

The Rise of Agentic AI

Reimagining Enterprise Execution and Intelligence



Modernizing enterprise operations isn't just about jumping on the latest tech trends or automating routine tasks anymore. Nowadays, true transformation means rethinking how work gets done. It's about moving beyond passive, assistive tools to intelligent systems that can act, adapt, and deliver business outcomes autonomously. Organizations are often held back not only by outdated infrastructure but also by the limitations of traditional automation, which tends to be rule-bound and siloed. These constraints slow down agility, decision-making, and ultimately stifle innovation. That's where Agentic AI comes in, offering a new way of doing things. Imagine intelligent agents that operate with contextual awareness and goal-driven autonomy across complex enterprise environments.

At the core of Agentic AI is a significant shift: AI systems evolving from supporting roles to active collaborators in enterprise workflows. Unlike conventional automation or early generative AI applications that simply produce outputs based on prompts, agentic AI systems are designed to understand high-level objectives, break them down into actionable tasks, make real-time decisions, and execute workflows with minimal human intervention. This capability is already transforming various industries—from banking and healthcare to manufacturing and retail—by speeding up response times, improving accuracy, and enabling continuous adaptation to change. The result is an enterprise that is not only more efficient but also more resilient and competitive in a rapidly changing marketplace.

This Everest Group whitepaper, supported by NTT DATA, dives into the transformative power of agentic AI, offering a strategic framework for adoption, governance, and value realization. It explores the evolution of agentic AI from earlier automation models, highlights industry-specific use cases, and maps the technology provider landscape. Importantly, it addresses the key considerations and challenges enterprises must navigate—from data readiness and integration to security, compliance, and workforce adaptation. By combining advanced agentic capabilities with a solid organizational strategy, enterprises can move beyond incremental gains and unlock sustained competitive advantage, ushering in a new era of intelligent, adaptable operations.



Nitin Bajaj
Vice President, Generative AI
NTT DATA, Inc.



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

September 2025

The Rise of Agentic AI

Reimagining Enterprise Execution and Intelligence

Contents

- 03 Introduction
- 04 Introduction to the gen AI evolution journey
- 06 Agentic AI – the next big shift
- 09 Enterprise adoption landscape
- 14 Navigating the agentic AI ecosystem
- 17 Roadblocks to adoption
- 19 Strategic considerations for enterprise adoption
- 23 The road ahead

Introduction

In recent years, gen AI has captured the enterprise imagination. It has reshaped how enterprises approach innovation, efficiency, and competitiveness by enabling faster idea development, streamlining operations, and enhancing customer experiences. Its early use cases, such as content generation, summarization, and conversational assistance, have demonstrated real gains in productivity.

But as adoption scales, business leaders are looking beyond assistive tools toward systems that can act. The next phase centers on systems that can deliver measurable business outcomes, not just generating passive outputs. This demand has catalyzed the emergence of agentic AI, an advanced class of AI systems that combines goal orientation, contextual awareness, and autonomous execution to operate across complex enterprise environments with minimal human intervention.

Agentic AI represents a strategic shift in how AI is applied across the enterprise. These systems extend beyond static workflows by enabling goal-oriented execution, continuous learning, and real-time responsiveness. This transformation results in a new enterprise capability, where intelligent agents perform tasks, manage workflows, and collaborate with humans in real time. Unlike traditional automation, agentic AI is not limited to

fixed rules or structured inputs. It thrives in environments where decision complexity, information asymmetry, and responsiveness matter. This makes it relevant not only for operations and IT, but for finance, customer experience, product design, and beyond.

This Viewpoint explores agentic AI's transformative potential and provides insights into:

- Agentic AI's evolution from earlier AI and automation models
- Industry-specific use cases where agentic AI is already delivering value
- Agentic AI technology providers ecosystem
- Key adoption considerations, including RoI, data readiness, security, and governance
- Roadblocks and challenges in agentic AI adoption
- Emerging themes in the agentic AI market

The rise of agentic AI is not just a technology milestone. It represents a new architecture of enterprise execution, where intelligent systems move from supporting roles to active collaborators.

Organizations that embrace this shift thoughtfully will unlock not just efficiency, but resilience, innovation, and competitive advantage in an increasingly AI-driven world.

Introduction to the gen AI evolution journey

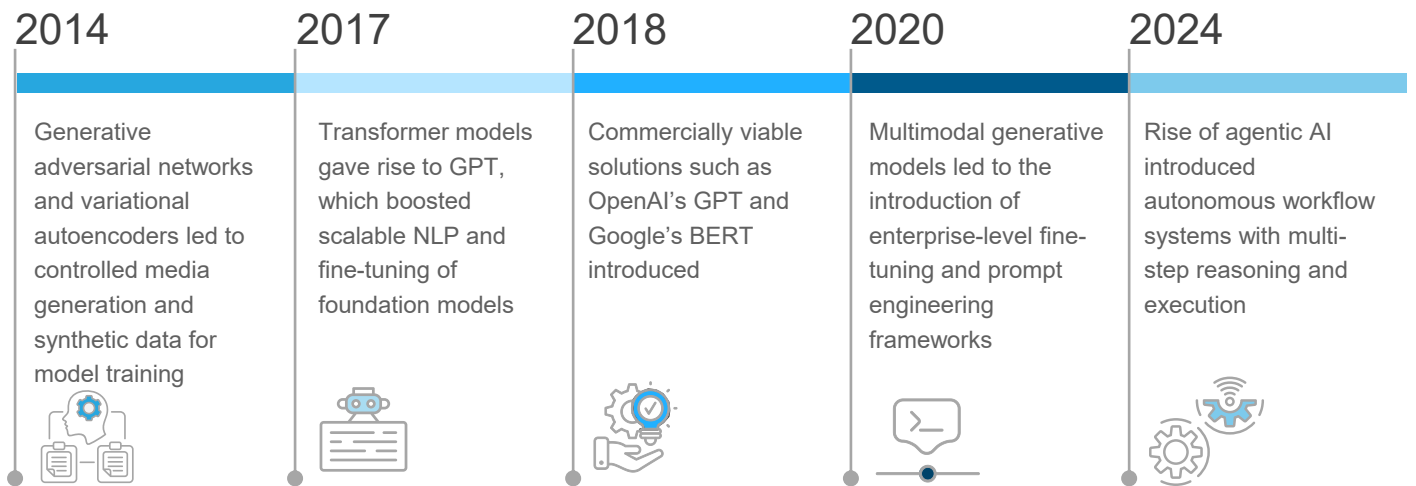
In less than three years, gen AI has gone from a breakthrough technology to a boardroom priority. Sparked by the launch of ChatGPT in 2022, it quickly gained executive attention for real business impact generation through Natural Language Processing (NLP), image generation, and automated decision-making.

Unlike traditional AI, gen AI's intuitive interfaces make it accessible across the enterprise, enabling non-technical teams to accelerate decisions, personalize experiences, and streamline workflows without relying heavily on IT teams. This shift marks a turning point in AI adoption, bringing innovation closer to the business and unlocking new levels of agility, speed, and creativity.

Exhibit 1 represents the evolution of gen AI capabilities over the past decade.

