experience in the IT engineering services industry, specializing in optimizing software development lifecycle (SDLC), DevSecOps transformation, application modernization, and quality engineering services. She has worked with large enterprise customers across BFSI, telecommunications, RCG, technology, and other industry verticals. She has a proven track record on DevSecOps transformation initiatives, application modernization (on-prem and cloud adoption) consulting, automated delivery pipeline framework design set-up for large enterprises, community of practice (COP) set-up for various niche areas like microservices and containerization, observability and SRE, digital integrations, cloud native engineering, and quality engineering.

TRENDS REPORT

How Agentic AI Reshapes Your Enterprise

A Framework For Evaluating Impact, Readiness, And Risk

October 15, 2025

By Leslie Joseph, Sam Higgins with Frederic Giron, Chiara Bragato, Bill Nagel

FORRESTER®

Summary

Agentic AI is entering the enterprise faster than leaders can assess its implications. Most conversations focus on capabilities or pilots rather than what changes when agents begin coordinating, deciding, or acting autonomously. This report introduces a four-mode decision-making framework, distinguishing between architectural factors and deeper organizational shifts, to help leaders classify agentic systems' impact and make well-informed investment and operational decisions. It offers a practical lens for cross-functional decision-makers to evaluate where real operating model change is happening and what it takes to make that change safe, effective, and scalable.

There's No Clear Lens For Evaluating Agentic AI's Operating Model Impact

Agentic AI is entering the enterprise via copilots, workflow agents, and embedded orchestration layers. But while the technical capabilities are advancing, most enterprise leaders lack a means by which to evaluate how these systems reshape work and what value, business capability, structural, architectural, cultural, or governance changes they demand. The dominant conversation around agents is still framed in terms of technology and its functionality, interoperability, and demos, not its operating model implications. As a result, firms struggle to distinguish productivity enhancements from the deeper coordination, control, and execution shifts required for true operating model reinvention and radical business model innovation (see Figure 1).

Figure 1

The Payoff Starts With Task Automation At Scale As A Stepping Stone To Reinvention

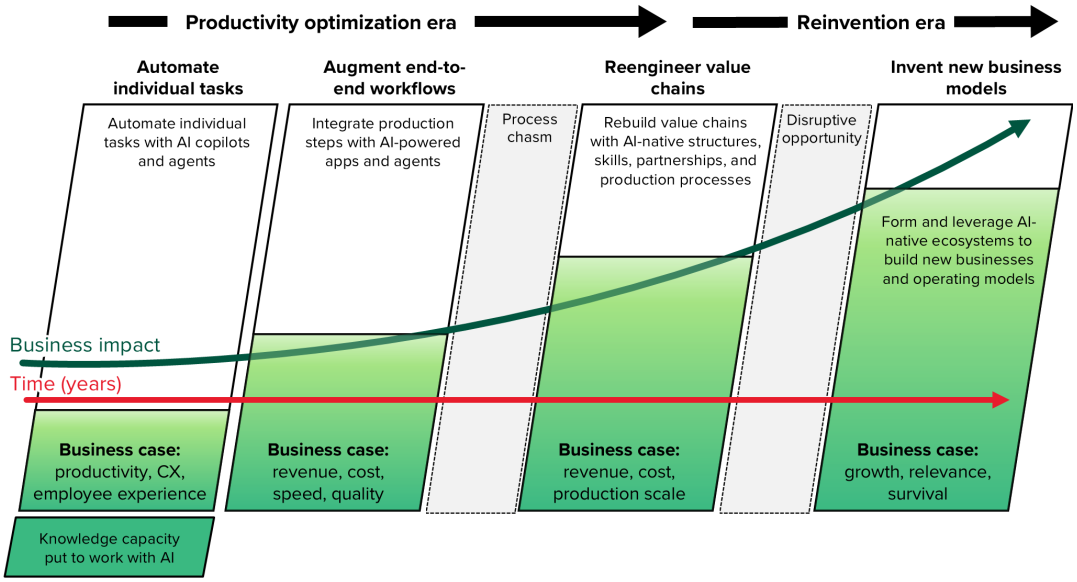

The Multiyear Journey To Agentic AI Requires A Clear Decision-Making Framework

Without a common vocabulary or evaluative framework, agent deployments remain fragmented, ambition outpaces readiness, and value capture stalls. Advances like this are not unique: In the 1990s, PCs put information power and action on every desktop; in the two succeeding decades, web and mobile interfaces put engagement and insight in

