The Rise of Agentic AI

From Experimentation to Enterprise Readiness

If the first phase of agentic AI was about possibility, the next phase is about proof. As Everest Group's January 2026 report, Agentic AI in Enterprises: Readiness, Realities, and Roadblocks, makes clear, we are now firmly in the early-but-accelerating stage of the agentic maturity curve. Most enterprises have begun experimenting, and nearly two in five surveyed have been on the agentic journey for over a year. Yet the majority remain in pilot mode, with only about one in five initiatives reaching production. This signals a pivotal inflection point: enterprises believe agentic AI is transformative, but they are still building the foundational capabilities—data, governance, operating models, and talent—required to scale it responsibly. In other words, the market has moved beyond curiosity, but not yet to institutionalization.

Agentic AI's promise is real—but scaling it demands discipline. Enterprises are approaching agentic AI pragmatically, prioritizing operational efficiency, productivity, and cost optimization as early proof points, with more than 90% expecting measurable performance improvements. However, success will not come from technology alone. It depends on aligning executive sponsorship, funding models, governance frameworks, and workforce transformation. Organizations that treat agentic AI as a long-term transformation—rather than a series of disconnected pilots—are the ones most likely to convert early experimentation into durable competitive advantage.

This Everest Group whitepaper, supported by NTT DATA, provides a comprehensive assessment of where enterprises stand—and what it will take to move forward. Based on a global survey of 112 large enterprises, Everest Group maps adoption maturity across three segments—Initial Explorers, Structured Experimenters, and Strategic Scalers—and outlines the operating models, funding approaches, and sourcing strategies shaping progress. Looking ahead, the next 12–18 months will define whether enterprises remain in controlled experimentation or successfully transition to scaled, governed AI autonomy—where human-agent collaboration becomes the new operating model for the enterprise.

Wendy Collins
Chief AI Officer for
North America
NTT DATA, Inc.

Everest Group®

April 2026

Agentic AI in Enterprises: Readiness, Realities, and Roadblocks

NTT DATA This document has been licensed to NTT DATA

Contents

Copyright © 2026, Everest Global, Inc. All rights reserved.

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

Introduction

The rapid evolution of generative and agentic AI has ushered in a new phase of enterprise transformation – one defined by intelligent systems that can reason, act autonomously, and collaborate with humans to deliver tangible business value. As technology advances at an unprecedented pace, the question is no longer what AI can do, but how effectively enterprises can harness and scale it to achieve measurable outcomes.

Agentic AI represents the next frontier in this journey, advancing beyond generative capabilities to systems that can independently reason, decide, and act. By understanding context, making informed decisions, and executing complex workflows, these systems are poised to redefine how work gets done, driving productivity, innovation, and agility across industries.

Yet for many enterprises, the agentic AI's promise remains elusive. Adoption remains concentrated in proofs of concept and small-scale pilots, and early obstacles – such as unclear RoI and high costs – often limit momentum. As organizations advance in their agentic AI journey, the challenge shifts from operational constraints to the broader realm of AI accountability, including concerns about ethical risks, bias, and governance. Against this backdrop, understanding how enterprises invest in and operationalize agentic AI becomes vital to distinguishing early wins from strategic inertia.

To explore this dynamic landscape, Everest Group conducted a global survey of 112 enterprises across industries and revenue segments, examining their readiness, adoption maturity, and priorities for scaling. The study provides a comprehensive view of how organizations are approaching agentic AI – from investment and sourcing strategies to talent development and change management.

In this viewpoint, we offer a grounded perspective on how enterprises are translating AI's potential into measurable business impact. The analysis aims to help enterprise leaders benchmark their maturity, identify practical pathways to scale, and accelerate responsible adoption. For technology and service providers, it highlights evolving enterprise expectations and outlines opportunities to build solutions that are not only innovative but also enterprise-ready in the era of agentic AI.

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

What is agentic AI and how is it different?

Agentic AI is an advanced form of AI that enables autonomous agents to make decisions, adapt to context, and act independently to achieve defined outcomes. These agents can interpret multimodal inputs, set goals, devise and adjust plans, and execute multi-step workflows with minimal human oversight.

What sets agentic AI apart is its ability to go beyond fixed-rule automation and prompt-driven generative AI. Instead of simply following predefined steps or generating content on request, agentic AI interprets intent, reasons across tasks, adapts in real time, and orchestrates processes end-to-end.

Exhibit 1 illustrates how agentic AI differs from traditional Intelligent Automation (IA) and generative AI across key dimensions.

Exhibit 1: Defining characteristics of agentic AI relative to existing automation technologies

Source: Everest Group (2026)

	Traditional IA	Generative AI	Agentic AI
Scope	Handles predefined, repeatable tasks	Creates new content such as text, images, or code based on inputs/prompts	Manages complex processes, requiring dynamic adjustments
Task vs. goal orientation	Task-oriented, works best for predictable outputs	Not inherently goal-oriented, mainly focuses on producing creative or informative outputs	Flexible, goal-oriented; works on broader objectives and decomposes them into actionable steps
Decision-making	Follows specific, rule-based instructions	No decision-making ability; generates outputs based on prompts and learned patterns	High degree of autonomy; makes autonomous, data-driven decisions
Real-time adaptability	Static in its processes; requires rule updates by humans	Limited adaptability; can generate variations but lacks self-directed learning	Adapts strategies dynamically based on real-time data and conditions
Human supervision	Requires significant human intervention	Requires limited human intervention in the form of prompts to generate response	Requires minimal human oversight
Analogy	Like using a microwave preset – it runs the same programmed sequence every time to help in meal preparation, with no ability to adjust based on context	Like a recipe writer – it can generate creative variations of recipes when prompted, but it does not cook the meal or manage the process	Like a smart personal chef – you describe the meal you want, and it plans the recipe, adjusts ingredients, adapts in real time, and completes the task end-to-end

Current state of agentic AI adoption

Currently, agentic AI adoption can be best described as nascent yet steadily gaining momentum. Most enterprises are taking a deliberate, learning-oriented approach, experimenting through controlled pilots, developing use cases, and closely tracking early outcomes before committing to scale. While enthusiasm for the technology is evident, adoption remains concentrated in the initial phases, as enterprises explore feasibility, business value, and integration challenges. This cautious optimism reflects both the transformative potential and the operational complexity of implementing agentic AI at scale.

