FROM PATTERNS TO ECOSYSTEMS:

THE EVOLUTION OF AGENTIC AI SYSTEMS



Infosys®
Navigate your next

A WORD FROM INFOSYS

Agentic AI is redefining the very core of how financial services operate and deliver value. At Infosys, our intent is to understand and guide the sector as it navigates the possibilities and challenges of a new era, one where agentic AI is poised to fundamentally change the industry's competitive landscape.

The rise of agentic AI signals a shift from traditional process automation to systems that can independently interpret complex objectives. These agentic systems go beyond mere efficiency gains; they empower financial services firms to reimagine decision-making and customer engagement at scale. For Infosys, investing in this exploration means equipping our clients and partners with the frameworks necessary to lead these changes.

Across the financial sector, banks are deploying AI agents for complex tasks such as risk evaluation and fraud detection. Insurers are using similar technologies to enhance claims management and customer interaction. These applications demonstrate the expanding role of agentic AI in daily operations, strengthening institutional agility and resilience.

Of course, this evolution brings new challenges. Successfully integrating agentic AI requires navigating complex regulatory environments and meeting heightened expectations for personalized experiences. It calls for a fundamental rethinking of processes, architecture, and governance.

We are pleased to present this Everest Group Report which offers executives frameworks to progress from pilot projects to enterprise-wide adoption of agentic AI. It focuses on identifying and scaling practical use cases, with an emphasis on effective governance. As agentic AI evolves from isolated patterns to dynamic ecosystems, this report equips financial services leaders with the insights and tools needed to shape the future: strategically, responsibly, and at scale.

70%

of global financial institutions are piloting or scaling AI initiatives beyond traditional automation

Integrating

AGENTIC AI

requires a fundamental fundamental rethinking of processes and governance.



Dennis Gada

Executive Vice President, Global Head of Banking & Financial Services, Executive Committee Member, Infosys

External Document © 2025 Infosys Limited

Everest Group®

September 2025

Agentic AI in Financial Services

Designing Trusted and Value-generating Systems

Infosys This document has been licensed to Infosys

Contents

03 Introduction

04 Agentic AI: a new paradigm for financial services

09 Prioritizing agentic AI use cases for BFS enterprises

13 Technology and architecture enablers for enterprise-scale agentic AI

21 Enterprise design strategies for agentic AI: from system patterns to multi-agent ecosystems

27 Strategic and commercial implications of scaling agentic AI

31 Conclusion

Copyright © 2025, Everest Global, Inc. All rights reserved.

www.everestgrp.com | this document has been licensed to **Infosys**

Introduction

AI adoption in Banking and Financial Services (BFS) has surged over the past 18-24 months, with over 70% of global financial institutions piloting or scaling AI initiatives beyond traditional automation.[1] Yet, most enterprise efforts continue to remain confined to task-level efficiencies, such as chatbots, document summarization, and narrow predictive models.

This shift is beginning to take root. Bank of America's Erica handles over 2 billion client interactions autonomously. Morgan Stanley's AI Next Best Action platform empowers 16,000 advisors with personalized client prompts. JPMorgan's COIN agent reviews 12,000 contracts in seconds. HSBC's real-time treasury optimizer, powered by generative AI and risk simulation, recalculates liquidity positions across markets 10 times faster than before. While these examples are largely generative AI-driven, they represent important steps in organizational readiness and signal a shift toward agentic systems. Together, they mark the beginning of a structural transformation in how financial institutions execute, serve, and govern at scale.

As regulatory complexity increases and customer expectations shift toward hyper-personalized experiences, forward-looking institutions are embracing agentic AI as the next frontier.

Agentic AI is a strategic leap. Unlike generative AI, which is prompt-based and primarily informational, or Robotic Process Automation (RPA), which is rules-based, agentic AI introduces goal-oriented systems that can act autonomously or semi-autonomously, learn continuously, and coordinate across workflows. These systems can interpret enterprise objectives, retrieve historical context, use enterprise tools, and deliver business outcomes without requiring constant human direction.

In this Viewpoint, we outline an enterprise readiness approach for adopting agentic AI in financial services. Rather than focusing on isolated use cases, the report emphasizes building a holistic foundation that integrates orchestration, memory, governance, and runtime design. Within this context, it introduces a use case prioritization framework to help institutions determine where to begin and how to scale. It also examines real-world applications across customer onboarding, compliance, capital markets, treasury, and wealth management.

1 Source: Everest Group Key Issues Survey 2025

The report also explores the anticipated rise of multi-agent ecosystems, where intelligent agents will collaborate across domains to enable distributed and adaptive execution in complex financial workflows. It includes the strategic and commercial implications for banks, insurers, and capital markets players, including how delivery models are evolving, what governance maturity is required, and which success factors separate pilots from production.

The message is clear: agentic AI is not an upgrade; it will become a re-architecture of execution. And in a sector defined by complexity, trust, and speed, it will be the institutions that build for scale, embed control, and lead with intent that define the next financial services era.

Agentic AI: a new paradigm for financial services

Agentic AI is redefining how enterprises pursue intelligent transformation. Unlike generative AI, which focuses on content creation, or traditional RPA, which automates tasks based on static rules, agentic AI is introducing systems that operate with memory, purpose, and autonomy. These intelligent agents interpret business goals, retrieve context, plan workflows, make decisions, and execute actions without constant human involvement. Importantly, in financial services, most early deployments are expected to operate in semi-autonomous or human-augmented modes, with full autonomy maturing gradually as governance structures and trust mechanisms evolve.

The enterprise significance of agentic AI

This shift will become highly relevant for financial institutions facing increased pressure to modernize. Agentic AI is more than a technological layer; it is a strategic enterprise capability. It functions as an embedded intelligence framework spanning the front, middle, and back office, enabling institutions to transition from fragmented operations to goal-driven orchestration.

Before diving into how agentic systems operate, it is important to understand how this enterprise capability has emerged. The evolution of enterprise automation, moving from static, rule-based workflows to predictive AI and now toward agent-led autonomy, is paving the way for agentic AI to become the next step in intelligent financial operations.

The following exhibit outlines this progression, contextualizing the emergence of agentic AI as the next step in enterprise maturity. It highlights the shift from rigid rule-based execution to intelligent, agent-led autonomy.

Exhibit 1: Evolution of enterprise operating models from automation to human-augmented autonomy

Source: Everest Group (2025)

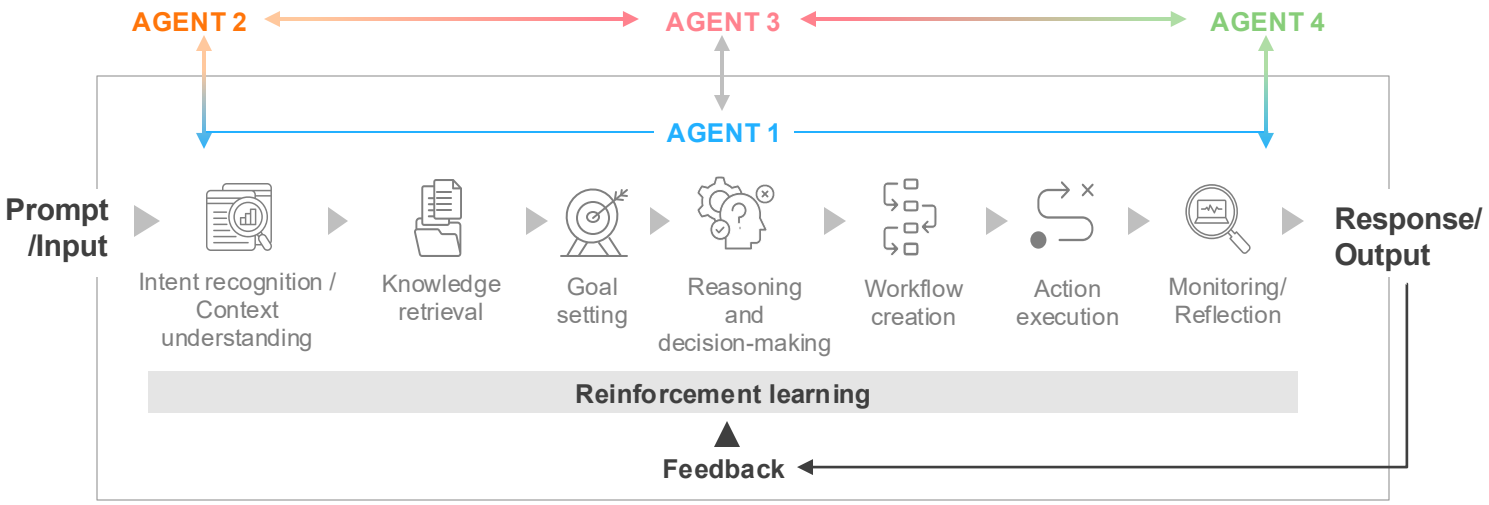
Phase 1 (2000 to 2015)	**Phase 2** (2015 to 2020)	**Phase 3** (2020 onwards)
Rule-Based Automation	**AI-Assisted engagement**	**Agentic execution**
Executes structured, repetitive tasks using hard-coded business rules. Systems lack adaptability or contextual awareness.	Supports predictive and conversational workflows that rely heavily on human oversight. AI augments decisions but does not act independently.	Drives goal-based, shaping into semi-autonomous execution. Agents operate with contextual memory, reasoning, and dynamic decision-making capabilities, while augmenting human oversight.