every pocket. Each advance empowered people with information, unlocked immediate connections, and unleashed waves of innovation and productivity improvement — but only after the path from the old to the new was understood. Today, enterprise leaders lack a clear lens to evaluate agentic AI's impact and create an effective roadmap because:

- **They focus more on what agents do than how they change work.** Vendors emphasize the benefits of what agents equipped with reasoning chains, tools, and autonomous workflows might do. But these use-case-based narratives rarely address how enterprises should adapt to accommodate agentic behavior. This emphasis on technical capability obscures consequential questions of how agentic systems change who makes decisions, how work is sequenced, and where accountability sits. Treating agents merely as more intelligent software leads firms to underestimate the design and governance shifts required to integrate them effectively, leaving business and tech teams to experiment with no clarity as to whether they're enabling isolated productivity tools or initiating transformational change. Without an operating model lens, firms risk misclassifying use cases and mistaking depth of capability for depth of impact.
- **There's no shared vocabulary for scaling agentic AI.** Enterprises pursuing agentic AI often struggle to align teams on common goals, as each function frames agent use differently. Data science teams may focus on model design, platform teams on orchestration, and business units on productivity or cost outcomes. One group may see a use case as augmentation, another as automation, and a third as a prototype for business model reinvention. This terminology sprawl creates friction in portfolio planning, investment decisions, and architecture design. Agentic capability often spans domains of enterprise impact across task-level augmentation, workflow execution, cross-domain orchestration, and operating model reinvention. Leaders need a structured vocabulary that links agentic capability to this impact; without it, initiatives fragment and strategy execution slows.
- **Leaders misjudge the readiness needed to move from pilots to structural impact.** Agentic AI projects can stumble not because the technology underperforms, but because firms misjudge the preconditions for success. Pilot use cases often benefit from manual scaffolding, direct human supervision, and narrow scope, but attempts to scale fail when critical dependencies like system observability, real-time data integration, escalation paths, and user trust are missing. Leaders mistake promising prototypes for scalable solutions and are surprised when agentic systems break down under real-world variability. To succeed, enterprises must assess both technical feasibility and the architecture, governance, and role design needed

to support higher-impact modes of agent behavior. Without that lens, scaling is guesswork.

- **Scaling agentic AI without systemic clarity creates risk.** When firms deploy agents without understanding the change they're introducing, they fall into predictable traps. Some overreach, extending autonomy to sensitive workflows without the right oversight, leading to breakdowns in trust or compliance. Others undercommit, treating agents as tactical enhancements and missing the chance to drive structural gains in speed, adaptability, or insight. In both cases, investment outpaces benefit. Worse, disconnected initiatives multiply across teams, leading to redundant tools, overlapping roles, and inconsistent controls. Leaders need a framework that distinguishes superficial productivity from systemic change and provides a roadmap for sequencing investment, designing safeguards, and targeting durable value.

Agentic AI Drives Four Distinct Modes Of Operating Model Impact

Agentic AI at scale reshapes enterprise value not with any one capability, but the degree to which it alters how work is structured, coordinated, and executed. Our framework identifies four distinct modes of agentic impact, from narrow task-level augmentation to full operating model reinvention. These are not stages in a maturity model — they're different kinds of impact needing different enabling conditions, risk controls, and architectural commitments. Understanding which mode a given initiative falls into can help leaders align investment with ambition, clarify ownership and risk posture, and avoid mistaking isolated gains for systemic transformation. This framework also creates a shared vocabulary across business and tech functions, helping enterprises coordinate decisions and scale with coherence. Each mode is analytically distinct but often coexists with others; a single domain may contain all four. Tech leaders must learn to distinguish them and intervene, govern, or evolve them deliberately.