Nearly two out of three enterprises view agentic AI not as a mere upgrade, but as a reinvention of enterprise workflows and workforce dynamics, a transformative gateway to sustainable competitive advantage.

Adoption timelines and scale

Adoption timelines indicate that agentic AI is steadily gaining ground across enterprises. Nearly two in five surveyed organizations (42%) began their journey within the past 6-12 months, while another 37% have been experimenting for over a year. Large organizations – with revenues exceeding US$10 billion – are further ahead, with nearly half having adopted agentic AI for more than 12 months. Their progress is supported by stronger data and infrastructure foundations and more mature governance structures. Exhibit 2 illustrates the adoption timelines across enterprises.

Exhibit 2: Split of enterprises based on their agentic AI adoption timeline

Source: Everest Group (2026)

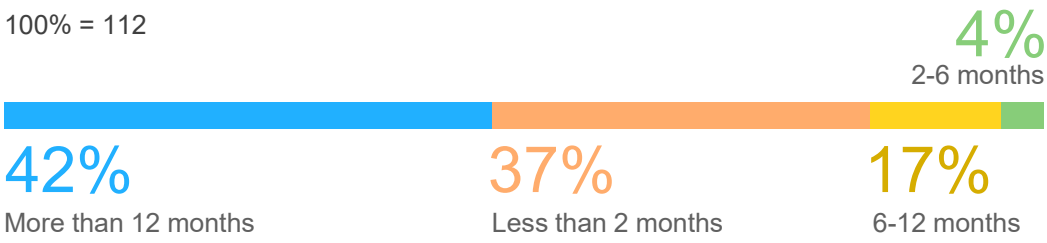

Beyond timelines, the nature and scale of ongoing projects provide a clearer picture of where enterprises stand in their agentic AI journey. To date, enterprise performance remains at an early stage, with only one in five pilots successfully moving into production. This underscores the experimental nature of current initiatives and the operational challenges of scaling autonomous systems. However, the outlook is markedly positive. As enterprises refine their strategies and align investments with measurable outcomes, positive momentum is expected over the next 12-18 months.

As enterprises move from pilots to early production, their focus is shifting from testing feasibility to identifying high-impact applications across business functions. A growing number of use cases are beginning to demonstrate tangible business value, paving the way for broader enterprise adoption.

Exhibit 3 highlights such use cases in various functional domains.

Exhibit 3: Key agentic AI use cases across business functions

Source: Everest Group (2026)

Finance and accounting	Invoice processing	Tax workflow automation	Collections optimization
HR	Employee onboarding	Training management	Resume shortlisting
Customer service	Customer engagement optimization	Self-service ticket resolution	End-to-end query handling
IT infra and operations	Debugging and code optimization	Iterative testing and validation	Cloud operations monitoring and patching
Supply chain management	Inventory management	Predictive maintenance	Logistics optimization
Sales and marketing	Campaign management	Smart pricing optimization	Intelligent email follow-ups
Cybersecurity	SOC automation	Threat hunting and hypothesis testing	Threat detection and response
BFSI[1] (industry-specific)	KYC processing	Policy underwriting	Portfolio management
HLS[2] (industry-specific)	Risk assessment of patient before screening/test	Appointment scheduling and follow-ups	Regulatory reporting for new drug application
Manufacturing (industry-specific)	Root-cause analysis on the line	Production scheduling and rescheduling	Knowledge-based engineering

1 Banking, financial services, and insurance
2 Healthcare and life sciences

Adoption maturity spectrum

The current mix of pilots and production deployments reveals distinct maturity stages among enterprises, influenced by variations in digital readiness, governance rigor, and strategic intent. Based on current adoption patterns, enterprises can broadly be classified into three segments, each representing a different level of progression in scaling agentic AI and realizing business impact. Exhibit 4 outlines these segments and their defining characteristics.

Exhibit 4: Enterprise segments based on agentic AI adoption maturity

Source: Everest Group (2026)

		Focus areas
Initial explorers Notable share (~25-30%)	Limited pilots or proofs of concept focused on feasibility exploration Low automation maturity with minimal governance frameworks in place	Defining business cases and measurable success metrics Building foundational data, infrastructure, and talent
Structured experimenters Majority of enterprises (~55-60%)	Multiple pilots underway with select use cases moving toward production Moderate automation maturity with evolving governance mechanisms	Scaling validated use cases across additional functions Strengthening data, infrastructure, and cross-functional collaboration
Strategic scalers Small but growing share (~10-15%)	Multiple agentic AI use cases in production delivering measurable business impact Advanced automation maturity supported by robust governance frameworks	Expanding adoption across functions and regions Institutionalizing unified orchestration and governance models

These maturity variations are not coincidental – they reflect deliberate strategic choices as enterprises align agentic AI investments with broader goals for efficiency, innovation, and growth.

Strategic drivers and enterprise value realization

As enterprises advance in their agentic AI journey, their focus is becoming increasingly outcome-oriented. Most organizations are prioritizing initiatives that deliver measurable improvements in operational performance, productivity, and cost efficiency, while gradually exploring opportunities in innovation and adaptability.

Exhibit 5 illustrates the top drivers of agentic AI adoption and the corresponding data around enterprises anticipating value realization from these drivers/business objectives.