Exhibit 1: Evolution of gen AI capabilities

Source: Everest Group (2025)



Enterprises have been actively investing in gen AI, launching pilots at a rapid pace and exploring use cases across functions. While enthusiasm remains high, most organizations are proceeding with a degree of caution. Concerns around RoI, computational costs, and exposure to security and compliance risks have made many enterprises favor controlled pilots and Proofs of Concepts (PoCs) to identify tangible use cases. Despite this measured approach, pilot-to-production conversion rates doubled to 30% in 2024, highlighting growing executive confidence in gen AI's business value.¹

¹ Based on Everest Group's proprietary data of over 30 AI service providers

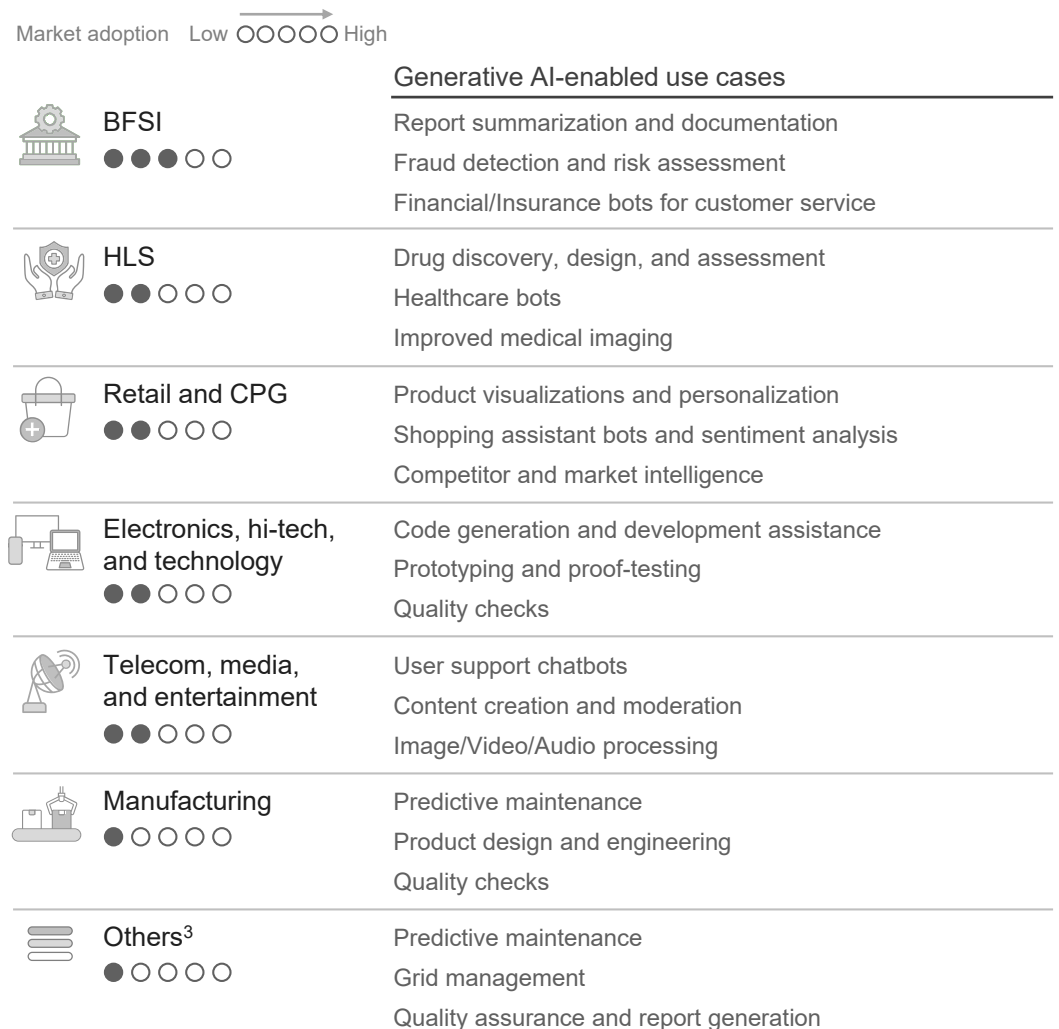
This confidence is translating into stronger commitments. According to Everest Group’s Mapping the Next: Key Priorities For 2025 and Key Issues for 2024 Survey, generative AI has emerged as the top investment priority for enterprises in 2025.² Notably, 54% of organizations are actively looking for partners to help embed gen AI within broader solutions, while 47% are investing in workforce upskilling to build internal gen AI capabilities.

Industries with mature digital ecosystems and strong data foundations, such as Banking, Financial Services, Insurance (BFSI), Healthcare and Life Sciences (HLS), and Retail and Consumer Packaged Goods (RCPG), are at the forefront of gen AI adoption. Their early successes highlight gen AI’s ability to automate routine work, enhance personalization, and reduce time-to-market for innovation.

Exhibit 2 highlights gen AI’s adoption across industries and key use cases.

Exhibit 2: Industry-wise gen AI adoption

Source: Everest Group (2025)



² Everest Group Surveys – Mapping the Next: Key Priorities For 2025 and Key Issues for 2024
³ Others include government & public sector, energy & utility, and professional services, among others

As enterprises move beyond initial gen AI deployments, their investment priorities are evolving. Attention is now shifting toward advanced capabilities that enhance performance, scalability, and impact. Organizations are experimenting with custom large and small language models tailored to domain-specific use cases, alongside efforts to fine-tune foundation models for greater accuracy and relevance. There is also growing interest in multimodal AI systems that combine text, image, and voice inputs to enable richer interactions and broader applicability.

Yet despite these advances, gen AI's key role in most enterprises remains assistive, focused on content generation, retrieval, and summarization. The next frontier lies in agentic autonomy, where AI systems evolve from passive tools into active participants capable of interpreting goals, making context-aware decisions, and autonomously executing tasks across functions.

Agentic AI – the next big shift

Enterprise leaders are increasingly asking a more advanced question: how can AI systems move from assisting employees to acting on behalf of them – securely, contextually, and at scale? The emergence of agentic AI addresses this shift.

Unlike traditional automation or gen AI applications that provide outputs based on prompts, agentic AI introduces a new operational model where systems can interpret objectives, break them into tasks, make decisions using real-time context, and execute workflows with minimal human input. These agents can reason, adapt, and coordinate actions across enterprise systems, bringing a layer of intelligent autonomy that was previously unattainable.

For enterprises, this marks a structural evolution in enterprise operations. Traditional processes, previously defined by fixed rules and handoffs, are being reimaged as outcome-oriented workflows where goals, not tasks, shape execution. Agentic AI acts as a connecting layer between strategy and execution, translating high-level business objectives into autonomous, responsive actions.

Agentic AI accelerates decision cycles, reduces manual workload, and scales operations without linear cost increases. It also improves responsiveness in rapidly changing conditions, whether rerouting a disrupted supply chain, triaging internal incidents, or optimizing resource allocation across markets. By embedding intelligence directly into operational layers, organizations can delegate complexities to machines while empowering human teams to focus on strategy, innovation, and governance.