- **Mode 1: augmentation.** In this mode, agents augment human roles by accelerating existing tasks, improving accuracy, or expanding access to insight. These deployments typically appear as copilots, embedded recommenders, or assistants that support but do not replace human workers or decision-makers. Process logic, role boundaries, and performance expectations are unchanged; value is created via improvements in incremental efficiency, decision-making, and experiences, often within a single team or function. Cisco used agents to reduce renewal times

and make its risk mitigation plan more effective. This mode requires minimal organizational change but depends heavily on trust, usability, and local integration; while it's often the fastest to implement, it's also the easiest to overestimate. Without structural redesign, the long-term impact remains constrained. Leaders should view this mode as a low-risk entry point rather than an end state.

- **Mode 2: automation.** Here, agents assume broader, often complete responsibility for narrow, rule-based workflows like triaging service tickets, provisioning access, or processing low-risk approvals. These functions were previously executed by humans but followed predictable logic and structured decision pathways. The agent fully automates these steps, often with exception handling or human override mechanisms in place. Epicor automated its RFQ process by implementing an agent that extracts pricing, lead times, and part details from supplier emails. The broader process remains intact; upstream inputs, downstream consumers, and governance structures are unchanged. This mode offers measurable efficiency gains but requires clearer accountability, traceability, and fallback design. It also shifts the human role from executor to supervisor, introducing new forms of cognitive and operational overhead. Scaling this mode successfully demands confidence in both performance boundaries and escalation logic.

- **Mode 3: enrichment.** In this mode, agents operate across traditional process or domain boundaries, orchestrating decisions, handoffs, and execution in response to real-time signals. Unlike automation, which executes fixed logic, enrichment adapts dynamically to predictable changes in the business environment; the structure of the workflow itself is no longer static. This improves speed, flexibility, and resilience but increases complexity and system dependency. It requires robust observability, cross-platform interoperability, and new forms of governance to monitor agent behavior and outcomes. Commonwealth Bank of Australia's creation of Lumos, a multiagent workflow, increased the number of applications it migrated to the cloud from 10 per year to more than 20 in a quarter. Organizationally, this mode often creates friction: Agents disrupt legacy ownership models and role boundaries, requiring deliberate change leadership. Done right, this unlocks enterprise agility; done poorly, it creates fragmentation and confusion.

- **Mode 4: reinvention.** This is the hardest yet most consequential form of agentic impact. Agents aren't executing tasks or coordinating workflows; they're operating as digital counterparts to human business functions. They have goals, make decisions, learn, and act continuously within defined boundaries. These agents can negotiate with external systems, reconfigure processes end to end, or optimize in real time. The operating model no longer relies solely on humans or predefined

logic. This demands platform-level investment, real-time data infrastructure, and a redefinition of accountability, control, and escalation. Human roles shift to strategic oversight, policy definition, or exception intervention. Value arises from adaptability, scalability, and emergent capability. At the same time, risk exposure increases. Enterprises that reach this mode without architectural and organizational readiness risk creating opaque, uncontrollable systems. Leaders should approach this mode with clear boundaries, strong safeguards, and aligned incentives.

Classify Agentic Systems And Align Decisions With Impact

Agentic AI initiatives vary widely in form, ambition, and impact; too often, enterprises approach them with inconsistent assumptions and mismatched expectations. To make agentic systems safe, scalable, and effective, leaders need a clear way to evaluate the change they're introducing. The four-mode framework is a tool for evaluating how agentic systems reshape enterprise business capabilities and structure (see Figure 2). Neither maturity model nor capability taxonomy, the framework helps classify agentic initiatives by the kind of change they introduce; then, what each initiative requires in terms of architecture, accountability, and governance can be assessed. The framework supports a range of enterprise decisions, from pilot evaluation and investment prioritization to architectural design and readiness planning, and offers a shared vocabulary for aligning stakeholders across domains. Most importantly, it prevents teams from mistaking a working demo for a scalable system or a smart assistant for a structural transformation. Leaders should treat the framework as a diagnostic lens that links ambition to design and operating discipline within interactive operating model optimization. They should:

- **Identify the kind of systemic change the agent introduces.** Classify the system according to the four structural modes of agentic impact. Ask if the agent supports human work, executes predictable logic, dynamically reconfigures workflows, or operates independently within policy boundaries. Focus not on surface functionality but on where decision rights shift, how workflows adapt, and what role humans still play. Many teams misclassify mode 3 systems as mode 2 because they focus on the agent's tasks rather than on its influence over workflow logic or role coordination — like Bank of America's AI Erica, which has expanded beyond its initial scope and become embedded in the business. The mode you select does more than define a label; it shapes which risks to govern, what dependencies to address, and how to measure success. Treat mode classification as the first structural design decision, not an afterthought.

- **Use the dimensions to surface what must change and what must stay.** Once you've identified the mode, use the framework's dimensions to evaluate how the system interacts with enterprise structure. These include workflow structure, human role, decision logic, integration depth, risk exposure, and governance needs. Each dimension surfaces practical considerations: What must be redesigned, monitored, or delegated? What fallback mechanisms or policy constraints must exist? This step could reveal critical misalignments such as deploying a mode 3 orchestration agent into an architecture built for mode 2 automation or assigning oversight responsibilities without role clarity. Working dimension by dimension forces a cross-functional conversation and helps ensure that architectural decisions, operational roles, and risk postures evolve together.

- **Assess organizational readiness to support the chosen mode.** Ambition often outpaces readiness. Many enterprises pilot agents that act with autonomy or adapt workflows dynamically without the architectural observability, policy enforcement, or control mechanisms required to sustain them. While this doesn't mean that you shouldn't explore advanced modes, it does mean that agent behavior must be grounded in enterprise readiness: escalation paths, data quality, cross-system orchestration, or the availability of trusted feedback loops. Readiness is not a binary checklist; it's a condition to evaluate continuously. The framework helps reveal when a firm is supporting what looks like mode 3 or mode 4 behavior with mode 1 infrastructure and oversight — a gap where failure modes often emerge.

- **Use the framework to coordinate, not classify.** The framework's greatest value is as a lens shared across disciplines (see Figure 3). AI teams see model performance; platform teams see workflow plumbing; business leaders see outcomes and risk. The four modes help unify these perspectives around structural impact. When used in planning, roadmap reviews, or architecture governance forums, the framework facilitates informed debate: Is the agent actually changing how work is done or just accelerating it? What does that require from adjacent systems or teams? What new failure points or design levers must we acknowledge? Used this way, the framework becomes a mechanism for cross-functional alignment, more than a mere diagnostic, as agentic systems scale.

Figure 2

Four Modes Of How Agentic Reshapes Enterprises

Dimension	Mode 1: augmentation	Mode 2: automation	Mode 3: enrichment	Mode 4: reinvention
Primary agent role	Assistant or recommender	Task executor	Dynamic orchestrator	Autonomous goal-seeker
Workflow structure	Fixed, human-driven	Defined, agent-executed	Adaptive, reconfigurable	Emergent, policy-anchored
Human role	Operator	Supervisor or fallback	Orchestrator or exception handler	Strategic overseer or policy-setter
Decision logic	Agent suggests, human decides	Agent decides within narrow bounds	Agent chooses flow paths contextually	Agent initiates, coordinates, and adapts independently
Process ownership	Team- or function-specific	Shared with agent	Cross-functional, dynamically distributed	Ecosystem-level or platform-distributed
Value type	Productivity, speed, quality	Cost efficiency, cycle time	Agility, responsiveness, resource optimization	Strategic advantage, scalability, emergent capability
Integration depth	Surface-level (UI, chat, embedded UX)	Workflow-level	System-level across domains	Platform-level with continuous state awareness
Feedback and learning	Human feedback (explicit)	Rules or thresholds	System signals, KPIs, and alerts	Runtime telemetry, learning loops, strategic feedback
Risk exposure	Low (disuse, inaccuracy)	Operational risk (quiet failure)	Coordination risk (handoff, inconsistency)	Strategic/systemic risk (alignment, autonomy drift)
Governance needs	Adoption tracking, usability guardrails	Escalation paths, auditability	Behavioral observability, override logic	Dynamic policy enforcement, platform-level guardrails

How to use:
The table lets execs spot patterns, stress-test assumptions, and sequence investments. Each row isolates one critical lever or dependency that varies meaningfully across the four modes. Each column represents a distinct mode of structural change, not a maturity tier. A single enterprise can operate in all four simultaneously across domains.