Exhibit 5: Agentic AI adoption's key drivers and the corresponding positive impact on business objectives

Source: Everest Group (2026)

100% = 112

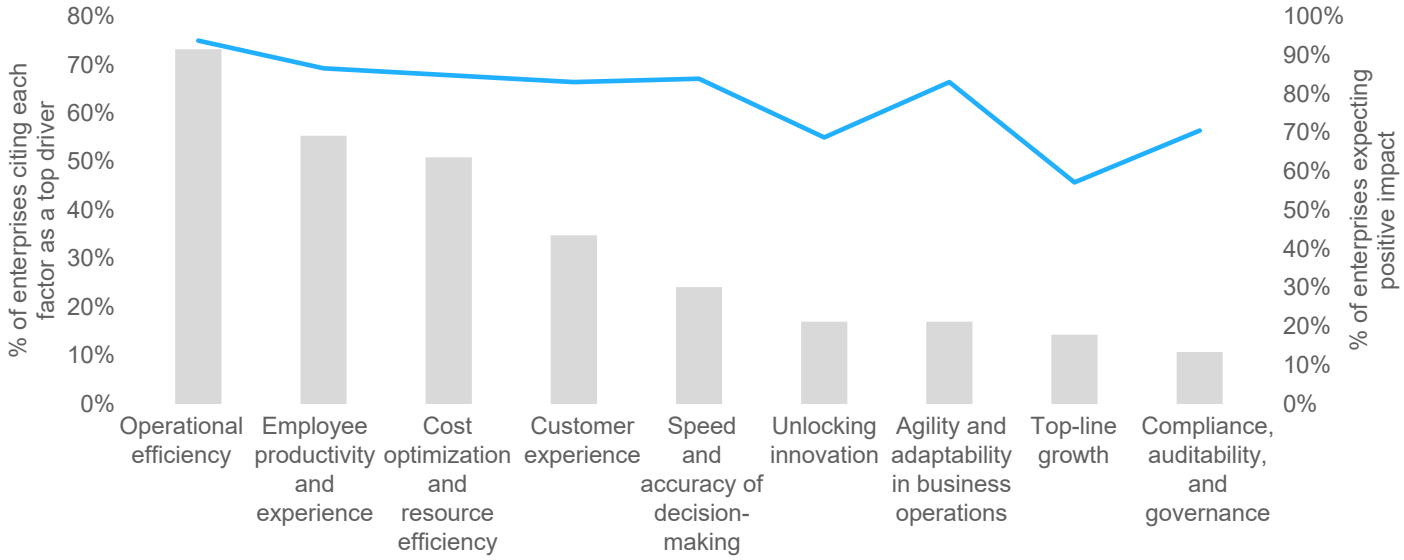

Survey findings reveal that enterprises are approaching agentic AI with a pragmatic focus on driving measurable, near-term outcomes. Efficiency, productivity, and cost optimization remain the top objectives, highlighting a clear intent to validate business value through tangible outcomes. More than 90% of respondents anticipate improvements in performance and workforce productivity, indicating that enterprises are prioritizing high-certainty, operational use cases that deliver visible impact and build organizational confidence. This early emphasis on operational excellence provides a critical foundation for trust and scalability – enabling enterprises to progress from "safe bets" to more strategic applications in areas such as innovation, customer experience, and business agility.

Operating models for agentic AI initiatives

As organizations transition from experimentation to structured deployment, defining the right operating model becomes a critical enabler of scale and control. Enterprises are testing different combinations of funding, governance, and execution to balance innovation with accountability. The sections below examine how enterprises are structuring ownership and decision rights to ensure that agentic AI investments remain strategically aligned and operationally sustainable.

Budget ownership

Enterprises are adopting varied approaches to agentic AI, ranging from centralized oversight to decentralized experimentation. Survey results reveal that for a vast majority (84%), agentic AI is now on the CXO agenda, reflecting a strategic intent to institutionalize it as part of broader enterprise transformation. CIOs and CTOs are the primary budget holders for agentic AI initiatives, controlling over half of the total investment authority across enterprises. This emphasizes the technology-intensive nature of current implementations, where infrastructure, integration, and governance readiness dictate progress. As illustrated in Exhibit 6, business unit and functional heads are key stakeholders following the CXOs.

Exhibit 6: Budget holders for agentic AI initiatives

Source: Everest Group (2026)

XX% Percentage of respondents

100% = 112

54%
CIO/CTO

46%
CFO

39%
CDO/CIO

36%
Business unit / GBS / Functional heads

35%
Head of AI / Data / Analytics

28%
CTO

22%
Automation CoE Heads

The dominance of CXO-level funding reflects a clear strategic intent to embed agentic AI within enterprise transformation efforts.

Execution ownership

Execution responsibility for agentic AI initiatives closely mirrors funding dynamics. Nearly 44% of enterprises have centralized execution under CIO or CTO leadership, underscoring a preference for strong internal control and consistent governance. Only a small share of organizations maintain decentralized or outsourced execution models, reflecting the view that agentic AI is a core capability requiring dedicated oversight and enterprise-wide integration. Exhibit 7 illustrates the different ownership models enterprises are adopting to drive agentic AI execution.

Exhibit 7: Ownership for implementing agentic AI initiatives

Source: Everest Group (2026)

XX% Percentage of respondents
100% = 112

44% Centralized ownership under the CIO/CTO organization

35% Jointly owned by business and technology teams (shared model)

8% Owned by the AI / data science / analytics function

6% Decentralized, with individual business units driving their own implementations

4% Led by a dedicated automation or digital transformation office / CoE

1% Outsourced or co-owned with a service provider or partner

The centralization trend indicates that enterprises view agentic AI as strategically important and have assigned formal responsibility to technology leaders to ensure unified governance. Yet, despite this centralization, many organizations still experience fragmented execution as business units deploy agentic capabilities independently or adopt Commercial Off-the-Shelf (COTS) tools without IT involvement. As deployments scale, enterprises will need to adopt more structured hybrid models that maintain centralized standards while establishing clearer engagement mechanisms for business-led innovation. This approach will accelerate adoption while strengthening oversight.