Consider a fraud detection system built on traditional automation that follows fixed rules and flags predefined anomalies. While effective in structured environments, such systems are less adaptable when faced with new or evolving threat patterns. An agentic system, by comparison, analyzes transaction patterns, learns from evolving behaviors,

applies contextual reasoning, simulates multiple scenarios, and autonomously triggers appropriate responses. Its ability to adjust in real time enhances performance in dynamic and high-risk environments. For enterprises, this capability translates to reduced response times, improved accuracy in strategic decisions, and minimized operational disruptions.

“AI agents will become our digital assistants, helping us navigate modern-world complexities. They will make our lives easier and more efficient.”

– Jeff Bezos, Founder and CEO, Amazon

Exhibit 3: Differences between traditional Intelligent Automation (IA) and agentic AI

Source: Everest Group (2025)

■ Business value/outcome

Traditional IA		Agentic AI
Executes defined, repeatable tasks	Autonomy Faster execution, reduced human intervention	Manages complex, adaptive processes
Task-driven with predictable outcomes	Advanced reasoning More accurate and context-aware decisions	Goal-driven; breaks down objectives into executable steps
Limited autonomy; depends on rules and human inputs	Adaptive planning Greater operational resilience	Makes independent, data-driven decisions with minimal oversight
Basic transactional interactions between bots	Multi-agent operation Streamlined and dynamic operations	Agents coordinate, prioritize, and delegate collaboratively
Follows static workflows and predefined rules	Workflow optimization Improved user and customer experiences	Continuously analyzes and optimizes workflows in real time
Static systems requiring manual updates or retraining	Continuous learning Continuous process improvement	Learns and evolves through feedback and experiences
Slow to adapt; changes depend on human intervention	Adaptability Increased agility and responsiveness	Responds dynamically to real-time data and changing environments
Inflexible integrations limited to specific platforms	System interoperability Enhanced scalability and flexibility	Connects seamlessly across environments

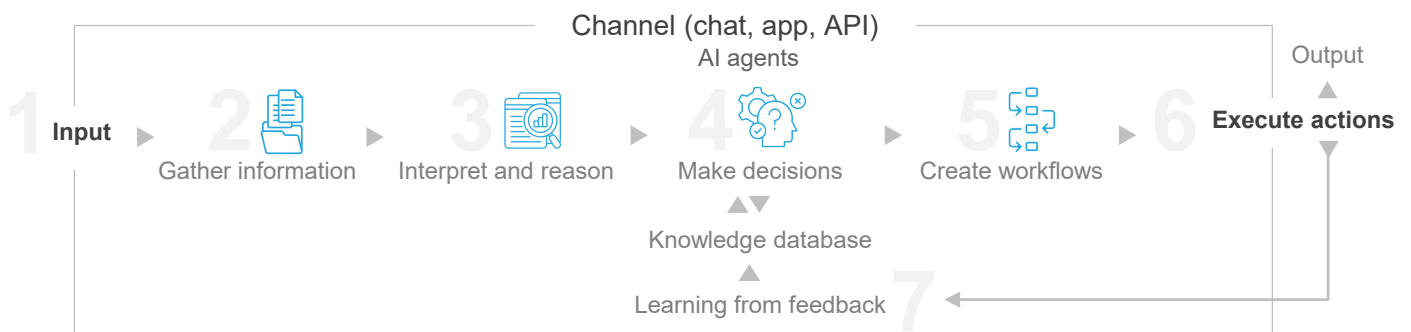
How agentic AI works

Agentic AI systems function through a dynamic, multiphase workflow designed to interpret goals, make context-aware decisions, and execute actions autonomously. Unlike earlier AI systems that relied on static inputs and linear logic, agents continuously learn, adapt, and respond in real time.

Exhibit 4 depicts a simplified workflow of an agentic AI system.

Exhibit 4: An agentic AI system's workflow

Source: Everest Group (2025)



At a high level, a typical agentic AI system operates through the following steps:

- 1. Receiving input:** The agent receives input through natural language (for example, chat, voice), APIs, or web applications. Using language models and intent recognition, it interprets the goal and extracts task requirements
- 2. Gathering information:** The system analyzes contextual signals and retrieves relevant data from memory systems such as vector databases and knowledge graphs. This process ensures informed decision-making grounded in enterprise knowledge
- 3. Interpreting and reasoning:** Once the context is understood, the agent defines high-level objectives using domain models and knowledge frameworks to convert instructions into actionable goals
- 4. Making decisions (AI agents):** It then evaluates options using deterministic rule-based engines and probabilistic AI models, simulates potential outcomes, and selects an optimal execution path
- 5. Creating workflows:** Task planning and orchestration tools convert the chosen decision path into an executable workflow, which the automation agents then carry out
- 6. Executing actions:** The agent executes the tasks, monitors results in real time, and adapts actions, if required, based on system responses or new data through a continuous feedback loop
- 7. Learning from feedback and updating knowledge database:** The system captures feedback from execution outcomes and updates its knowledge base, enabling continuous learning and improved decision-making over time

This workflow enables agentic systems to operate with precision and flexibility while responding to new information and evolving objectives without manual reprogramming.

Enterprise adoption landscape

As agentic AI transitions from concept to capability, enterprise adoption is unfolding along two parallel dimensions: maturity over time and applicability across industries and business functions. Together, these lenses help clarify where value is being realized today and how adoption is likely to scale in the future.

Adoption trajectory: from pilots to embedded intelligence

Enterprise agentic AI adoption is expected to follow a clear maturity curve, evolving from isolated experiments to deeply embedded capabilities within enterprise operating models. The pace of adoption will depend on internal readiness, partner ecosystem maturity, and measurable business outcomes.

Exhibit 5: Agentic AI maturity curve

Source: Everest Group (2025)


Low High

Intelligence: cognitive capabilities, including understanding, reasoning, and planning

Autonomy: the ability to act independently and perform tasks without human input

Scope: degree of integration and application breadth of agentic AI within an enterprise

Near term (12-24 months)



- Point solutions / PoCs to assess agentic AI's potential
- Limited budget and risk-taking due to uncertainties
- Growing ecosystem of open-source and commercial frameworks
- Major tech companies integrating agentic capabilities into their platforms

Medium term (2-3 years)



- Transitioning beyond isolated PoCs to enterprise-wide adoption
- More plug-and-play platforms lowering barriers to adoption
- Early agent collaboration implementations (multi-agent orchestration)
- Confidence in the market to drive tangible ROI
- Increased executive buy-in and funding for agentic AI initiatives

Long term (3-5+ years)



- Structural integration of agentic AI into enterprise architectures
- Reimagining work beyond traditional task execution toward end-to-end workflows
- Sophisticated orchestration of multiple specialized agents
- Dynamic collaboration between humans and agentic AI in decision-making and operations
- Organizational strategy shaped by agentic capabilities
- Regulatory frameworks and standards for ethical autonomy, safety, and explainability

In the near term, most enterprises will remain focused on PoCs aimed at evaluating agentic AI's potential, driven by point solutions and use cases that demonstrate quick wins without overhauling architectures. C-suite stakeholders will prioritize business impact over technical experimentation, making use cases tied to efficiency, personalization, and decision augmentation more attractive. Agents in this phase are typically task-specific and operate in low-risk environments.