Figure 3

Agentic Lens Across Disciplines

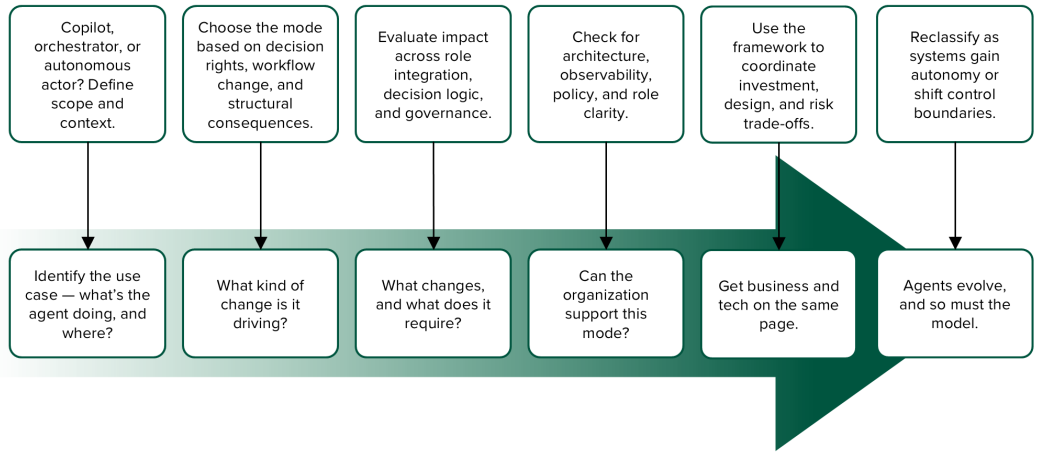

Agentic Modes Complement And Extend Existing Enterprise Models

The agentic operating model framework is not a substitute for enterprise architecture (EA), operating model design, AI governance, or automation. It complements these practices by providing a modal lens through which to evaluate how agentic systems change the way work is organized and executed within an operating model. Most current frameworks model "static" capabilities, processes, or systems that take human workers for granted and don't account for goal-seeking agents that reason, adapt, and act as people do. As enterprises deploy more advanced agents, the assumptions baked into architecture — who makes decisions, how workflows adapt, and where accountability sits — all begin to shift. This framework introduces a way to classify and assess those shifts without discarding the models that enterprises already rely on. It offers a bridge to connect high-level ambition with operational design and align strategy with control. Used well, it helps EA teams, platform owners, and business leaders clearly reason about where agents fit, what they affect, and how existing business capabilities must evolve in response. In this way, the framework:

- **Extends enterprise architecture by treating agents as persistent, adaptive actors.** EA frameworks are optimized to map systems, data, processes, and capabilities but not to model intelligent agents that make decisions, coordinate workflows, or operate continuously. Most EA approaches assume that human roles, workflow logic, and decision rights are fixed and systems execute predefined logic under human oversight. Agentic AI disrupts this, introducing software entities that act

independently, update state based on new inputs, and dynamically reconfigure coordination patterns. The framework doesn't conflict with EA — it extends it. It allows architects to classify the type of structural change a given agent introduces and surface the additional enablers, policies, or safeguards needed to support it. For EA leaders, this provides a missing vocabulary for describing and governing systems that cannot be reduced to static swim lanes or service maps.