The current distribution of funding and execution ownership shows that enterprises are still determining how agentic AI should fit within their broader investment and governance frameworks. Technology leaders may control the budget and delivery mandate, but growing involvement from finance and business stakeholders reflects a rising need for shared accountability, particularly as decentralized experimentation

creates governance gaps. As adoption expands, enterprises will need to formalize investment models that define ownership, prioritization, and guardrails to ensure funding aligns with enterprise-wide objectives. The next section examines how organizations are shaping these investment strategies, including funding mechanisms and planned spending, to support scaled, well-governed agentic AI deployment across the enterprise.

Funding model

Enterprises are currently experimenting with different funding approaches as they determine where and how to invest effectively in agentic AI. The absence of standardized models reflects varying levels of organizational readiness, governance maturity, and confidence in value realization.

Most enterprises (71%) follow a project-based funding approach, allocating budgets to individual PoCs or use cases without a recurring enterprise-wide process. Only 22% have dedicated budgets for agentic AI, while 7% rely on function-led funding through independent business units. Exhibit 8 depicts the split between different funding models being used for agentic AI initiatives.

Exhibit 8: Current funding models for agentic AI projects

Source: Everest Group (2026)

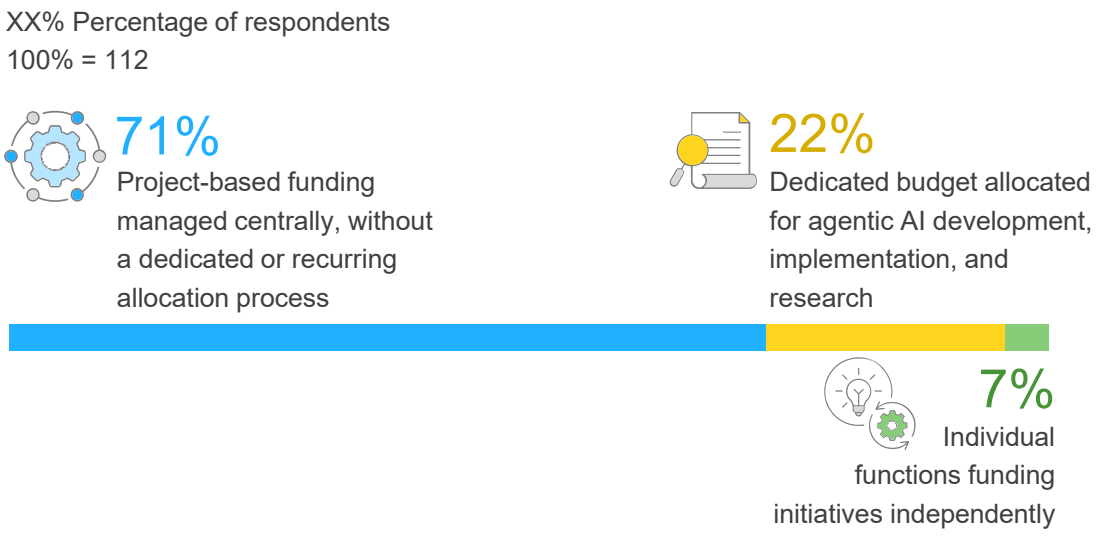

The findings reveal a clear link between funding maturity and RoI expectations:

- Enterprises with dedicated allocations tend to report shorter RoI horizons (6-12 months) and are more likely to have formal Centers of Excellence (CoEs) driving governance and capability building
- Project-based funding models, on the other hand, correspond with longer RoI expectations (12-24 months) and more distributed governance through ad hoc or function-led working groups

The dominance of project-based funding indicates that agentic AI remains in a learning and validation phase, with enterprises prioritizing controlled experimentation before institutionalizing budgets. As confidence builds, funding is expected to shift from fragmented, project-level spending toward structured, enterprise-level allocations, enabling sustained scaling and integration.

61% of enterprises believe that agentic AI is positioned to deliver faster RoI than previous intelligent automation technologies.

Scale of investment

Investment levels mirror the cautious yet strengthening enterprise stance toward agentic AI. Most organizations are adopting a measured approach, directing limited capital to test feasibility and validate use cases while preparing to expand funding for proven deployments.

However, the outlook points to a decisive inflection. As highlighted in Exhibit 9, nearly 65% of enterprises plan to double or more their current investments over the next 12 months, with many moving from the sub-US$1 million range into the US$2-5 million bracket. Furthermore, 95% intend to sustain or increase investments within the same period, signaling growing conviction in agentic AI's business potential.

Exhibit 9: Current and future investment levels in agentic AI initiatives

Source: Everest Group (2026)