In the medium term, adoption will move beyond isolated PoCs toward more coordinated, enterprise-wide deployments. As early implementations yield tangible ROI and operational reliability, executive confidence and investments will increase. Workforce readiness and trust in autonomous systems may become key barriers, requiring intentional change management. Key developments during this phase will include the emergence of multi-agent orchestration, where agents collaborate across functions and systems to support complex workflows. More modular and plug-and-play platforms will reduce implementation barriers, enabling faster scaling and wider adoption across business units.

In the long term, agentic AI will be structurally embedded in the enterprise operating model. Organizations will shift from automating individual tasks to designing goal-driven workflows, where agents operate autonomously, apply contextual reasoning, and deliver business outcomes with minimal oversight. At this stage, agentic AI will become not just a productivity tool but a differentiator in strategic agility and innovation. This phase will require new approaches to enterprise architecture, governance, and talent management. Agentic AI will no longer be seen solely as an enabler, but as a collaborative layer in strategic decision-making and execution.

As enterprises progress in their adoption journey, the role of partners, particularly service providers, will remain essential across all phases. In the near term, providers will help accelerate capability-building and manage early-stage pilot initiatives. As needs evolve in the medium term, enterprises may deepen their collaboration with both service and technology providers to co-develop integrated solutions tailored to specific use cases. In the long term, these partners will continue to play a strategic role, supporting complex transformation efforts and helping design and operationalize multi-agent systems that align with enterprise-wide governance, compliance, and performance objectives.

Agentic AI adoption across industries and functions

Agentic AI is moving from concept to capability across a range of industries and enterprise functions. While still in the early stages of adoption, its impact is becoming visible in areas where scale, decision complexity, and responsiveness are vital.

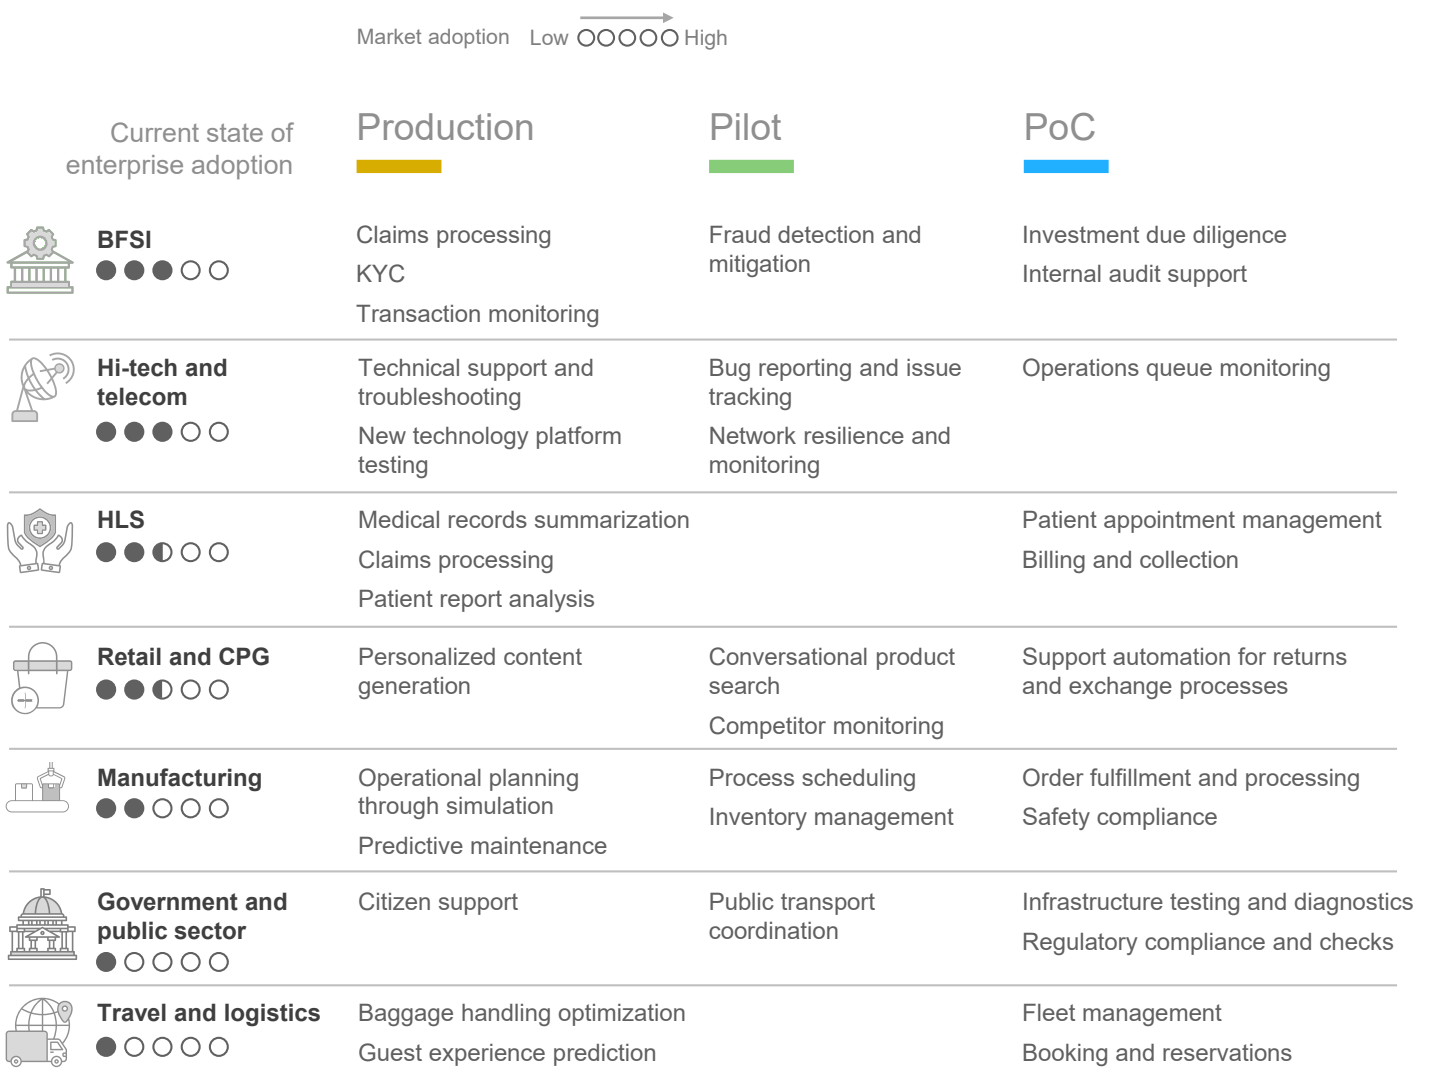
Industries with mature digital ecosystems, high operational complexities, and strong innovation imperatives are at the forefront of agentic AI adoption. Sectors such as BFSI,

HLS, and manufacturing are translating early investments into tangible business value. Their implementations are not only streamlining core functions but also setting strategic benchmarks for scalable, agentic AI-enabled transformation. A leading US bank, for instance, is taking a structured approach to embedding agentic AI into high-impact areas such as fraud detection, credit risk modeling, and customer interactions to enhance operational efficiency, customer satisfaction, and employee experiences.