This evolution sets the stage for understanding the defining capabilities of agentic systems. With that context, the structure and value of agentic AI become more evident when we look at how these systems function in practice.

Agentic AI systems dynamically interpret context, retrieve knowledge, reason across tasks, and orchestrate workflows either independently or in collaboration with other agents.

Exhibit 2: How do agentic AI systems operate?

Source: Everest Group (2025)



This evolution moves AI from assistive to agentic. It transitions from enabling tasks to independently orchestrating results.

Building on this operational understanding, agentic AI systems create the most enterprise value when embedded across defined business personas in banking and financial services.

Exhibit 3 maps on how agentic AI is being modularly embedded across key banking and financial personas, aligning with role-specific goals and decisions. Agentic AI drives intelligent autonomy across BFS by embedding reasoning, orchestration, and contextual memory into persona-specific use cases spanning customers, advisors, operations, compliance, and partners.

Exhibit 3: Agentic AI in banking and financial services – a modular lens across use case types

Source: Everest Group (2025)

	Use case
Customer-facing agents Autonomously interprets customer intent and recommends optimal loan products.	Loan eligibility resolution (Retail Lending)
Operations-facing agents Detects unmatched or failed trades in real-time and suggests reconciliation paths.	Trade settlement exception handler (Capital Markets)
Partner-facing agents Handles credentialing, schema validation, and token exchange	Embedded finance integration manager (Payments)
Advisor-facing agents Continuously monitors market shifts and suggests asset reallocation.	Portfolio rebalance advisor (Wealth)
Compliance-facing agents Reviews historical underwriting decisions and identifies potential bias.	Fair lending audit bot (Risk & Compliance)

The persona mapping illustrates how agentic AI can be deployed modularly across roles, each aligned to specific goals and decisions. This view reinforces the need for enterprises to shift from fragmented workflows to intelligent, goal-driven architecture built for agility and scale.

Three structural shifts will drive enterprise adoption of agentic AI

Enterprise adoption of agentic AI will be powered by three structural changes. First, foundation models will continue to mature, with modern language models advancing in their ability to reason across steps, plan tasks, and coordinate with other agents and tools. Second, enterprise-grade orchestration will evolve to provide long-term memory, secure execution, and structured tool use across distributed systems. Finally, business demand will intensify, as financial institutions face growing pressure to deliver personalization at scale, manage real-time risk, and improve operational resilience.

Beyond technology readiness, agentic AI is beginning to unlock transformative capabilities that directly link to the structural pain points and strategic aspirations of the financial services industry. It has the potential to unify fragmented workflows by orchestrating processes across origination, underwriting, servicing, compliance, and risk, turning disconnected systems into cohesive, intelligent execution layers. Agentic AI can also enable real-time risk management and compliance by continuously scanning market data, regulatory updates, and transactional behavior, allowing institutions to adjust strategies dynamically and embed governance directly into workflows. It will further deliver personalization at scale by analyzing behavioral signals and financial history to autonomously tailor recommendations, plans, and offers across retail, wealth, and lending. In operations, it can dramatically reduce cost-to-serve by offloading complex, rule-driven tasks such as KYC, reconciliation, and document analysis to adaptive agents that improve with every cycle.

Most importantly, it will future-proof the enterprise. With a modular, goal-driven architecture, agentic AI positions BFS firms to navigate fast-evolving shifts, such as open banking, regulatory mandates, and decentralized finance, with resilience and agility that static systems cannot match. This foundational shift in architecture paves the way for a reimagined approach to value delivery, one where intelligence is embedded across every layer of execution.

Agentic AI will fundamentally reshape how BFS firms deliver value. Advisors will be supported by intelligent copilots that continuously learn from client goals and shifting market dynamics, enabling more precise, real-time engagement. Risk managers will deploy agents that proactively detect and mitigate cross-border compliance breaches, transforming risk management into a proactive, data-driven capability. Operations leaders will architect environments where workflows can evolve, self-correct, and scale autonomously without human intervention.

In short, agentic AI will separate the digitally competent from the digitally sovereign. Institutions that embed these systems today while aligning governance, architecture, and trust will build a structural advantage that scales and compounds over time.

This broader view is further supported when we look at agentic AI through the lens of enterprise functions. Use cases are now emerging across the front, middle, and back office, each focused on improving accuracy, reducing cost, and accelerating execution.

www.everestgrp.com | this document has been licensed to **Infosys**

- In the **front office**, agents support onboarding journeys, personalize customer interactions, and assist advisors with next-best recommendations
- In the **middle office**, they monitor transactions, detect anomalies, adjust allocations, and maintain compliance alignment in real time
- In the **back office**, they automate document review, manage reconciliation workflows, and perform exception handling with limited oversight

This transition from early experimentation to enterprise deployment marks a critical inflection point. As adoption expands across functions, institutions must build foundational readiness by developing platforms, governance, and collaboration capabilities. Only then can they transition toward strategic prioritization and scalable execution.

Enterprise foundations for agentic AI adoption

The path to prioritizing use cases begins with recognizing the foundational elements required for enterprise-scale adoption. These are not deep technical details, but core aspects that institutions must understand to ensure adoption evolves in a structured and scalable way rather than through fragmented pilots. Scaling agentic AI will require a coherent framework that integrates platform design, governance, and enterprise-wide coordination.

At a minimum, BFS institutions will need to prepare in four areas:

- **Architecture and platform design:** define an enterprise-level approach to agentic AI platforms, including orchestration, memory, and integration layers, to ensure scalability and interoperability
- **Governance framework:** establish oversight mechanisms, such as auditability, escalation protocols, and transparency standards, to manage agent behavior in highly regulated environments
- **Collaboration and teaming:** adapt operating models so that humans, agents, and systems work together seamlessly, with accountability and decision-making remaining aligned to enterprise objectives
- **Agent-to-agent (A2A) coordination:** develop principles for how multiple agents will interact, share context, and delegate responsibilities – capabilities that will be increasingly important even if early deployments begin with single-agent pilots

These are crucial areas for enterprises to understand early on the journey. They are not intended to slow adoption, but rather to ensure that initial initiatives can scale with governance and enterprise alignment. With these building blocks in place, institutions will then be better positioned to prioritize agentic AI use cases in a structured manner, balancing feasibility with long-term enterprise impact.

www.everestgrp.com | this document has been licensed to **Infosys**

Prioritizing agentic AI use cases for BFS enterprises

As agentic AI moves from experimentation to enterprise execution, institutions must address a practical challenge: where to begin, and how to sequence adoption. Agentic AI presents diverse applications across financial services, but these use cases' impact and implementation readiness will vary significantly. Some use cases align closely with existing processes and data infrastructure and can be scaled quickly. Others remain aspirational, requiring stronger orchestration capabilities, deeper integration, or clearer business ownership.

Now we discuss a prioritization framework that helps enterprises determine where to begin, how to expand, and what to reserve for future investment. The goal is to build momentum with feasible, high-value use cases while laying the foundation for long-term transformation.