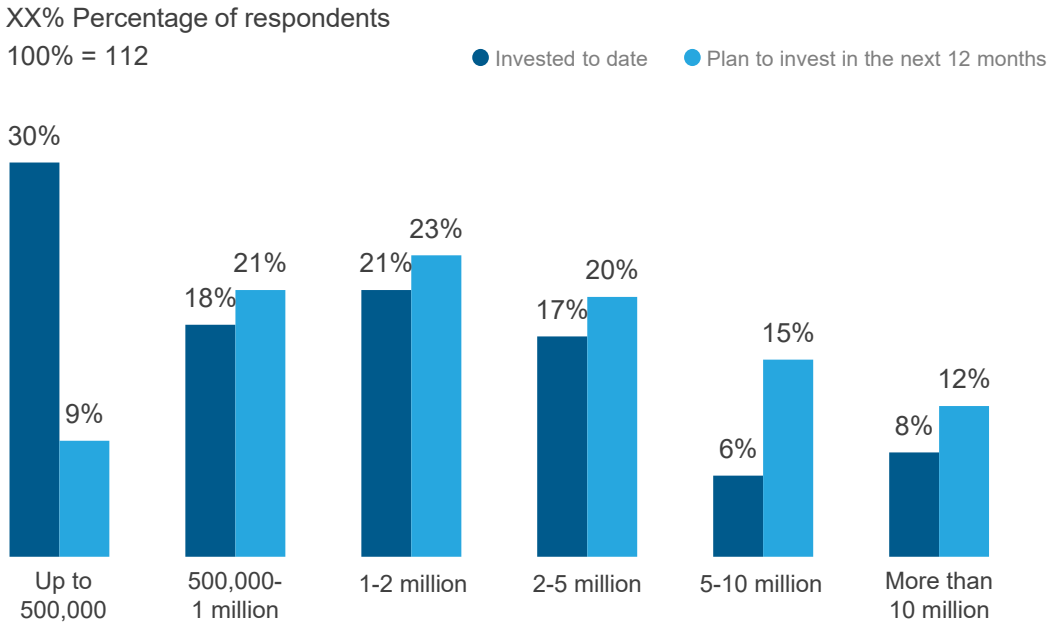

Sourcing strategies

With growing investment momentum, enterprises are exploring various ways to source and operationalize emerging capabilities. Most organizations are still determining the right balance between in-house development and external partnerships, influenced by factors such as control, scalability, data security, and internal AI maturity.

Two dominant sourcing models are emerging:

- Technology provider-led: Leveraging products from technology providers and implementing them internally or through provider support
- Service provider-led: Engaging end-to-end solutions from service providers offering proprietary platforms and managed services

In addition, many enterprises are experimenting with hybrid sourcing models to accelerate learning while managing cost and performance risks. Exhibit 10 summarizes the range of sourcing approaches currently in play.

Exhibit 10: Sourcing approaches for agentic AI solutions

Source: Everest Group (2026)

XX% Percentage of respondents
100% = 112

50%

Product from a **technology provider, implemented** in-house

33%

Service provider delivers **services on third-party technology procured directly by enterprise**

44%

Service provider delivers an **end-to-end solution**

37%

Open-source/Self-hosted **solution implemented and maintained** in-house

31%

Platform fully built and **implemented** in-house

29%

Service provider delivers **services using third-party technology it has procured**

Role of ecosystem partners: technology and service providers

As illustrated above, both technology providers and service providers are playing critical roles in enabling enterprises' early agentic AI initiatives. Technology providers supply the platforms, infrastructure, and tools needed to design, deploy, and scale AI agents, while service providers bridge the gap between capability and execution through advisory, implementation, and operational support. Together, they form the backbone of the evolving agentic AI ecosystem, helping organizations move from conceptual exploration to tangible business impact.

Technology providers are shaping the ecosystem by delivering foundational platforms, infrastructure, and tools that enable enterprises to design, deploy, and scale AI agents. Leading hyperscalers such as Microsoft, AWS, and Google Cloud are establishing this foundation through integrated platforms, pretrained models, cloud infrastructure, and development environments that support secure scaling. At the same time, enterprise software providers such as Salesforce, ServiceNow, and SAP are embedding agentic capabilities directly into business applications, allowing contextual and workflow-level adoption.

Among technology providers, hyperscalers dominate enterprise adoption, with Microsoft leading through integrated offerings such as Copilot Studio, M365 Copilot, and Azure OpenAI Services. Its strong enterprise ecosystem and focus on secure contextualization make it the preferred partner for many early implementations. AWS and Google Cloud continue to serve as key anchors, while enterprise software platforms such as ServiceNow and Salesforce are embedding agentic capabilities into their core products, enabling faster and domain-specific adoption.

With respect to capabilities offered by technology providers, seamless integration with enterprise systems emerged as the most critical requirement, followed by function-specific pre-built agents, multi-agent orchestration, and industry-specific solutions. Low-code build capabilities and agent monitoring and observability also rank highly, reflecting enterprises' need for operational control and rapid customization.

Complementing the role of technology providers, service providers are increasingly pivotal in helping enterprises strategize, identify use cases, and implement custom solutions. They often act as integration and orchestration partners, combining technologies from multiple providers, adapting them to business-specific contexts, and ensuring responsible deployment at scale. As enterprise adoption deepens, these partnerships are expected to evolve toward strategic co-creation models, where service providers collaborate closely on innovation, governance, and continuous optimization of agentic AI capabilities.

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

Enterprise expectations from ecosystem partners

As sourcing approaches mature, enterprises' expectations from ecosystem partners are becoming more defined. Enterprises are looking beyond generic AI capabilities; they expect partners to deliver secure, interoperable, and domain-adapted solutions that integrate seamlessly within existing technology and business frameworks. Exhibit 11 summarizes these expectations.

Exhibit 11: Enterprise expectations from ecosystem partners

Source: Everest Group (2026)

XX% Percentage of respondents
100% = 112

Featured in the Top 3 list		Ranked 1	Ranked 2	Ranked 3
46%	Proven agentic AI use cases and implementation playbooks	15%	18%	13%
44%	Domain expertise	14%	18%	12%
42%	Availability of skilled and accessible talent	17%	15%	10%
37%	Quick and error-free implementation	23%	5%	9%
35%	Ability to act as a business partner with a shared vision	7%	13%	15%
27%	Effective governance frameworks	13%	8%	6%

As enterprises strengthen their sourcing partnerships, they are recognizing that the success of these initiatives depends not only on technology and delivery models but also on the availability of skilled talent capable of designing, managing, and governing autonomous systems. The next section explores how enterprises are preparing their workforces, addressing skill gaps, and enabling effective human-AI collaboration.

Talent and workforce readiness

Enterprises increasingly recognize that talent is the defining enabler of agentic AI success. While most remain in the early deployment stages, many are already investing in capability development and role realignment to strengthen readiness. Talent gaps are emerging across both technical and strategic dimensions, spanning orchestration engineers, AI designers, data specialists, governance experts, and integration architects. In many cases, the challenge extends beyond acquiring niche skills to building multidisciplinary teams that connect technology development with business strategy and responsible governance.