Exhibit 6: Agentic AI adoption across key industries

Source: Everest Group (2025)

[ILLUSTRATIVE]



In healthcare, agentic AI is streamlining claims processing, triage, and appointment scheduling. These applications enhance service delivery while maintaining compliance and clinical oversight. Manufacturers are applying agentic AI to improve predictive maintenance, quality control, and supply chain responsiveness. These systems help

reduce downtime and manage operational risks, particularly in asset-heavy and regulated environments. Similarly, RCPG enterprises are using agentic AI to consolidate consumer insights and enable real-time decisioning support.

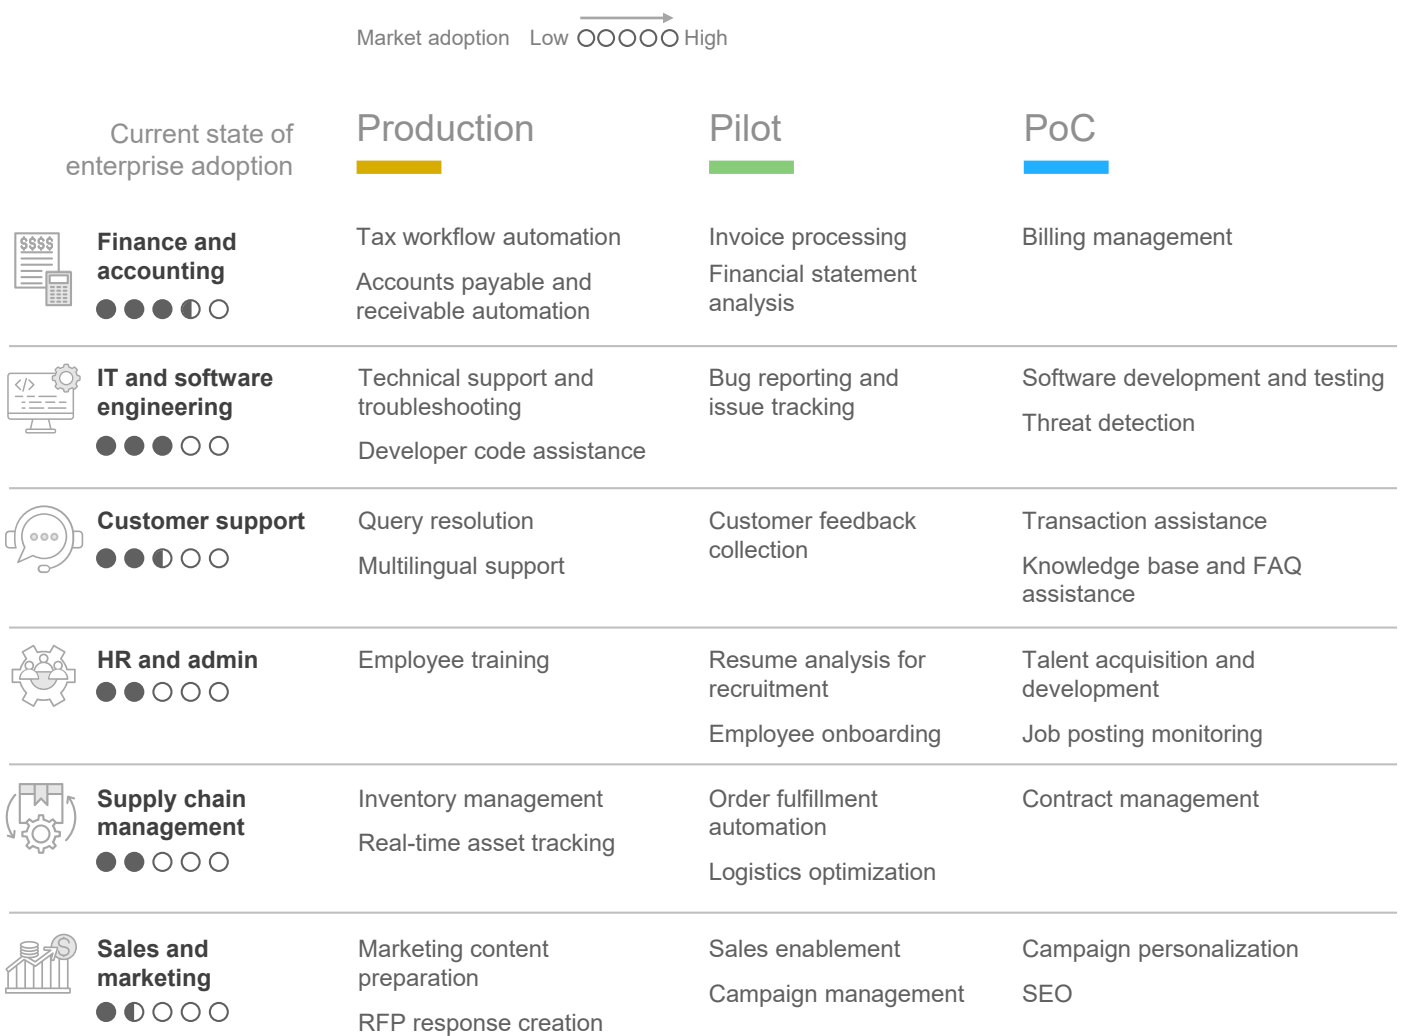
Other sectors, including hi-tech and telecom, public, and travel and logistics, are progressing through early-stage adoption focused on targeted use cases such as software testing, citizen support, and baggage handling optimization. These efforts are laying the foundation for broader integration.

Beyond industry-specific adoption, agentic AI is gaining traction across core business functions, transforming how enterprises manage operations, interact with customers, and enable workforce productivity.

Exhibit 7: Agentic AI adoption across key business functions

Source: Everest Group (2025)

[ILLUSTRATIVE]



Functions such as finance and accounting, IT and software engineering, and customer support are seeing practical deployments, especially where work is high-volume, repetitive, or requires quick decision-making.

Enterprises are deploying agents to support activities such as document processing, anomaly detection, technical troubleshooting, and customer interactions. These systems operate across applications, learn from context, and provide consistent support, reducing manual workload while improving speed, accuracy, and user experiences. For example, in finance and accounting, agents reconcile records and support audit readiness; in customer service, they resolve multi-turn queries with contextual intelligence; and in HR, they streamline onboarding and accelerate training by guiding employees through tasks and information.

These functional deployments signal a broader shift in how work is structured and scaled within the enterprise. By taking on routine tasks and first-level decision-making, agentic AI systems enable teams to redirect effort toward exception handling, strategic analysis, and continuous improvement. This shift not only drives efficiency but also enhances the organization's ability to scale operations without a corresponding increase in cost or headcount.

Other enterprise functions are also beginning to explore agentic AI, particularly where personalization, responsiveness, or coordination across multiple systems is required. In supply chain operations, agents are being used to track orders, adjust logistics in real time, and manage inventory more proactively. In sales and marketing, they support campaign optimization, generate tailored content, and assist with lead qualification. While adoption in these areas is still emerging, early implementations suggest strong potential for driving agility and improving time-to-market.