A framework for impact and execution readiness

Everest Group assessed diverse agentic AI use cases based on two core criteria:



Strategic impact
The potential to transform customer experience, improve compliance, reduce operational costs, or create new revenue opportunities



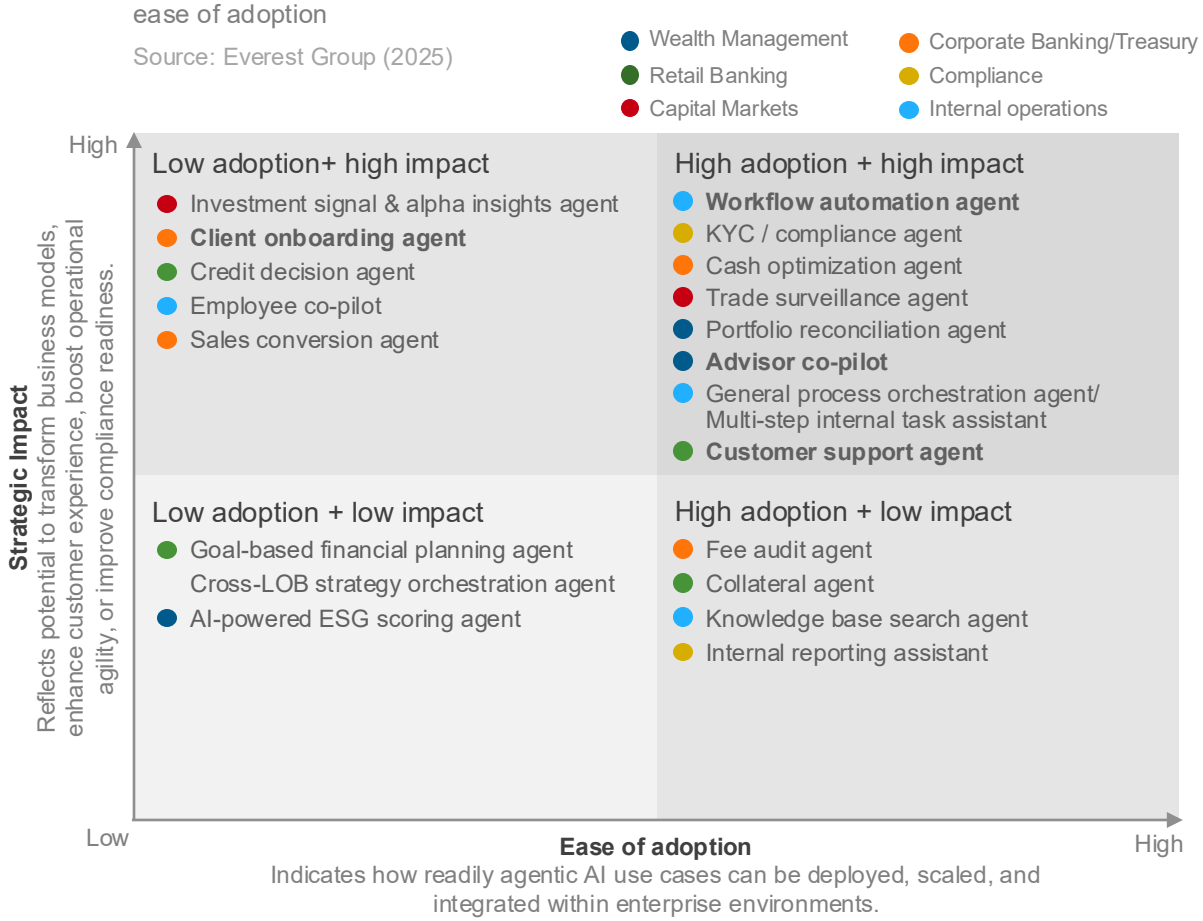
Adoption complexity
The orchestration maturity, data readiness, system integration effort, and agent runtime capability required to implement the solution

Exhibit 4 builds on the above framework to map agentic AI use cases across financial services, offering a structured view to guide enterprise adoption and investment focus.

The goal is to build momentum with feasible, high-value use cases while laying the foundation for long-term transformation.

Exhibit 4: Prioritizing Agentic AI use cases in financial services by strategic impact and ease of adoption

Source: Everest Group (2025)



Note: **Use cases marked in bold** represent frequently cited patterns across multiple real-world BFS deployments and case studies (2021–2025). These carry higher adoption confidence.

The resulting framework categorizes use cases into four action zones:

- **High-value bets (high complexity and high impact):** These are transformational use cases that require mature orchestration, integration effort, and organizational readiness. While harder to execute, they represent the intelligent, autonomous financial operations' future
- **Quick wins (low complexity and high impact):** These are lightweight, easy-to-deploy agents that deliver measurable value. Often used in targeted back-office or operational workflows, they are ideal for accelerating proof of value
- **Incremental use cases (low complexity and low impact):** These use cases improve process efficiency or accuracy but have limited scope. They are often deployed in support or administrative functions where cost savings or speed improvements are the primary goal
- **Moonshots (high complexity and low impact):** These are exploratory or long-horizon use cases that carry high execution risk. While limited in current feasibility, they can become strategic differentiators as orchestration capabilities and data maturity improve

This framework will help institutions align high-potential use cases with the operational maturity required for scale. By sequencing adoption deliberately, enterprises can maximize near-term impact while building the foundation for intelligent, agent-led execution models.

As institutions operationalize this framework, emerging deployment patterns are beginning to surface. BFS enterprises are clustering their agentic AI deployments along the very dimensions of strategic impact and execution readiness that this structure highlights.

Enterprise adoption patterns and implications

Across the BFS landscape, enterprise agentic AI's adoption is beginning to indicate a clear use case cluster, shaped by complexity, business value, and execution readiness. We have observed a few quick-win deployments, such as KYC agents, advisor copilots, and reconciliation agents, are gaining momentum in structured, well-governed environments where institutions can realize immediate gains in accuracy, speed, and cost efficiency.

The following exhibit presents real-world examples of how BFS enterprises are adopting agentic AI across functions. These deployments varying from payments and trading to compliance and product innovation - demonstrate how agentic systems are moving from pilots to live environments, delivering measurable outcomes.

Exhibit 5: Enterprise adoption of agentic AI across banking and financial services

Source: Everest Group (2025)

BNY MELLON	BANK OF AMERICA	citibank
AI "digital employees" autonomously validate payments and write code under supervision; have real user credentials.	"Maestro" agent proactively delivers trading insights from internal data without human prompting.	Devin agent pilots at Citi demonstrate how autonomous digital employees can draft research, generate code, and summarize documents under enterprise supervision
J.P.Morgan	**VISA mastercard**	**Goldman Sachs**
LOXM autonomous agent conducts equity trades using deep reinforcement learning—end-to-end.	Agentic commerce pilots where AI agents autonomously execute digital payments using tokens and personal rules.	AI Assistant + Devin agent autonomously draft research, write code, and summarize documents.

More examples of agentic AI deployment are starting to emerge across leading institutions. Large banks such as Scotiabank and Morgan Stanley are already demonstrating the value of deploying these agents in production, particularly in high-volume or regulatory workflows. In the next tier, high-value strategic bets, including client onboarding orchestrators, alpha signal agents, and cross-channel conversion engines, are being piloted by digital-forward institutions, such as Betterment and RBC Capital Markets. These use cases offer transformative upside but demand robust orchestration environments, cross-functional coordination, and deep business ownership. Simultaneously, several institutions are rolling out incremental agents, such as fee audit bots and internal search assistants, to target specific pain points in administrative or support functions, with JPMorgan Chase driving measurable efficiency gains. Finally, moonshots, such as ESG scoring agents or agents that orchestrate across lines of business, remain in exploratory phases due to their technical and governance complexity.

However, these applications represent intelligent operating models' future, where decision-making and execution are continuously distributed across an enterprise-wide network of collaborating agents. This layered adoption pattern reinforces a key insight. Agentic AI will need to be sequenced deliberately, with institutions balancing near-term returns with long-term architectural transformation.

The following patterns observed in early implementations offer practical signals for institutions seeking to scale with purpose:

- First, **impact-driven design is non-negotiable.** Agents are being introduced in high-volume, outcome-sensitive workflows and not in experimental sandboxes. Whether it is accelerating client onboarding, streamlining compliance reviews, or augmenting financial advice, the common thread is that these agents are built to reduce friction, increase precision, and enhance speed where it matters most.
- Second, **agentic AI is augmenting and not replacing human roles**. In current deployments, agents are functioning as digital copilots, assisting advisors, supporting analysts, and automating sub-tasks in operations and risk. This augmentation model drives adoption and ensures that accountability, judgment, and escalation remain human-led.
- Third, **governance is foundational and not optional**. Every effective agent deployment incorporates oversight mechanisms, including intent logging, escalation protocols, and decision traceability. These controls are particularly key in regulated functions, such as credit, trading, and compliance, where explainability and reversibility must be built into the system's every layer.
- Finally, **cross-functional agents are still rare**, with most current implementations confined to a single business function or platform. The vision of agents that coordinate across advisory, compliance, and servicing is taking shape but remains aspirational for most institutions.

www.everestgrp.com | this document has been licensed to **Infosys**

As BFS enterprises transition from early pilots to more scaled implementations, the focus has shifted from validating agentic AI's potential to understanding what makes deployments successful. The imperative is no longer just about sequencing use cases; it is about architecting the right foundation. Moving forward, institutions must now build the enabling technologies and enterprise architecture required to embed agentic systems securely and at scale.

Technology and architecture enablers for enterprise-scale agentic AI

For financial institutions, deploying agentic AI will not be a plug-and-play exercise. It will represent a foundational redesign of how enterprise systems operate, integrate, and reason. Traditional automation and predictive AI were built on fixed rules, static thresholds, and siloed infrastructure. Agentic AI will demand something that is fundamentally different, with a dynamic, modular, and memory-enabled architecture that allows autonomous systems to pursue goals, coordinate across functions, and operate within enterprise guardrails.