Evolving roles to support agentic AI

To address these needs, enterprises are introducing new roles and expanding responsibilities within existing teams. Nearly half (48%) plan to add new AI-focused roles such as AI Product Managers and Knowledge Engineers, while 40% are creating agentic AI-specific positions such as AI Agent Engineer or Decision Engineer. A few organizations are also considering establishing a Head of Agentic AI role to oversee enterprise-wide strategy and governance. In parallel, organizations are emphasizing governance and risk oversight by adding roles such as AI Ethicists and Model Supervisors to ensure responsible adoption. These shifts reflect a move toward structured, accountable talent frameworks that balance innovation with control.

Exhibit 12 highlights the different approaches enterprises are considering to redefine roles and responsibilities to effectively support agentic AI adoption.

Exhibit 12: Organizational approaches to roles and responsibilities for agentic AI

Source: Everest Group (2026)

XX% Percentage of respondents

100% = 112

48%

Plan to introduce new AI roles

40%

Intend to establish agentic AI-specific roles

35%

Expect business stakeholders to take on broader responsibilities

29%

Will be augmenting existing AI roles

28%

Do not plan new roles but are likely to expand technical teams' responsibilities

Building and acquiring talent

Developing the right talent has become the most critical lever for progress. Most enterprises are adopting a hybrid approach that combines internal reskilling with selective external hiring and partnerships. As highlighted in Exhibit 13:

- Reskilling existing employees (78%) remains the dominant strategy, underscoring enterprises' focus on retaining institutional knowledge while enhancing AI fluency
- Hiring from other companies (57%) and partnering with third-party providers to build CoEs (54%) complement internal efforts, providing access to specialized expertise and enabling rapid capability expansion
- Outsourcing partnerships (53%) are also being leveraged to address near-term skill shortages in high-demand areas

This blended model enables enterprises to balance speed with sustainability, strengthening core capabilities while leveraging ecosystem expertise as adoption scales.

Exhibit 13: Reskilling and sourcing approaches shaping enterprise talent strategies

Source: Everest Group (2026)

XX% Percentage of respondents
100% = 112

Despite growing investment in talent and capability development, workforce readiness remains a key constraint to scaling agentic AI. As enterprises advance, bridging skill gaps, embedding governance, and aligning roles with evolving business needs will be essential to sustaining momentum. The following sections examine the barriers and enablers shaping enterprise progress – and what differentiates early movers from those still navigating adoption challenges.

Four in five enterprises believe the future of work will be defined by human-agent collaboration, with AI agents serving as digital teammates that enhance human judgment and accelerate execution.

Barriers and catalysts shaping enterprise adoption

Enterprises are actively engaging with agentic AI and see strong potential in the technology, yet the path to scaled adoption remains constrained by several challenges. Survey findings indicate that data security, privacy, and compliance risks (69%) remain the most pressing concerns, followed by lack of clarity on success metrics, RoI, and business value (64%), and integration challenges with existing systems (59%). These results show that enterprises' current focus is primarily on value realization and trust-building, while challenges related to scalability, talent, and governance are emerging as the next critical frontiers. Exhibit 14 summarizes the leading challenges constraining agentic AI adoption.

Exhibit 14: Key challenges that enterprises face in their agentic AI journey

Source: Everest Group (2026)

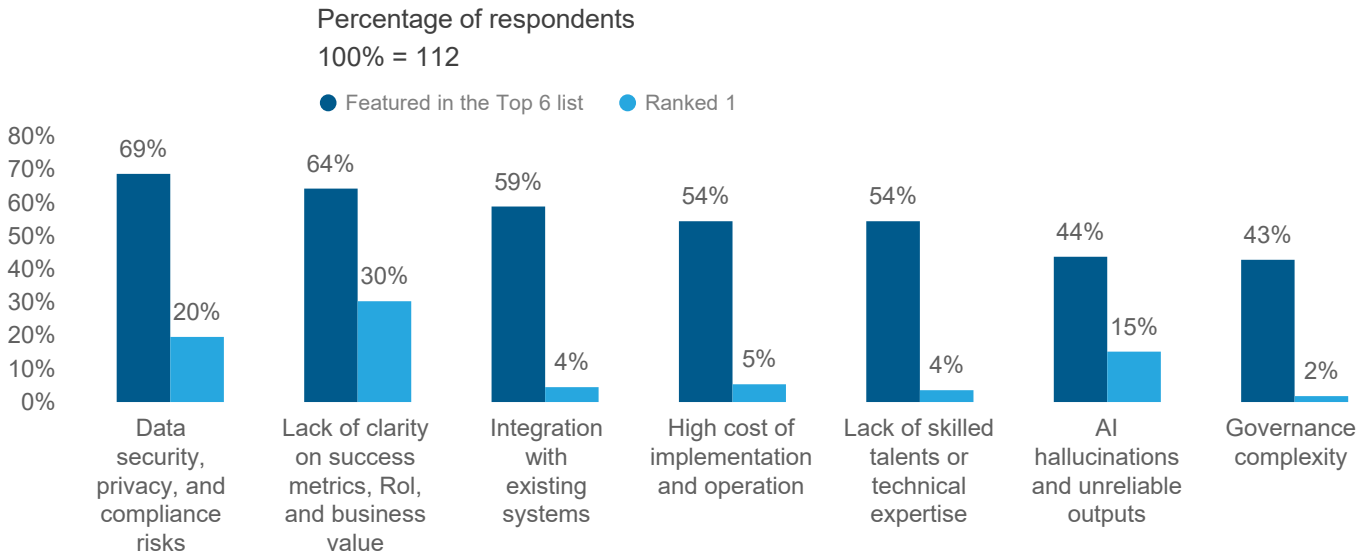

Note: Other challenges, arranged in descending order of how frequently they appeared in the Top 6 list, include: non-availability of quality data for training (39%), ethical risks, bias, and workforce impact concerns (35%), lack of explainability of agentic AI solutions (34%), high compute requirements (34%), need for contextualization and fine-tuning of foundation models (34%), latency in outputs (29%), and ESG/sustainability concerns (11%)

Survey findings indicate that data security, privacy, and compliance risks remain the most pressing concerns, followed by lack of clarity on success metrics, RoI, and business value, and integration challenges with existing systems.