Navigating the agentic AI ecosystem

As adoption accelerates, enterprise leaders are faced with a fast-evolving ecosystem of agentic AI providers, each varying in capabilities, maturity, and integration approach. Selecting the right partner requires not only technical alignment but also a clear understanding of how each category fits within the broader enterprise architecture and AI strategy.





















Overview of major technology providers shaping the agentic AI ecosystem

The agentic AI ecosystem comprises several distinct cohorts. Hyperscalers and BigTech firms deliver scalable, end-to-end agentic capabilities through their cloud and AI platforms. Enterprise platforms and IA providers are embedding agents into their core offerings to enhance existing workflows. Pure-play agentic AI companies are focused on building flexible, customizable agent frameworks. Native AI platforms cater to organizations with strong internal AI capabilities and a need for highly specialized solutions. In addition, service providers are playing a key role in facilitating adoption, integration, and change management.

Exhibit 8: Agentic AI provider ecosystem

Source: Everest Group (2025)

[ILLUSTRATIVE]

Provider type	Service providers			
BigTech	 Bedrock Agents	 AgentSpace, Vertex AI Agent Builder	 watsonx	 AutoGen
Enterprise platform players	 Joule	 Agentforce	 Xanadu	
Pure-play agentic AI providers	 taskade	 Ema	 newo.ai	 Relevance AI
IA providers	 AI Agent Studio	 kore.ai	 OneReach.ai	 Agent Builder
Native AI platforms	 ANTHROPIC  H2O.ai	 c3.ai Agentic AI Platform	 dataiku	 DataRobot

There are five main types of agentic AI providers

1 **BigTech providers:** such as Bedrock by AWS, AutoGen by Microsoft, Watsonx by IBM, and AgentSpace by Google

Differentiators

BigTech companies differentiate themselves through their access to vast computational resources, foundational models, extensive security capabilities, and cloud-native infrastructure. They also offer integrated services that support enterprises across agentic AI development, deployment, and scaling phases, along with providing some pre-built agents to accelerate adoption

Challenges

These solutions may lack domain specificity or vertical depth and often require significant internal expertise to customize and govern them

2 **Enterprise platforms:** such as Joule by SAP, Agentforce by Salesforce, and Xanadu by ServiceNow

Differentiators

They stand out by integrating agentic capabilities natively into specific domain areas such as CRM, ERP, and ITSM, enabling faster adoption and operational scalability across departments. Their low/no-code environments also empower business users to create and manage agents more easily

Challenges

While integration into existing workflows is relatively straightforward, these providers offer limited flexibility outside their platforms and are constrained to predefined use cases

3 **Pure-play agentic AI providers:** such as newo.ai, taskade, Relevance AI, and Ema

Differentiators

Their core strength lies in domain specialization and agility, rapidly introducing new capabilities and offering platforms that support highly customized agents tailored to specific enterprise use cases

Challenges

Pure-play providers often face integration challenges, as their solutions typically operate outside native enterprise environments. As a result, integration, tuning, and customization require careful execution to ensure reliability and alignment with existing systems and workflows

Value proposition

They offer comprehensive, end-to-end AI solutions that enable enterprises to build highly scalable and secure agentic AI systems, making them ideal for large organizations seeking to integrate AI deeply into existing digital ecosystems

Enterprise platforms have embedded agentic AI directly into their core business applications, allowing enterprises to enhance workflow efficiency without needing to overhaul existing systems. Their solutions make AI adoption seamless by extending familiar enterprise software.

Pure-play agentic AI providers offer specialized, highly customizable agent solutions tailored to specific industries and functions. They offer flexibility, orchestration, and customization to enterprises, often at a lower cost compared to traditional providers

4 **IA providers:** such as UiPath, Automation Anywhere, and Kore.ai

Differentiators

These providers offer a unique advantage by blending traditional RPA with intelligent agent orchestration, enabling a natural progression from rule-based automation to dynamic, AI-driven workflows. For organizations with mature automation programs, they provide familiar deployment models and fast time-to-value

Challenges

These providers are relatively new entrants to the agentic AI market. Building strong differentiation beyond their established RPA roots and gaining enterprise trust for broader agentic AI deployments remains a key challenge

5 **Native AI platforms:** such as C3.ai, DataRobot, H2O.ai, and Anthropic

Differentiators

They offer full control over model life cycles, deployment, and compliance, and are well-suited for highly regulated or specialized use cases. Their agentic AI offerings are a natural extension to enterprises already leveraging these platforms for their previous AI and gen AI projects

Challenges

These platforms may have high technical entry barriers, requiring deep in-house AI expertise and sustained investment in talent, infrastructure, and governance

Value proposition

IA providers extend their strong automation heritage into the agentic AI space, enabling enterprises to seamlessly orchestrate traditional bots and decision-capable AI agents

Native AI platforms provide customizable environments to develop, train, and scale sophisticated agentic models. They focus on explainability, transparency, and fine-grained control, allowing enterprises to deploy agentic solutions faster

Choosing the right provider cohort depends on several factors, including the organization's digital maturity, existing technology stack, AI strategy, and preferred pace of adoption. Enterprises seeking scale and stability may prioritize hyperscaler or platform-based solutions that integrate seamlessly with existing infrastructure. Those looking for flexibility, experimentation, or specialized capabilities may find stronger alignment with pure-play agentic AI or native AI platforms. In all cases, success will hinge on the ability to ensure interoperability, establish robust governance, and align agentic capabilities with clear business outcomes.

Exhibit 9 provides a simplified view of how different types of agentic AI solution providers stack up against the key enterprise priorities they are best suited to address.

Exhibit 9: Provider selection decision matrix

Source: Everest Group (2025)

[ILLUSTRATIVE]

Provider type	Best fit for (enterprise priorities)
BigTech providers	Scalability, multi-agent orchestration, cloud-native infrastructure, cross-system orchestration, security, and mature support and developer ecosystems
Enterprise platforms	Ease of integration with ERP/CRM, faster deployment, lower integration complexities, and mature support and developer ecosystems
Pure-play agentic AI providers	High customization, niche/domain-specific use cases, explainability, agile multi-agent orchestration, and lower ownership costs
IA providers	Easier evolution from RPA systems, lower integration complexities, and familiar tooling
Native AI platforms	High customization, high degrees of control, explainability, and niche/domain-specific use cases

This matrix is intended to serve as a reference point for evaluating different types of agentic AI providers. It offers a structured framework to compare providers across key decision-making dimensions and aligning their strengths with enterprise priorities. This matrix will enable decision-makers to compare provider categories against key enterprise priorities such as scalability, speed of deployment, customization, control, and cost efficiency. In some cases, a hybrid strategy may be optimal, allowing different parts of the organization to prioritize different provider types based on their unique requirements.