This re-architecture will need to support both performance and trust. Agents in financial services will not be isolated chatbots; they will evolve into entities that monitor risk, engage clients, or trigger transactions. To operate safely and at scale, agentic systems must be built on a technology stack that ensures transparency, compliance, control, and adaptability. Enabling this shift will require five core architectural layers, each of which is distinct in enabling enterprise-grade autonomy.

Now we discuss the core technology layers required to operationalize agentic AI and explain how enterprise teams can design an architecture that balances autonomy with control.

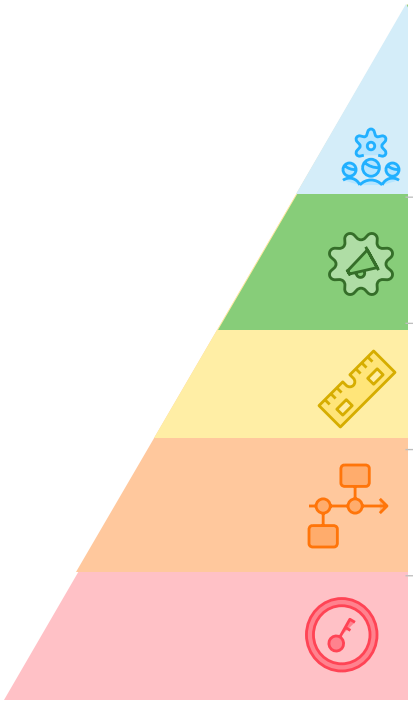
Key layers that define the architectural stack for agentic AI

Agentic systems will require a modular architecture that enables agents to interpret intent, plan and sequence tasks, interact with tools, and improve performance over time.

Exhibit 6 outlines the foundational architecture needed to operationalize agentic AI in financial services, detailing key system layers and supporting technologies required for scalable, secure, and goal-directed execution.

Exhibit 6: Architectural capability stack for enabling agentic AI in financial services

Source: Everest Group (2025)



	Key capabilities	Supporting technologies
Governance and control Ensures transparency, trust, and human oversight	Policy enforcement, explainability, audit trails	NeMo Guardrails, enterprise observability platforms
Tool Integration Connects agents to enterprise systems and APIs	Tool invocation, task chaining, I/O handling	OpenAI Function Calling, Salesforce Agentforce, UiPath connectors
Memory layer Maintains context across workflows and sessions	Long-term recall, dynamic context construction	Pinecone, Weaviate, LlamaIndex
Orchestration layer Plans, sequences, and executes multi-step tasks	Goal decomposition, task planning, multi-agent workflows	LangChain, AutoGen, Semantic Kernel
Agent runtime Provides execution context and identity for agents	Persistent state, agent profiles, secure execution environments	Azure AI Foundry, Vertex Agent Engine, Amazon Bedrock

Note: The technologies mentioned are illustrative. The ecosystem is evolving rapidly, with new tools and capabilities emerging continuously.

Building on the architectural layers illustrated as above, the following five capabilities define how agentic systems operate, starting from goal interpretation to secure execution and governance.

- **Intent recognition and goal planning**
  At the agentic execution's core is the ability to translate high-level objectives into sequenced tasks. Whether triggered by natural language prompts, business events, or data thresholds, agents must interpret goals and break them down into discrete, actionable steps. This process requires advanced planning and reasoning models, often built on Large Language Models (LLMs) with capabilities for chain-of-thought processing, instruction following, and logic-based inference. These planning layers form the agent's cognitive engine that determines what to do, when, and why.
- **Memory and contextual continuity**
  Unlike generative AI, which operates in stateless, single-session bursts, agentic systems need persistent memory. Agents must recall prior interactions, past decisions, user preferences, and contextual constraints across sessions and workflows. This is achieved through vector databases, such as Pinecone, Weaviate, and Chroma. Additionally, temporal knowledge graphs that enable semantic retrieval and longitudinal state tracking support contextual continuity. In banking, where context defines compliance and personalization, memory is not a convenience but a core requirement for accuracy and trust.

- **Secure tool use and system integration**
  For agents to deliver outcomes, they must interact with enterprise systems by accessing client data, invoking APIs, triggering workflows, and retrieving documents. In this context, tool use refers to the agent's ability to securely invoke and coordinate external enterprise services and systems, making it a central enabler of outcome delivery. This is where tool use becomes central. Secure tool integration frameworks, such as LangChain, AWS Lambda, and Microsoft Semantic Kernel, allow agents to connect to banking systems, CRM platforms, risk engines, and data services. Each interaction must respect identity boundaries, permission protocols, and audit requirements, ensuring that agents act only within predefined authority scopes.
- **Execution environment and runtime control**
  Agents must operate in environments that can monitor, validate, and control execution in real time. This layer includes orchestration logic, sandboxed execution environments, and observability modules that capture agents' every step. Hyperscaler platforms, such as AWS Bedrock, Azure AI Foundry, and GCP Vertex AI, provide runtime environments with built-in guardrails, memory management, and identity-based access control allowing agents to execute securely across hybrid and cloud-native infrastructure.
- **Governance, auditability, and explainability**
  In financial services, autonomous decision-making must be explainable by design. Institutions must track what an agent did, why it did it, and what data or logic informed that decision. This process requires intent logging, decision path traceability, real-time override mechanisms, and post-hoc audit trails. Specialized platforms, such as Stratyfy, SolasAI, and FairPlay, provide fairness overlays and model validation for lending, underwriting, and transaction monitoring, thus ensuring agents meet regulatory and ethical standards.

Each layer must ensure safety, control, and enterprise fit while delivering consistent, reliable performance. However, technology alone cannot deliver trusted autonomy at scale – it must be complemented by process-level redesign.

Business process re-engineering as a parallel enabler

While technology provides the architectural foundation for agentic AI, realizing its full potential also requires rethinking core business processes. Business Process Re-engineering (BPR) will be a vital complement to technology enablement, ensuring that agentic systems are embedded into workflows designed for autonomy and collaboration.

For BFS institutions, this means revisiting how work is structured across functions, how accountability is distributed between humans and agents, and how processes can evolve dynamically rather than remain static. Key opportunities include:

Onboarding flows: redesigned to incorporate agent-led verification and document handling

Compliance checks: restructured to align with continuous monitoring rather than static reviews

Reconciliation processes: re-architected to enable A2A interaction and exception handling

Without process-level redesign, agentic systems risk being layered on top of legacy workflows, limiting their ability to deliver enterprise-scale outcomes. BPR and technology adoption must therefore advance in tandem to ensure intelligent execution is not only technically feasible but also operationally sustainable.

Together, technology enablers and process redesign will represent a clear break from legacy approaches. They will illustrate why agentic AI will require a fundamentally different design philosophy than traditional enterprise AI systems.

How agentic AI differs from traditional AI system

The leap from traditional AI to agentic AI is not incremental; it will be architectural. Conventional AI in BFS relies on predictive models embedded in narrow workflows, such as a model to score credit, a bot to answer FAQs, or a script to extract document fields. These systems may function well in isolation, but they will lack the flexibility to manage dynamic goals, respond to unexpected conditions, or coordinate across enterprise processes.

Agentic infrastructure, by contrast, will be built to support distributed intelligence. Agents not only automate but also observe, reason, and decide. They must be capable of:

- **Dynamic goal planning:** breaking down objectives into sequenced tasks that adapt to changing conditions
- **Runtime delegation:** coordinating across workflows and systems in real time
- **Multi-turn memory:** retaining past interactions and decisions for continuity and compliance
- **Error recovery:** identifying breakdowns and re-routing actions without human escalation

This requires an infrastructure that is modular, composable, and API-first. Agents must integrate seamlessly with dozens of enterprise systems while ensuring role-based access, data security, and real-time compliance. These needs make traditional, rules-based environments unsuitable for agent-driven execution at scale.

For BFS institutions, this shift requires a deliberate transition toward composable, API-first ecosystems where intelligence can be injected into the technology stack's every layer. It is not about replacing existing systems but about enabling agents to interoperate across them with precision and control.

These infrastructure shifts raise a critical next question: how should BFS enterprises design agentic systems that can scale from pilots to production while maintaining trust and control?

www.everestgrp.com | this document has been licensed to **Infosys**

Designing trusted agentic AI systems for financial enterprises from sandbox to scale: the real design challenge

Deploying agentic AI at scale is not simply building advanced models or integrating new tooling; it is a fundamental design challenge. Financial institutions must make deliberate choices around what roles agents play, how much autonomy they are granted, and what oversight mechanisms ensure responsible execution. In regulated environments such as banking, trust is a design principle, not a feature that can be layered in later.