These challenges highlight enterprises' dual struggle to build trust and demonstrate measurable value, while simultaneously establishing the technical and organizational foundations needed for scale. Early-stage barriers, such as unclear RoI and high costs, tend to diminish as enterprises mature and refine their adoption models. In contrast, ethical risks, bias, and governance concerns grow more prominent over time, signaling a shift from operational constraints to broader questions of responsibility, control, and explainability.

Best practices for adoption and scaling

Everest Group take

Agentic AI's impact will compound over time, rewarding enterprises that focus on sustained progress rather than short-term results. Organizations that adopt a long-term vision, invest consistently, and scale in a structured manner are best positioned to unlock transformative outcomes.

Enterprise survey data reinforces this perspective: organizations anticipating longer RoI timelines tend to perceive lower short-term impact from agentic AI. This pragmatic outlook reflects a deeper understanding that transformational value requires time for use cases to mature, integrate, and scale across the enterprise.

Despite ongoing barriers, enterprises are beginning to navigate complexity and unlock measurable impact from agentic AI. As they advance from curiosity to capability, success increasingly depends on balancing innovation with discipline. Achieving sustained value requires a structured roadmap that begins with readiness, advances through experimentation, and culminates in scalable transformation. Below, we outline two enterprise scenarios and the corresponding strategic actions organizations should prioritize to advance their agentic AI journey:

A. Organizations that are yet to start their agentic AI journey	**B.** Organizations that have begun their agentic AI journey and are currently in one of the three stages outlined in Exhibit 4

A. Organizations that are yet to start their agentic AI journey

For organizations preparing to embark on their agentic AI journey, the priority lies in establishing a solid foundation before scaling. Early decisions related to data, talent, and technology readiness often determine whether initiatives can scale effectively later. Enterprises should therefore evaluate their readiness across key dimensions – data and infrastructure maturity, governance frameworks, workforce capability, and integration with existing enterprise systems. A structured readiness checklist, illustrated in Exhibit 15, helps identify capability gaps, define ownership, and prioritize investments that enable experimentation with control and confidence.

Exhibit 15: The pre-journey agentic AI readiness checklist for enterprises starting their agentic journey

Source: Everest Group (2026)

1. **Evaluate necessity**: assess if agentic AI is truly needed or if existing automation or analytics tools can meet the same goals
2. **Align with business strategy:** ensure the agentic AI vision directly supports core enterprise objectives and transformation priorities
3. **Secure executive sponsorship:** appoint a senior leader to champion the initiative and allocate time and budget for exploration
4. **Establish a responsible AI stance:** define early principles for ethics, transparency, and accountability before starting pilots
5. **Assess current AI maturity:** map existing automation, analytics, and generative AI capabilities to identify realistic starting points
6. **Audit data quality and accessibility:** evaluate available data, its readiness for AI use, and existing silos or gaps
7. **Evaluate infrastructure:** review cloud, integration, and security readiness to ensure scalability for agentic models
8. **Identify skill gaps:** gauge internal AI literacy and plan for targeted upskilling or external partnerships
9. **Form a cross-functional task force:** involve business, data, IT, risk, and compliance teams to ensure aligned decision-making
10. **Create governance scaffolding:** define ownership, oversight, and controls for AI exploration, even before pilot launch
11. **Map potential ecosystem partners:** identify technology providers, consulting allies, and start-up collaborators
12. **Benchmark industry peers:** study early adopters to learn from their agentic AI implementation patterns
13. **Establish data-sharing and API policies:** set up integration and exchange frameworks to enable secure, collaborative innovation
14. **Start with small co-innovation projects:** launch joint discovery sessions or labs to ideate use cases and test feasibility

B. Organizations that have begun their agentic AI journey

For enterprises already underway in their agentic AI journey, the focus shifts from readiness to scaling and institutionalization. These enterprises are building on early learnings to formalize operating models, strengthen governance, and embed measurable performance frameworks. The following segments outline best practices that help organizations move decisively from exploration to transformation.

Segment 1: Initial explorers

1. Assess readiness: identify gaps in data, talent, and infrastructure vital to enabling agentic AI	2. Connect ecosystems: align internal teams, technology partners, and data pipelines to establish a cohesive foundation for scalable experimentation	3. Test and learn: run small, low-risk pilots that build confidence, demonstrate value, and create internal advocacy

Segment 2: Structured experimenters

1. Align leadership and strategy: foster enterprise-wide alignment through cross-functional collaboration among business, technology, and governance teams	2. Institutionalize governance: implement modular, repeatable frameworks that balance agility with control, enabling confident scaling, reuse of best practices, and consistent compliance standards	3. Measure with precision: redefine success metrics to include agent adaptability, reasoning accuracy, and learning progression, alongside traditional business impact

Segment 3: Strategic scalers

1. Reimagine enterprise workflows: move beyond wrapping agents around existing processes to redesign workflows natively for agentic intelligence, enabling dynamic, outcome-driven operations	2. Institutionalize responsible AI: build adaptive governance that evolves alongside agents, ensuring ethical autonomy, transparency, and accountability at scale	3. Synergize across ecosystems: develop cross-functional and cross-enterprise agent networks that exchange intelligence across partners, functions, and industries

4. Educate for continuous capability: focus on workforce enablement by developing AI-fluent teams that evolve with technology and collaborate seamlessly with intelligent systems for strategic decision-making

The experiences of leading adopters demonstrate that successful agentic AI transformation rests on a foundation of readiness, governance, and disciplined execution. As enterprises refine their operating models and scale responsibly, they are becoming more deliberate in what they expect from technology partners, seeking solutions that are not only powerful but also adaptable, secure, and aligned with enterprise ecosystems.