Roadblocks to adoption

What is holding enterprises back

Agentic AI offers significant promise, but scaling it effectively requires addressing enterprise-wide challenges that stem from both organizational constraints and broader ecosystem limitations. Understanding these barriers is essential to mitigating risks and accelerating adoption:

- **Unclear RoI and business case challenges:** Agentic AI often delivers value in areas difficult to quantify using traditional cost-savings models, such as faster decision-making, reduced process latency, and enhanced user experiences. These benefits contribute to business agility and risk mitigation but rarely map to simple before-and-after comparisons. Without clear value frameworks, executive buy-in can stall

- **Security and compliance risks:** The autonomy of agentic systems introduces new risks related to access control, auditability, and regulatory exposure. Enterprises must account for scenarios where agents access sensitive systems or trigger downstream processes without direct human oversight. This raises the stakes for identity and access management, breach containment, and real-time monitoring, especially when agents interact with third-party APIs or cross internal system boundaries. According to Everest Group's Survey, data privacy and security is among the top three challenges that enterprises face in scaling their AI initiatives for 2025.⁴ Beyond data access, risks also stem from agents making decisions that, while technically permissible, could violate policy or compliance standards if not properly governed
- **Data fragmentation and quality:** Agentic systems depend on real-time, structured, and trustworthy data. Legacy systems, siloed data stores, and inconsistent governance undermine the agents' ability to reason, act, and adapt
- **Infrastructure and integration limitations:** Real-time orchestration requires scalable, API-first infrastructure and event-driven workflows. Enterprises with legacy IT stacks may find it difficult to operationalize agents at scale
- **Lack of internal expertise:** Designing, deploying, and monitoring autonomous agents requires new capabilities such as prompt engineering, agent orchestration, behavior tuning, and risk governance – capabilities most enterprises are still building. In many cases, enterprises also face broader knowledge gaps, particularly among frontline employees unfamiliar with how agentic systems operate or how to effectively supervise their decisions. These gaps can slow adoption, reduce trust in outputs, and lead to underutilization even when the technology is in place
- **Workforce readiness and resistance:** When AI agents begin taking on decision-making roles traditionally held by humans, they can trigger fear, resistance, or disengagement among employees. Perceived threats to job security and professional relevance often lead to passive resistance, low engagement during pilot phases, and distrust in agent-led recommendations. This reluctance to collaborate with AI systems can undermine change management efforts and stall broader adoption
- **Regulatory uncertainties:** From the EU AI Act to US executive orders, evolving legislation creates ambiguity around agentic AI system governance. Enterprises face a patchwork of overlapping, and at times conflicting, mandates, making legal and compliance teams cautious and slowing executive decision-making on agentic AI investments

Enterprises must take a proactive approach to overcome these roadblocks – investing in clear communication, cross-functional alignment, and continuous learning. Addressing foundational issues such as data readiness, infrastructure maturity, workforce capabilities, and governance is vital. These challenges should be viewed not as blockers, but as design inputs. Enterprises that tackle them head-on will be better positioned to scale agentic AI confidently and unlock its full potential.

⁴ Everest Group Surveys – Mapping the Next: Key Priorities For 2025 and Key Issues for 2024

Strategic considerations for enterprise adoption

For agentic AI to scale from isolated pilots to enterprise-wide capabilities, organizations must address key considerations across business alignment, technical readiness, operational oversight, and governance maturity. Doing so helps enterprises not only avoid failure but maximize value and scale with confidence.

Exhibit 10 highlights the key factors enterprises must address to effectively adopt agentic AI.

Exhibit 10: Key considerations for agentic AI adoption

Source: Everest Group (2025)

Organizational considerations



Business alignment and use-case fit
Leadership buy-in and change readiness
Total cost of autonomy

Technical considerations



Build versus buy versus co-create strategy
Data and infrastructure readiness
Scalability and enterprise integration

Security and compliance considerations



Security, access control, and fallback planning
Regulatory compliance and legal defensibility
Governance and autonomy management

Operational considerations



Agent orchestration
Human-in-the-loop design and workforce collaboration
Training and upskilling

Organizational considerations

These factors form the foundation for integrating agentic AI into existing ecosystems while mitigating risks and optimizing outcomes.

- **Business alignment and use case fit:** Agentic AI should be deployed where autonomy delivers distinct value. That includes workflows with high decision volume, complexity, or latency. Use case selection must align with enterprise priorities, whether reducing operational drag, accelerating resolution cycles, or enabling round-the-clock responsiveness

- **Leadership buy-in and change readiness:** Endorsement from business and functional leaders is essential. Adoption efforts often stall due to resistance from middle management or uncertainty around accountability. Successful implementations involve not just leadership approval, but active change stewardship
- **Total cost of autonomy:** Beyond implementation costs, enterprises must consider ongoing investments in agent retraining, performance monitoring, drift correction, and compliance oversight. RoI models should include both strategic value and long-term ownership costs

Technical considerations

Technical readiness is a vital pillar in the successful enterprise deployment of agentic AI. From infrastructure decisions to data preparedness, each factor directly affects implementation efficiency and long-term scalability.

- **Build versus buy versus co-create strategy:** Organizations must weigh the benefits of building in-house agentic AI systems versus adopting commercial platforms or co-creating solutions with partners. While pre-built platforms may offer faster deployment, they can limit flexibility, especially for enterprises with highly specialized processes. In-house builds enable deep customization and tighter integration, but require mature capabilities in agent design, orchestration, and ongoing maintenance. Co-creation with providers may offer a middle path by balancing control, speed, and innovation. This decision should reflect operational context, use case complexity, available expertise, and long-term strategic goals. Partner selection must align with enterprise data architecture, security policies, and governance standards
- **Data and infrastructure readiness:** Agentic AI relies on high-quality, contextual, and connected data ecosystems. This foundation requires not just data access, but real-time integration, semantic consistency, and the use of memory architectures such as vector databases and knowledge graphs. To support dynamic decision-making and autonomy at scale, organizations also need scalable infrastructure and modular, interoperable APIs
- **Scalability and enterprise integration:** While pilots help validate feasibility, enterprise value requires scaling agentic AI across teams, regions, and functions. This effort involves not only functional integration and workflow redesign, but also infrastructure readiness, performance optimization, and monitoring frameworks. Enterprises must ensure agentic systems can scale without introducing exponential complexity or operational fragility. The ability to support wide adoption while maintaining control and adaptability is essential to long-term success

Governance and compliance considerations

With agentic AI, security concerns extend beyond unauthorized access. The real risk lies in autonomous agents making context-sensitive decisions that affect interconnected systems. A single failure or breach can trigger the domino effect, compromising legal defensibility, eroding user trust, and disrupting workflows often without immediate detection.

- **Security, access control, and fallback planning:** Agents often require broad system access and deep integration with sensitive environments. Enterprises must predefine access boundaries, monitoring frameworks, and fallback protocols to prevent unintended actions and support incident recovery
- **Regulatory compliance and legal defensibility:** Autonomous agents must operate within established legal and regulatory frameworks, especially in sectors such as BFSI, healthcare, and public services. Enterprises should ensure agents comply with jurisdiction-specific regulations, can be audited retrospectively, and provide decision traceability that stands up to legal scrutiny
- **Governance and autonomy management:** Effective deployment requires clearly defined policies for agent oversight, escalation paths, and inter-agent coordination. Governance models must evolve to support continuous learning, performance tracking, and alignment with enterprise risk appetite

Operational considerations

Operationalizing agentic AI goes beyond deploying models, it requires carefully designed systems that integrate seamlessly into enterprise workflows.