The journey begins with narrow-scope agents operating in sandboxed environments or low-risk workflows.

Exhibit 7 shows how BFSI enterprises are progressing from early, narrow-scope pilots to broader agent deployment across customer-facing, operational, and high-risk environments, reinforcing the transition from experimentation to production-grade execution.

Exhibit 7: Enterprise deployments of agentic AI across BFSI, ranging from internal pilots to customer-facing systems

Source: Everest Group (2025)

- **Semi-autonomous pilots** – moderate-risk, partially supervised, not customer-facing
- **Internal operations pilots** – internal copilots, automation of back-office tasks
- **Customer-facing/high-autonomy** – agents in live, external-facing environments
- **Potential use cases–** Exploratory or early-stage agentic AI applications with limited rollout.

BANKING

BFS segment	Process domain	Use case	Enterprise example	Agentic AI tech / platform
Retail Banking	Customer engagement	Always-on customer assistant via mobile app	Bank of America	Erica (in-house NLP agent)
		Proactively predicts customer needs and flags fraud risks	Capital One	Eno AI assistant
	Loan processing	Autonomous review of commercial loan contracts	JPMorgan	COIN (Contract Intelligence platform)
	Workflow Orchestration	Back-office operations, mortgage automation	ING Bank	In-house agentic AI (pilot)
Commercial Banking / Lending	Underwriting automation	Flagging anomalies, loan data analysis	Large US retail bank (unnamed)	Automation Anywhere agentic automation platform
	Predictive analytics	Autonomous insight generation in operations	HSBC	CausaLens agentic predictive analytics tools
	Digital operations	Conversational AI agents handling internal tasks	Citigroup	Citi AI Suite- Stylus agent-Voice & doc-based copilots
	Compliance & fraud ops	Core ops automation, fraud case handling	HSBC	Enterprise agentic AI platform
	Workflow automation	Content generation, translation, and customer support	Standard Chartered	SC GPT

- **Semi-autonomous pilots** – moderate-risk, partially supervised, not customer-facing
- **Internal operations pilots** – internal copilots, automation of back-office tasks
- **Customer-facing/high-autonomy** – agents in live, external-facing environments
- **Potential use cases–** Exploratory or early-stage agentic AI applications with limited rollout.

BFS segment	Process domain	Use case	Enterprise example	Agentic AI tech / platform
Cards & Payments	Customer onboarding	AI agent guiding issuers/merchants through setup	Mastercard	Onboarding AI Agent (LLM + RAG)- Built on Databricks with human-in-loop
	Dispute & transaction intelligence	AI for commerce security, dispute handling	Visa	Intelligent Commerce AI platform
	Fraud analytics	Fraud detection and CX automation	American Express	Autonomous agents
	Payment security	Agentic AI for fraud detection & risk routing	Mastercard	Agentic Pay
Treasury & Operations	Payment processing	Digital agent verifying and processing transactions	BNY Mellon	Eliza AI platform (GPT-4 on Azure)- Agents with system login credentials
	Compliance workflows	AI compliance assistant	BNY Mellon	Behavox Quantum + autonomous agents
CAPITAL MARKETS				
Capital Markets / Investment Banking	Trade execution	Self-learning algorithm trading equities	JPMorgan	LOXM (AI engine for equities)
	Trade execution	Agentic AI for high-frequency trading optimization and adaptive market strategies	Goldman Sachs	Neural network–powered trading system with autonomous agents
	Research & sales	Market content generation, sales/trader support	BofA, Goldman, Citi, UBS, Morgan Stanley	Maestro (BofA) - GS AI Assistant - AskResearchGPT (MS) - Client360 (UBS)
Custody & Fund Administration	Ops automation	Processing reconciliations, compliance workflows	HDFC Bank, SBI, Wells Fargo	Agentic workflow automation agents - Document-to-decision orchestration
CROSS-FUNCTIONAL OPERATIONS ENABLEMENT				
	Workflow orchestration	AI across IT, fraud, PDF automation	Independent Bank	SnapLogic agents
	Enterprise-wide AI agents	Internal copilots at large scale (2000+ users) for customer experience, automate routine tasks	Wells Fargo	Google Cloud Agentspace
	Customer engagement	AI agents in onboarding and security ops	State Bank of India	In-house pilots
	Security ops center, fraud analytics	Agentic bots in fraud and threat ops	HDFC	Internal agents
	Software development	AI writing/debugging code autonomously	Goldman Sachs	Devin (Cognition Labs – autonomous coding agent)

www.everestgrp.com | this document has been licensed to **Infosys**

These deployments reflect more than isolated technology adoption. They represent a shift toward intentionally designed agentic systems, underpinned by three foundational enterprise design imperatives: the definition of clear agent purpose, the embedding of governance mechanisms, and the establishment of structured autonomy boundaries. Delivering this shift often requires not only technical redesign but also process-level changes, including how workflows are structured, how decisions are escalated, and how responsibility is distributed across teams.

Beyond technology providers focused on domain-specific workflows, the broader agentic AI ecosystem includes hyperscalers, orchestration toolkits, enterprise platforms, and emerging agent runtime environments. These players are shaping the technical foundation required to scale agent-based execution across regulated BFS environments.

Designing for purpose, constraints, and control

Designing for scale requires agents to operate with clearly defined business objectives. These goals might include reducing customer onboarding time, increasing reconciliation accuracy, or improving loan decision turnaround. Importantly, each agent must operate within predefined constraints, including what data it can access, what systems it can act upon, and when human sign-off is required. Beyond constraints, override mechanisms must also be in place. These mechanisms allow human supervisors to pause, modify, or reverse agent decisions in real time, ensuring compliance and mitigating risk.

Leading financial institutions are already adopting these principles. For instance, Morgan Stanley deployed an AI-powered Next Best Action engine that is an advisor co-pilot, surfacing client-specific recommendations while keeping final decisions in the hands of human wealth managers. Similarly, JPMorgan Chase uses its COIN platform to automate legal contract reviews, but with structured workflows for human escalation when confidence thresholds are not met. These systems demonstrate that agentic models can operate autonomously in a framework that ensures transparency, auditability, and human accountability.

As financial institutions begin operationalizing agentic AI, a small but growing technology provider ecosystem is addressing high-priority agentic AI use cases. These providers do not offer broad infrastructure or specific Agentic AI tools. They focus on embedding agentic intelligence directly into domain-specific workflows, such as lending, onboarding, compliance, and treasury operations.

Embedding governance and risk mitigation

Institutions must embed governance into every stage of agent design and deployment. This integration includes logging agent intent, documenting the logic behind decisions, setting approval thresholds, and conducting pre-deployment simulations to stress-test behavior. These measures ensure agents operate safely in high-stakes workflows while being fully auditable and reversible.

For example, before rolling out its AI-driven capital risk optimizer, HSBC subjected the system to rigorous simulation under varied market scenarios. These tests validated that the agent could manage capital buffers dynamically, while being compliant with Basel III constraints and audit-ready in real time. Such simulations are not only optional but also a prerequisite for trust in regulated environments.

www.everestgrp.com | this document has been licensed to **Infosys**

Design for autonomy, govern for trust

In summary, designing agentic systems for financial enterprises demands integrating intelligent automation, domain expertise, and governance-first architecture. This system's success will depend not on what agents are capable of but on how clearly their goals are defined, how responsibly their boundaries are managed, and how well monitored they are in real-world conditions.

This transition from design principles to operational execution introduces a pivotal requirement: the role of technology providers in enabling enterprise-grade agentic AI. As BFS institutions shift from isolated use cases to enterprise-wide deployments, they must rely on a coordinated ecosystem of infrastructure and platform providers that enable autonomy, governance, orchestration, and compliance by design.

Exhibit 8 illustrates this evolving provider landscape and highlights how hyperscalers, enterprise platforms, and domain-specific enablers are operationalizing agentic AI across BFS.

Exhibit 8: The agentic AI provider landscape for banking and financial services

Source: Everest Group (2025)

Provider category	Representative providers	Primary capabilities enabled	Enterprise impact
Hyperscalers	Google Cloud (Vertex Agent Engine) Amazon (Bedrock Agents)	• Agent runtime environments • Secure execution layers • Memory persistence • Tool orchestration frameworks	• Enables scalable agent deployment and chaining • Supports BFS use cases in service automation, underwriting, and internal ops • Ensures enterprise-grade security, governance, and observability
Technology platforms and enterprises	LangChain, Semantic Kernel, AutoGen, OpenAI (Function Calling) Salesforce (Agentforce), UiPath	• Goal-based task planning • Multi-agent orchestration • API and tool invocation • Workflow automation	• Powers dynamic agent behavior across customer onboarding, KYC, servicing, and document workflows • Bridges LLMs with enterprise systems and APIs
Core Agentic AI Technology Providers	Multimodal., Ema, lyzr, Neurofin.ai	• Multi-agent orchestration and task routing • Domain-specific agent modeling (finance, risk, compliance) • Embedded tools for memory, tool use, reasoning • BFS-specific UI/API connectors for workflows	• Powers autonomous agents in underwriting, KYC, onboarding, and compliance • Agent deployment frameworks tailored to BFSI systems • Enables composable financial workflows through no-/low-code agent builders • Aligns with regulated BFS needs (e.g., AML, credit risk, servicing automation)

With the foundational infrastructure established and ecosystem players now supporting runtime, orchestration, reasoning, and compliance, the strategic focus for financial institutions turns to enterprise design. This shift calls for deliberate design decisions around how intelligent agents are structured, governed, and coordinated to deliver cohesive execution across complex financial environments.