- **Complements AI governance frameworks by adding process and operational context.** Most AI governance models focus on model behavior, fairness, transparency, and lifecycle management. They provide essential guidance on how to evaluate and manage AI risks at the system or model level but do not typically account for how AI systems reshape the broader operating environment. When agents start deciding, escalating, or orchestrating, the impact is not just model-level; it's also structural. Our framework helps bridge that gap. It surfaces what governance needs look like at each mode of agentic impact, whether that means ensuring traceability in routine automation or enforcing policy boundaries and oversight in autonomous systems. Rather than duplicating governance efforts, it enhances them by helping leaders understand which risks emerge when AI shifts from prediction to action and how governance must evolve in response.

- **Reframes automation strategy around decision and coordination.** Enterprise automation strategies have long focused on optimizing execution by automating repetitive tasks, integrating systems, or accelerating known workflows. Agentic AI shifts the focus from execution logic to decision logic; instead of asking "What tasks can we automate?", leaders must now ask "What decisions or goals can we delegate, and under what conditions?" Forrester's 2025 data shows that half of enterprise data and analytics decision-makers say their organization will focus on automation and AI over the next 12 months, so the distinction is critical. Our framework makes that shift visible. It distinguishes between systems that automate task steps (mode 2) and those that orchestrate workflows or operate against goals (modes 3 and 4). This distinction is often absent from automation roadmaps, which lump agent-based systems in with robotic process automation or workflow tools. The framework helps automation leaders reassess where autonomy actually sits, what coordination it demands, and what failure looks like when agents drive the flow.

- **Bridges the gap between digital ambition and design reality.** Digital strategy frameworks often focus on aspirational outcomes — autonomous operations, real-time responsiveness, and adaptive processes — that rarely translate cleanly into operating model design. Most firms lack a structure for reasoning about

how such aspirations affect role design, decision rights, systems integration, or governance posture. Our framework closes that gap; it connects ambition to design by classifying what kind of change each agentic initiative introduces and its implications for architecture, investment, and control. It gives strategy and architecture teams a common tool to stress-test use cases, align roadmaps, and ensure that transformation efforts are grounded in operational feasibility. This connection is critical for organizations navigating the leap from digital prototypes to scalable platforms.

FORRESTER®

We help business and technology leaders use customer obsession to accelerate growth.

FORRESTER.COM

Obsessed With Customer Obsession

At Forrester, customer obsession is at the core of everything we do. We're on your side and by your side to help you become more customer obsessed.

Research

Accelerate your impact on the market with a proven path to growth.

- Customer and market dynamics
- Curated tools and frameworks
- Objective advice
- Hands-on guidance

Learn more.

Consulting

Implement modern strategies that align and empower teams.

- In-depth strategic projects
- Webinars, speeches, and workshops
- Custom content

Learn more.

Events

Develop fresh perspectives, draw inspiration from leaders, and network with peers.

- Thought leadership, frameworks, and models
- One-on-ones with peers and analysts
- In-person and virtual experiences

Learn more.

Contact Us

Contact Forrester at www.forrester.com/contactus. For information on hard-copy or electronic reprints, please contact your Account Team or reprints@forrester.com. We offer quantity discounts and special pricing for academic and nonprofit institutions.

Forrester Research, Inc., 60 Acorn Park Drive, Cambridge, MA 02140 USA
Tel: +1 617-613-6000 | Fax: +1 617-613-5000 | forrester.com

ABOUT TECH MINDRA

Tech Mahindra (NSE: TECHM) offers technology consulting and digital solutions to global enterprises across industries, enabling transformative scale at unparalleled speed. With 150,000+ professionals across 90+ countries helping 1100+ clients, TechM provides a full spectrum of services, including consulting, information technology, enterprise applications, business process services, engineering services, network services, customer experience & design, AI & analytics, and cloud & infrastructure services. It is the first Indian company in the world to have been awarded the Sustainable Markets Initiative's Terra Carta Seal in recognition of actively leading the charge to create a climate and nature-positive future. Tech Mahindra is part of the Mahindra Group, founded in 1945, one of the largest and most admired multinational federations of companies.

For more information on how TechM can partner with you to meet your scale at speed imperatives, please visit Tech Mahindra | Scale at Speed