- **Agent orchestration:** As agentic AI deployments scale, enterprises must manage how multiple agents interact, share context, and sequence decisions. Without clear orchestration, agents with overlapping or conflicting objectives can generate inefficiencies or unintended outcomes. Operational models should define how agents communicate, escalate, and defer to each other to maintain coordinated execution across workflows
- **Human-in-the-loop design and workforce collaboration:** Successful agentic AI adoption depends on employees knowing when and how to engage with autonomous systems. Agentic AI must enable meaningful human oversight, especially in high-stakes scenarios where ethical, regulatory, or judgment-based intervention is required. This approach includes onboarding frontline staff to agent behaviors, defining handoff protocols, and establishing clear points for human review, correction, or escalation
- **Training and upskilling:** Enterprises must continuously invest in reskilling employees to understand, architect, and supervise agentic AI systems. Building internal expertise ensures agents are trusted, correctly used, and continuously improved

Successful agentic AI enterprise adoption requires a multidimensional strategy that aligns organizational goals, ensures operational and technical preparedness, incorporates regulatory vigilance, and fosters strong ecosystem partnerships. Addressing these dimensions cohesively creates the foundation for scalable, sustainable adoption and increases the likelihood of achieving meaningful business outcomes.

The following checklist illustrates the critical organizational, technical, governance, and operational considerations enterprises must evaluate when adopting agentic AI. It provides organization leaders with a starting point to identify readiness gaps, prioritize investments, and guide structured agentic AI implementation. This is not intended to be a one-time evaluation, but an ongoing reference to guide adoption maturity over time.

Exhibit 11: Agentic AI adoption checklist

Source: Everest Group (2025)



Organizational considerations

- ✓ Do we have use cases where autonomy delivers measurable value (for example, high-volume or latency-sensitive workflows)?
- ✓ Do we have committed sponsorship from both business and functional leaders?
- ✓ Do we have clarity on costs beyond initial deployment (such as retraining, monitoring, and compliance updates)?



Technical considerations

- ✓ Do we have a clear build versus buy strategy tailored to organization's AI maturity and use case needs?
- ✓ Do we have a vetted list of technology and service partners?
- ✓ Do we have scalable and modular infrastructure with interoperable APIs to support dynamic agent operations?
- ✓ Do we have a strategy to scale agentic AI pilots to enterprise-wide deployment without introducing complexity?



Operational considerations

- ✓ Do we have well defined interaction rules between agents to ensure coordinated execution?
- ✓ Do we have decision points where human oversight is required and how agents will defer or escalate accordingly?
- ✓ Do we have training programs to upskill employees on maintaining and supervising agentic AI solutions?



Governance and compliance considerations

- ✓ Do we have access boundaries, monitoring protocols, and recovery mechanisms in place?
- ✓ Do we have clear agent logging and traceability aligned with regulatory needs?
- ✓ Do we have governance frameworks in place to monitor, escalate, and adjust agent behavior over time?
- ✓ Do we have inter-agent coordination to ensure alignment with enterprise goals and risk appetite?

The road ahead

Future of agentic AI in enterprises

Agentic AI is not a short-term trend, it is a long-term architectural evolution in how enterprises operate. While most organizations are still in the early stages of adoption, the trajectory is clear. As capabilities mature, governance frameworks evolve, and trust builds, agentic AI will move from supporting back-end tasks to shaping how decisions are made, work is delegated, and outcomes are delivered across the enterprise. These systems will become integral to the modern workplace, helping businesses innovate, scale, and stay competitive.

Customer-facing adoption is on the horizon

Thus far, most enterprises limited the use of agentic AI to internal tasks, citing concerns around trust and reliability for customer-facing applications. However, as these systems demonstrate value in low-risk scenarios, organizational confidence is growing. Forward-thinking enterprises have started experimenting with agentic AI in external-facing roles such as personalized customer service and intelligent sales support. This trend signals a turning point in enterprise readiness to entrust agentic AI with high-impact, customer-facing functions.

New organizational capabilities will define competitive advantage

As adoption scales, businesses must invest in new capabilities to responsibly and effectively manage agentic AI. These capabilities will be vital in maximizing the value of agentic AI while minimizing operational and reputational risks. Some of these evolving capabilities may include:

- **Agent orchestration layers** to securely and efficiently coordinate complex, multi-agent workflows across systems and functions
- **Dedicated teams** focused on monitoring agent behavior, ensuring compliance, and fine-tuning autonomy boundaries
- **New human-agent roles** such as AI interaction designers, agent trainers, and collaboration ethicists, to ensure productive and ethical integration of agentic AI with the human workforce
- **Budgeting for end-to-end agentic AI life cycle** as costs shift from one-time deployment to sustained management, demanding ongoing investments in oversight, iteration, and cross-functional coordination

New workplace: realigning the human workforce in the age of agentic AI

As agentic AI becomes more deeply embedded into core business functions, a profound realignment of the human workforce is underway. This shift is not just about augmenting roles, it involves redefinition and, in some cases, replacement of traditional roles with fully autonomous agentic systems. Some enterprises such as Amazon and CrowdStrike have already announced headcount reductions as the reliance on evolving AI capabilities increase.^{5, 6}

To navigate this transition effectively and ethically, businesses must invest in a new generation of hybrid human-agent roles, such as:

- **AI interaction designers** who shape intuitive, context-aware interfaces for agentic collaboration
- **Agent trainers** responsible for refining the behavior of agents through feedback loops and contextual learning
- **Collaboration ethicists** who define and uphold ethical standards for human-agent interactions, ensuring fairness, transparency, and psychological safety
- **Agentic workflow orchestrators** who design, manage, and optimize the flow of work between human teams and autonomous agents, ensuring efficient collaboration, clear task handoffs, and cohesive operations in hybrid environments

Enterprises that proactively realign their workforce, retrain talent, and embed these new roles into their operating model will be best positioned to lead in the agentic era.

RoI will shift from efficiency to adaptability

While early agentic AI deployments are often justified by efficiency gains, the real value lies in adaptability: faster responses to market signals, scalable problem-solving, and reduced operational latency. This value creation demands a reframing of RoI from static savings to dynamic resilience.

Stakeholder trust will shape pace of adoption

As agents take on more autonomous roles, stakeholders, including customers, employees, and regulators, will demand clarity on how decisions are made, when humans are involved, and how errors are resolved. Trust will hinge not only on visibility, but on embedding reversibility, accountability, and compliance into agent design and governance.

The trajectory of agentic AI reflects a familiar pattern in enterprise innovation: cautious exploration, scalable development, and ultimately, a transformation of organizational operations. Yet, what differentiates agentic AI is the scale and depth of transformation it enables. Agentic AI introduces a new class of digital actors that can engage dynamically across workflows and functions. The path forward is not solely about developing more advanced agents but about preparing enterprises to engage with them effectively.

Organizations that build the right capabilities and foster strong human-agent collaboration will be well positioned to lead in the era of intelligent autonomy.

⁵ Based on article [Seattle-based Amazon announces 'workforce reduction' due to AI](#) published by Fox13 in June 2025

⁶ Based on the article [CrowdStrike announces 5% job cuts, says AI is 'reshaping every industry'](#) published by Reuters in May 2025

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):

Abhivyakti Sengar
abhivyakti.sengar@everestgrp.com

Shubham Yadav
shubham.yadav@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.