Enterprise design strategies for agentic AI: from system patterns to multi-agent ecosystems

Building agentic AI into financial services operations is not only a technical challenge but also a design issue. Institutions must define how agents are created, what roles they play, and how much autonomy they are allowed to exercise. Success depends on thoughtful decisions around goal formulation, architecture, and control frameworks that align with the enterprise's risk, compliance, and business priorities.

Most financial institutions initiate their agentic AI journey by deploying narrow agents in sandboxed environments or within low-risk, well-defined workflows. These early-stage pilots automate sub-tasks, accelerate documentation processes, or support frontline staff with recommendations, offering a safe proving ground for new capabilities. However, as organizations scale agents into more complex and mission-critical functions, new design questions will emerge: What level of autonomy should agents have? When should they act independently versus escalate decisions? And how can institutions ensure transparency and oversight of agent behavior?

Designing intelligent agents: from principles to production

Answering the questions above requires a shift from experimentation to intentional, architecture-driven design. Effective agentic systems will rely on the three foundational components definition that will shape their behavior, boundaries, and trustworthiness.

- **Clear agent goals** - ensuring every agent has a defined business objective to guide its actions
- **Constraints and permissions** - limiting authority to approved datasets, tools, and systems
- **Override mechanisms** - embedding human escalation and intervention points to maintain accountability

Exhibit 9 illustrates three core elements that form the foundation for scalable and trusted autonomy in BFS.

Institutions must define how agents are created, what roles they play, and how much autonomy they exercise.

Exhibit 9: Foundational design components of intelligent agent design

Source: Everest Group (2025)

A Clear agent goals

- Define specific objectives (e.g., reconciliation, compliance anomaly detection, onboarding)
- Provide operational logic to guide decisions
- Ensure alignment with business outcomes

B Constraints and permissions

- Set authority limits on data, tools, and systems
- Apply role-based access controls
- Prevent scope creep or unauthorized actions

C Override mechanisms

- Establish escalation paths for human supervisors
- Allow pausing, modification, or reversal of actions
- Act as safety valves for risk and compliance triggers

As enterprises expand deployments across departments and business lines, these design components will define how safely agents operate and how confidently they can be scaled. Institutions that address them early will be better positioned to move from pilots to enterprise-wide adoption without compromising governance, compliance, or trust.

Approaches to structuring agent-based workflows

As financial institutions mature in their agentic AI deployment, they are adopting distinct architectural patterns tailored to the function, complexity, and risk profile of the workflows involved. These design strategies will allow enterprises to scale agent capabilities while maintaining control, accountability, and alignment with regulatory and operational priorities.

The following three design patterns are emerging as particularly relevant for BFS.

Exhibit 10: Emerging design patterns for structuring agent-based workflows

Source: Everest Group (2025)

Pattern	How it works	Best-suited workflows
Human-in-the-loop	Agents handle most workflow steps autonomously but defer critical or final decisions to human reviewers, ensuring oversight where regulatory scrutiny is highest.	High-risk domains where human judgment is essential, such as credit adjudication, trade surveillance, and compliance reviews
Swarm agents	Multiple agents work in parallel on discrete tasks within a larger process, enabling speed and scale.	High-volume, repetitive workflows such as onboarding, document classification, and fraud detection
Supervisory agents	Meta-agents oversee and coordinate subordinate agents, managing dependencies and ensuring consistency across workflows.	Complex, cross-functional areas such as wealth management, corporate lending, and treasury operations

These patterns are not mutually exclusive. In fact, as institutions scale agentic AI across business lines, they are combining these approaches, layering swarm architectures in departments and supervisory agents across them, all with human-in-the-loop oversight at key decision points. This modular and composable strategy allows institutions to tailor agent behavior based on workflow sensitivity, data complexity, and governance requirements, maximizing both value and trust. As these patterns mature, the next focus for institutions is embedding safeguards into agent behavior, ensuring that trust, auditability, and control are hardwired into system design from the outset.

Embedding trust: control, audit, and governance in agent design

Trust in agentic systems is not built by design but earned through strong governance. For financial institutions, embedding risk mitigation directly into the agent architecture is non-negotiable. This includes intent logging and explainability to track what the agent did and why, approval thresholds and override mechanisms for high-risk actions such as payments or credit decisions, and simulation environments to stress-test behavior under real-world edge cases.

These safeguards will be foundational in regulated sectors, such as banking, where transparency, auditability, and human reversibility are essential for responsible deployment. As institutions embed these governance capabilities, the next challenge becomes orchestrating agents not just individually but as coordinated networks. This sets the stage for the emergence of multi-agent ecosystems that can reason, collaborate, and scale intelligently across interconnected financial workflows.

Scaling intelligently: designing multi-agent ecosystems

As enterprises mature in how they design agentic systems, the next evolution will be systems comprising intelligent agent networks that collaborate, coordinate, and adapt dynamically instead of isolated agents. This multi-agent paradigm will become the foundation for complex enterprise workflows, particularly those that require distributed reasoning, system-wide decisions, and continuous adaptation.

For financial services institutions, the shift toward agentic ecosystems will define how autonomy is scaled responsibly across teams, products, and regulatory boundaries. This evolution will mark a transition from agent design and governance at the individual level to dynamic systems of agents that interact and adapt across the enterprise. The focus now moves to how institutions can architect such collaborative networks to enable distributed intelligence at scale.

Understanding multi-agent collaboration in enterprise architecture

In traditional AI deployments, agents are limited to narrow, single-function tasks, such as extracting data from documents, retrieving client records, or generating recommendations. However, the next frontier lies in building multi-agent ecosystems where intelligent agents operate in collaboration, not isolation. In this model, agents are designed with role-based specialization, each configured for a specific function, such as fraud detection, tax advisory, or onboarding triage.

These agents can then delegate and coordinate tasks among themselves, mirroring human team dynamics to improve efficiency and adaptability. More advanced systems enable dynamic routing of responsibilities, where supervisory agents oversee workflows and reassign tasks in real time based on capacity, context, or exception handling.

This architecture will be particularly valuable in complex, high-stakes BFS workflows, such as corporate treasury, trade reconciliation, and portfolio management, where distributed intelligence and synchronized execution are essential for speed, accuracy, and compliance. As agent networks become central to execution, it will become essential for financial institutions to enable secure collaboration, dynamic reasoning, and reliable orchestration across multi-agent systems that operate within regulated environments. Unlike single-agent architectures, these multi-agent systems introduce greater complexity, requiring enhanced technical enablers for distributed decision-making, inter-agent coordination, and system-level monitoring. This will mark a strategic inflection point where isolated agent functions give way to coordinated, enterprise-wide intelligence.