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

The next frontier in agentic AI

Enterprise expectations from technology

Agentic AI remains in a formative phase for most enterprises, and the next 12-18 months will be a defining period. As organizations continue experimenting, they are increasingly seeking technology that can help them move from pilots to production, and from assisted execution to intelligent autonomy. Exhibit 16 illustrates the top capabilities enterprises expect from agentic AI over the next 12-18 months.

Exhibit 16: Enterprise expectations from agentic AI (next 12-18 months)

Source: Everest Group (2026)

Percentage of respondents

100% = 112

38%

Expect more autonomy in decision-making and task execution

38%

Seek seamless integration across enterprise application

30%

Look for improved collaboration across multiple agents

23%

Anticipate enhanced cognitive capabilities and contextual intelligence

20%

Highlight the need for embedded governance and policy enforcement frameworks

These findings indicate that the enterprise focus for the next 12-18 months will center on scaling autonomy responsibly, ensuring that integration, oversight, and trust evolve alongside intelligence and automation. This reflects a shift from curiosity to clarity, as enterprises place integration, autonomy, and trust at the core of agentic AI success.

Emerging trends and future outlook

While enterprise expectations for agentic AI are becoming clearer, the technology itself is rapidly advancing to meet them. The ecosystem is evolving toward unified, intelligent, and governed agentic systems capable of driving scalability, interoperability, and responsible autonomy at an enterprise level. Key themes defining this evolution are:

- **Multi-agent orchestration:** agentic AI is moving beyond isolated workflows to enable coordination across business functions, such as operations, finance, IT, and customer experience, supported by emerging multi-agent communication protocols that standardize collaboration and goal-sharing among agents
- **Neuro-symbolic reasoning:** agentic intelligence is blending neural learning with symbolic logic, enabling agents to reason transparently, follow structured rules, and reduce hallucinations in decision-intensive domains
- **Cognitive memory stacks:** memory architectures are evolving into multi-layered cognitive stacks that allow agents to retain experiences, learn across interactions, and self-reflect to improve task performance
- **AgentOps infrastructure:** standardized orchestration and AgentOps platforms are advancing to provide observability, control, and lifecycle management, making agent systems more reliable, scalable, and enterprise-ready. This evolution is reflected in offerings such as LangSmith (LangChain), AgentCore (AWS), and Mosaic AI (Databricks)
- **Pre-built verticalized agentic systems:** pure-play providers such as Lyzr and Ema are developing pre-built, domain-specific agentic systems enriched with contextual knowledge and workflows, accelerating adoption across industries such as BFSI, healthcare, and retail
- **Embedded governance and explainability:** governance is becoming an embedded capability within agentic systems, featuring real-time monitoring, policy enforcement, and auditability to ensure transparency and accountability as autonomy deepens. Microsoft's Azure AI Foundry's Agent Service and Azure AI Content Safety exemplify this trend

As these advances unfold, agentic AI is poised to evolve from experimentation to an intelligent connective layer within the enterprise, linking systems, processes, and decisions through adaptive, context-aware collaboration.

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

Conclusion

Findings from Everest Group's enterprise survey affirm that while agentic AI adoption is still developing, enterprise momentum is accelerating. The current investment landscape reflects a pragmatic shift from experimentation to institutionalization. Organizations are moving beyond pilots to establish governance structures and invest in data, infrastructure, and talent foundations that will underpin long-term success. What began as isolated pilots funded through innovation budgets is now maturing into strategic investment blueprints.

However, challenges such as unclear RoI, data readiness, and integration complexities continue to temper progress. Overcoming these barriers will require discipline, cross-functional collaboration, and governance rigor to align financial strategy with measurable business outcomes.

Looking ahead, enterprises are seeking secure, interoperable, and adaptable platforms that embed agentic AI seamlessly into business workflows with minimal disruption and maximum scalability. The coming phase of enterprise adoption will be defined by how effectively organizations institutionalize discipline around investment, governance, and ecosystem orchestration, transforming early experimentation into enduring transformation and competitive advantage.

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

Appendix

Research methodology

The survey covered 112 enterprises, each with annual revenues exceeding US$1 billion, with the majority operating on a global scale. Among these, 88 enterprises had embarked on their agentic AI journey more than six months ago, indicating a growing level of engagement and experimentation in this emerging domain. Respondents included senior business executives responsible for digital transformation, AI and analytics, intelligent automation, and innovation functions. The participant group also comprised C-suite roles, such as Chief Information Officers (CIOs), Chief Technology Officers (CTOs), Chief Financial Officers (CFOs), and Chief Transformation Officers, as well as business unit heads, all possessing deep knowledge of and visibility into their organizations' automation and agentic AI initiatives.

The organization size, industry split, and geographic distribution of respondents are illustrated in Exhibit 17.

Exhibit 17: Respondents' profile by organization size, industry split, and geographic distribution

Source: Everest Group (2026)

Split of respondents by **global annual revenue** (100% = 112)	
US $1-4.99 billion	31%
US $5-9.99 billion	13%
US $10-24.99 billion	21%
US $25-50 billion	17%
More than US$50 billion	18%

Split of respondents by **geography** (100% = 112)	
North America	52%
Europe, Middle East, and Africa	21%
Asia Pacific	19%
UK and Ireland	08%

Split of respondents by **industry** (100% = 112)			
Healthcare and life sciences	20%	Manufacturing	10%
Banking, financial services, and insurance	18%	Professional services	09%
Hi-tech, telecom, media and entertainment	17%	Travel, transportation, and logistics	06%
Retail and consumer packaged goods	14%	Others[3]	06%

3 Others include energy and utilities, oil and gas, metals, mining, and other natural resources, public sector and real estate

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 NTT DATA

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):

Samikshya Meher, Practice Director
samikshya.meher@everestgrp.com

Divya Chandak, Senior Analyst
divya.chandak@everestgrp.com

Abhay Sharma, Analyst
abhay.sharma@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.