Enabling collaboration through multi-agent infrastructure

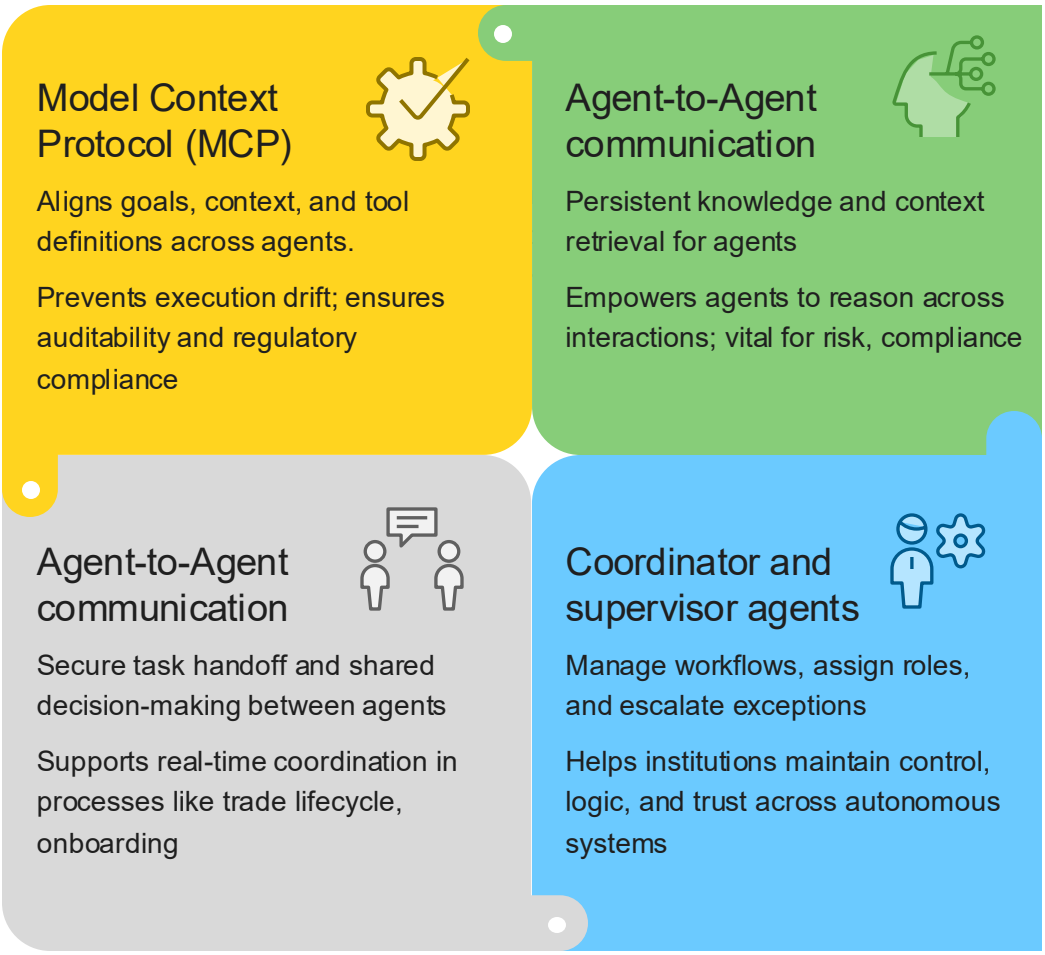
Multi-agent architecture will require more than orchestration. Unlike single-agent systems, which rely on localized reasoning and narrow execution paths, multi-agent environments demand a more robust technical foundation. These foundational elements

will ensure that intelligent agents can collaborate, reason, and scale securely, making them viable for complex, compliance-heavy financial workflows.

Exhibit 10 outlines the foundational components required to enable multi-agent agentic AI systems, highlighting the capabilities needed for secure, collaborative, and role-specialized execution in complex financial environments.

Exhibit 11: Key building blocks of multi-agent agentic AI systems in financial services

Source: Everest Group (2025)



As enterprises expand their agentic AI maturity, the move from isolated agents to multi-agent ecosystems will be a defining shift. These systems introduce new coordination, transparency, and autonomy layers that align with financial operations' complexity. However, their success will depend on how well institutions integrate planning, memory, security, and collaboration into their architectural foundations. Designing with these elements from the outset is not only a technical exercise but also a strategic choice that will determine whether agentic AI remains a point solution or evolves into a scalable operating model.

www.everestgrp.com | this document has been licensed to **Infosys**

Challenges and limitations in scaling agentic AI

While architectural foundations provide a pathway toward scalable agentic AI, adoption will not be without hurdles. Several challenges are likely to surface as enterprises move from pilots to production:

- **Integration complexity:** embedding agentic systems into legacy BFS infrastructure will require significant re-architecture and investment
- **Governance maturity:** oversight, auditability, and explainability standards remain in flux, creating uncertainty in highly regulated environments
- **Operational risks:** without strong monitoring, agents may escalate errors or trigger compliance concerns
- **Cost and talent barriers:** high development costs and limited availability of specialized talent may restrict adoption beyond tier-1 institutions
- **Organizational readiness:** shifting accountability from human-driven to semi-autonomous systems will demand new operating models and cultural alignment

These challenges highlight that adoption will not follow a linear trajectory. Success will come to institutions that treat governance, risk oversight, and organizational readiness as equal priorities to technology – enabling agentic AI to scale responsibly and sustainably.

As organizations progress from architectural readiness to enterprise-wide application, strategic and commercial considerations will become increasingly essential in guiding scalability, value realization, and operational alignment.

www.everestgrp.com | this document has been licensed to **Infosys**

Strategic and commercial implications of scaling agentic AI

As agentic AI moves from labs to enterprise deployment, financial institutions face tough, new strategic and operational questions. Success will not come from the most advanced models; it will come from scaling agents with clear goals in a governed, goal-aligned, and outcome-driven way.

Based on Everest Group's review of early adopters, emerging provider strategies, and evolving delivery models, this section outlines the key strategic and commercial implications of deploying agentic AI at enterprise scale. These implications span four strategic dimensions:

Exhibit 12: Dimensions shaping enterprise readiness for agentic AI deployment

Source: Everest Group (2025)

Dimension	What it means	Enterprise implication
Enterprise-wide readiness	Legacy systems, fragmented workflows, and regulatory ambiguity will slow adoption	Business processes will need to be reconfigured for modular execution and cross-functional ownership
Delivery model alignment	Centralized AgentOps vs. domain-packaged approaches create trade-offs in governance and flexibility	Delivery model choices will determine scalability, risk control, and compliance integration
Technology maturity	Hyperscaler platforms and data ecosystems define the pace and scope of viable deployments	Ecosystem partnerships and orchestration design will become critical for long-term success
Governance and regulation	Accountability, explainability, and auditability are non-negotiable in BFS	Alignment with regulatory expectations will be as important as technical capability

To succeed at scale, financial institutions must first understand the systemic blockers and enablers. Each of the strategic dimensions outlined above offered a lens into what is required to move agentic AI from localized pilots to operational systems. Yet even with this clarity, institutions can encounter entrenched challenges that delay momentum. These barriers form the next critical layer of consideration and directly influence whether institutions can realize the value potential of agentic AI in production environments.

www.everestgrp.com | this document has been licensed to **Infosys**

What's holding back enterprise adoption

Despite growing enthusiasm, enterprise-wide agentic AI adoption has been slower and more complex than anticipated, particularly in core financial workflows. The gap between pilot success and scaled deployment is not due to the lack of potential but due to systemic barriers that institutions must tackle directly. Some of these barriers are:

- **Legacy architecture creates friction:** Most legacy systems in BFS were never designed to support goal-driven, autonomous execution. Integrating intelligent agents into these environments requires more than surface-level APIs. It demands orchestration frameworks, secure data access, and modular, fault-tolerant deployment strategies that can span decades-old infrastructure
- **Value realization remains uncertain:** While agentic systems have demonstrated tangible benefits, such as reducing manual effort, improving decision accuracy, and accelerating turnaround time, these gains are often dispersed across teams or buried in complex workflows. Thus, business leaders struggle to quantify value in financial terms, delaying further investment and buy-in
- **Regulatory ambiguity persists:** Compliance leaders remain cautious about introducing agents that make autonomous or semi-autonomous decisions, particularly when explainability, traceability, and audit readiness are not guaranteed. Without standardized guidelines around AI governance further complicates confidence in agent-led processes
- **Lack of organizational readiness:** Agentic systems challenge traditional operating models, requiring new governance frameworks, cross-functional collaboration, and clearly defined ownership across business, technology, and compliance. Without this structural alignment, agent deployment risks becoming fragmented, underutilized, or blocked by internal inertia

To overcome these barriers, institutions will need to treat agentic AI not just as a technology initiative, but as a strategic transformation that reshapes enterprise architecture, redefines risk oversight, and reimagines how work gets done at scale.

Once institutions begin addressing these foundational roadblocks, the next layer of focus becomes how to design and operationalize the delivery model - determining where agents live, how they are governed, and who owns their lifecycle across business lines.

Evolving from sandbox pilots to structured enterprise-wide delivery

As agentic AI deployments start to move beyond experimentation, financial institutions would need to rethink how agents are delivered, governed, and managed across the enterprise. The shift from isolated lab pilots to production-ready systems would reshape internal operating models and delivery structures. Three dominant patterns are emerging.

Exhibit 13 illustrates three emerging delivery patterns for scaling agentic AI, highlighting how domain-specific agents, centralized AgentOps, and co-innovation with partners differ in their balance of agility, control, and scalability.

Exhibit 13: Which delivery pattern should be adopted for scaling agentic AI?

Source: Everest Group (2025)

	Domain-Specific Agents	AgentOps Structures	Co-Innovation with Partners
How it works	Focused units (e.g., KYC bots, treasury monitors) deployed independently; agile and easier to test.	Centralized teams manage deployment, monitoring, overrides, and compliance across functions.	Collaborations with hyperscalers, providers, or SIs to create modular agent stacks tailored to enterprise needs.
Strategic trade-offs	Risk of duplication and fragmented oversight without broader coordination	Standardized and auditable but may reduce flexibility for line-of-business needs.	Accelerates adoption but risks provider lock-in without strong internal ownership.

Each delivery approach carries strategic implications for ownership, risk management, compliance alignment, and scalability. As agent networks expand, the design and execution of delivery models will directly influence trust, speed, and scale. With these foundations in place, the focus shifts to the success enablers that determine whether institutions can convert pilots into production-ready, governed agentic AI systems.

Enabling success through enterprise readiness

As agentic AI becomes more embedded in enterprise operating models, success will depend on how well institutions align supporting enablers across data architecture, integration, governance, and business value. This alignment must be built on a foundation of operational readiness, including the reconfiguration of existing processes to support modular execution and coordinated decision-making. These enterprise-level capabilities must move in lockstep with delivery models to ensure agentic systems are secure, explainable, and accountable. The exhibit below outlines the foundational elements institutions must address to progress from pilots to production-grade deployments.

Exhibit 14 demonstrates the core enterprise dimensions that BFS enterprises must address to transition from agentic AI pilots to production-ready, governed, and outcome-driven deployments.

www.everestgrp.com | this document has been licensed to **Infosys**

Exhibit 14: The strategic framework for enterprise-wide agentic AI deployment

Source: Everest Group (2025)

Enterprise dimension	Strategic focus	Implications of not addressing the dimension	What leading enterprises are doing
Data and memory architecture	Move from prompt chains to persistent context and retrieval	Stateless agents, hallucinations, no decision traceability	Implementing vector stores, RAG setups, private APIs
Governance and oversight	Establish centralized AgentOps with monitoring and override control	Ad hoc agent behavior, lack of auditability or compliance trace	Standing up AgentOps teams with dashboards and control panels
Orchestration and integration	Enable agents to operate across tools, workflows, and human checkpoints	Agents stuck in silos or unable to escalate/coordinate	Using orchestration frameworks (e.g., LangChain, AutoGen)
Technology provider and partner strategy	Align build/assemble/buy across agent runtime, orchestration, and memory	Tech provider lock-in or duplicated investments	Partnering for orchestration/memory, building internal governance
Use case and value alignment	Prioritize measurable agent outcomes tied to business KPIs	Low ROI visibility, tech-first pilots with no path to scale	Starting with high-friction processes: KYC, onboarding, mid-office ops

While the strategic framework outlines the core enterprise dimensions that must be addressed, financial institutions can use the checklist mentioned in Exhibit 13 to assess their operational readiness for scaling Agentic AI. This assessment ensures that key controls, architecture, and ownership mechanisms are in place before moving to production.

Exhibit 15: Key readiness checkpoints for enterprise-wide agentic AI deployment

Source: Everest Group (2025)

- Have we implemented a persistent memory layer (e.g., vector database) beyond prompt chaining?
- Is there a centralized AgentOps team managing agent lifecycle, override rights, and compliance?
- Are orchestration frameworks in place to support multi-step workflows and agent fallback logic?
- Have we clearly mapped vendor roles across agent runtime, orchestration, and governance layers?
- Are agent use cases tied to measurable KPIs such as speed, cost reduction, or accuracy improvement?
- Do compliance teams have visibility into how agent decisions are logged, explained, and auditable?
- Have override paths and escalation controls been defined for high-risk or customer-facing agents?

Conclusion

Agentic AI is no longer a conceptual innovation. It is beginning to reshape how financial institutions approach execution, autonomy, and transformation. Unlike traditional automation or model-driven decision-making, agentic systems operate with intent, memory, and adaptability, positioning them as the next-generation enterprise architecture's foundational component.

However, realizing this potential requires more than deploying agents in isolated workflows. Financial institutions will need to define a clear path to value, supported by orchestration maturity, governance frameworks, and aligned stakeholder ownership. The focus must shift from pilot use cases to system-level design, where agents are embedded with guardrails, auditability, and enterprise accountability.

Institutions that lead this shift will not necessarily be the fastest adopters. They will be the ones that scale agentic systems with clarity, control, and business alignment. For enterprise leaders, the question is no longer whether agentic AI is viable. It is around – Where does it fit? How does it perform? And how do we integrate it into the operating model's core?

Looking ahead, the agentic paradigm is likely to evolve toward more adaptive, cross-enterprise ecosystems - where agents not only coordinate within institutions but also interact across markets, platforms, and regulatory networks. The convergence of agentic systems with tokenized assets, real-time data infrastructure, and AI-native platforms may give rise to entirely new business models. As adoption deepens, success will hinge on how well institutions design for both near-term control and long-term extensibility.

www.everestgrp.com | this document has been licensed to **Infosys**

Everest Group®
With you on the journey

Everest Group is a leading research firm helping business leaders make confident decisions. We guide clients through today's market challenges and strengthen their strategies by applying contextualized problem-solving to their unique situations. This drives maximized operational and financial performance and transformative experiences. Our deep expertise and tenacious research focused on technology, business processes, and engineering through the lenses of talent, sustainability, and sourcing delivers precise and action-oriented guidance. Find further details and in-depth content at www.everestgrp.com.

This study was funded, in part, by Infosys

For more information about Everest Group, please contact:

+1-214-451-3000

info@everestgrp.com

For more information about this topic please contact the author(s):

Pranati Dave, Vice President
pranati.dave@everestgrp.com

Ketan Kumar, Senior Analyst
ketan.kumar@everestgrp.com

Notice and Disclaimers

Important information. Please read this notice carefully and in its entirety. By accessing Everest Group materials, products or services, you agree to Everest Group's Terms of Use.

Everest Group's Terms of Use, available at www.everestgrp.com/terms-of-use, is hereby incorporated by reference as if fully reproduced herein. Parts of the Terms of Use are shown below for convenience only. Please refer to the link above for the full and official version of the Terms of Use.

Everest Group is not registered as an investment adviser or research analyst with the U.S. Securities and Exchange Commission, the Financial Industry Regulation Authority (FINRA), or any state or foreign (non-U.S.) securities regulatory authority. For the avoidance of doubt, Everest Group is not providing any advice concerning securities as defined by the law or any regulatory entity or an analysis of equity securities as defined by the law or any regulatory entity. All properties, assets, materials, products and/or services (including in relation to gen AI) of Everest Group are provided or made available for access on the basis such is for informational purposes only and provided "AS IS" without any warranty of any kind, whether express, implied, or otherwise, including warranties of completeness, accuracy, reliability, noninfringement, adequacy, merchantability or fitness for a particular purpose. All implied warranties are disclaimed to the extent permitted by law. You understand and expressly agree that you assume the entire risk as to your use and any reliance upon such.

Everest Group is not a legal, tax, financial, or investment adviser, and nothing provided by Everest Group is legal, tax, financial, or investment advice. Nothing Everest Group provides is an offer to sell or a solicitation of an offer to purchase any securities or instruments from any entity. Nothing from Everest Group may be used or relied upon in evaluating the merits of any investment. Do not base any investment decisions, in whole or part, on anything provided by Everest Group.

Everest Group materials, products and/or services represent research opinions or viewpoints, not representations or statements of fact. Accessing, using, or receiving a grant of access to Everest Group materials, products and/or services does not constitute any recommendation by Everest Group to (1) take any action or refrain from taking any action or (2) enter into a particular transaction. Nothing from Everest Group will be relied upon or interpreted as a promise or representation as to past, present, or future performance of a business or a market. The information contained in any Everest Group material, product and/or service is as of the date prepared and Everest Group has no duty or obligation to update or revise the information or documentation.

Everest Group collects data and information from sources it, in its sole discretion, considers reliable. Everest Group may have obtained data or information that appears in its materials, products and/or services from the parties mentioned therein, public sources, or third-party sources, including data and information related to financials, estimates, and/or forecasts. Everest Group is not a certified public accounting firm or an accredited auditor and has not audited financials. Everest Group assumes no responsibility for independently verifying such information.

Companies mentioned in Everest Group materials, products and/or services may be customers of Everest Group or have interacted with Everest Group in some other way, including, without limitation, participating in Everest Group research activities.



ABOUT INFOSYS

Infosys is a global leader in next-generation digital services and consulting. Over 320,000 of our people work to amplify human potential and create the next opportunity for people, businesses, and communities. We enable clients in 59 countries to navigate their digital transformation. With over four decades of experience in managing the systems and workings of global enterprises, we expertly steer clients, as they navigate their digital transformation powered by cloud and AI. We enable them with an AI-first core, empower the business with agile digital at scale and drive continuous improvement with always-on learning through the transfer of digital skills, expertise, and ideas from our innovation ecosystem. We are deeply committed to being a well-governed, environmentally sustainable organization where diverse talent thrives in an inclusive workplace.

Infosys®
Navigate your next

© 2025 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.