Agentic AI 2026

A Mid-market Playbook for Adoption and Scale

Nitesh Bansal
Managing Director & CEO,
R Systems

Foreword

Agentic AI marks a decisive shift in the enterprise technology journey - from systems that inform decisions to systems that execute them. We believe this is not simply the next feature in the AI evolution; it is the right direction for enterprises seeking durable competitive advantage.

The capabilities available today are extraordinary. Tools are abundant. Models are increasingly powerful. Practical use cases span engineering, IT operations, finance, customer service, and beyond. In theory, the opportunity appears limitless.

In practice, enterprises face real constraints.

Mid-market organizations operate within complex, legacy-heavy environments. Core systems are deeply embedded. Data architectures are fragmented. Governance models are evolving. Skills in AI oversight, orchestration, and risk management are still maturing. At the same time, teams are experiencing tool fatigue, navigating an expanding ecosystem without clear guidance on what truly drives durable value.

Scaling agentic AI in this environment requires more than deploying new technology. It demands architectural discipline, defined ownership models, embedded trust frameworks, and a deliberate progression of autonomy aligned to business risk. Most importantly, it requires a focus on measurable value in clearly defined, high-impact use cases. Execution, not experimentation, is the differentiator.

And as systems move from assistance to action, autonomy must advance in step with accountability. Greater autonomy does not diminish the role of people, it elevates it, making human oversight, domain expertise, and decision accountability even more critical.

This report, commissioned by R Systems and produced by Everest Group, reflects those realities. It does not assume greenfield environments. It does not assume unlimited budgets. It addresses the operational friction enterprises actually face, and outlines practical pathways to scale agentic AI within those constraints.

What distinguishes the leaders highlighted in this report is not the number of tools they deploy, but how deliberately they confront integration complexity and technical debt. They are embedding agentic AI into high-volume, well-governed workflows within legacy-heavy environments, rather than isolating it in pilots. They are addressing architectural fragmentation, strengthening governance in parallel with capability, and modernizing their enterprise fabric while continuing to run mission-critical operations.

Agentic AI will increasingly move from supervised assistance to controlled execution across core business functions. Our objective with this report is to equip enterprise leaders with the clarity and practical guidance needed to convert AI ambition into governed, measurable, and sustained execution.

Everest Group®

March 2026

Agentic AI 2026

A Mid-market Playbook for Adoption and Scale

Contents

03 Introduction

04 Adoption and maturity – where organizations stand today

08 Crossing the implementation chasm

13 Human-AI collaboration and oversight

18 Planning for scale – governance, ecosystem roles, and commercial models

22 The way ahead

25 Appendix

Introduction

The center of gravity in enterprise AI is shifting from assistance to action. Leaders are no longer asking how AI can help their teams – they are asking where AI can safely act on their behalf with context and at scale. Agentic AI is the engine behind this shift. It introduces a new operating model: instead of returning an output to a prompt, these systems can interpret a business goal, break it into tasks, use live context to decide and execute workflows – escalating to people only when judgment, policy, or accountability requires it. A perfect mix of technological readiness, enterprise demand, and operational pressure makes this an exceptionally strong moment to adopt agentic AI.

To understand how organizations are responding, we surveyed approximately 200 global enterprise leaders (see appendix for details on the sample). Most enterprises are currently in the pilot stage of their agentic AI journey, with several key patterns emerging:

- Agentic AI adoption is fastest in technology-heavy, data-rich areas such as IT operations, software engineering, supply chain, and customer support, while functions such as finance, HR, and sales/marketing remain mostly in early exploration
- Unclear business cases, complex system integration, an immature ecosystem of partners and tools, organizational change resistance, and ongoing security and privacy concerns slow down scaling
- Autonomy is limited and task-specific, but enterprises expect a noticeable increase over the next two years, particularly in customer support, IT operations, IT engineering, and sales/marketing – while retaining human oversight for higher-risk actions
- Only a small share of organizations has well-defined agentic AI policies, leaving a gap between scaling ambitions and the maturity of their governance frameworks
- To build confidence, enterprises are investing in strong data governance, transparency, human-in-the-loop controls, default logging and audit, phased pilots, and policies embedded directly into daily workflows
- Success hinges on a coordinated ecosystem – hyperscalers providers, orchestration platforms, enterprise software, service partners, and academia – working together to provide the tools, integrations, and skills enterprises need

This report also lays out a practical playbook for adopting agentic AI, helping enterprises move from scattered pilots to durable, scaled deployments. It distills where to start by function, how to prioritize use cases, and which guardrails and operating models are needed to build trust. In addition, it outlines the roles of ecosystem partners and internal stakeholders, providing leaders with a structured path to capture near-term value while laying the foundations for longer-term autonomy.

Adoption and maturity – where organizations stand today

As agentic AI advances from concept to capability, its adoption is unfolding along two complementary dimensions: maturity over time and applicability across industries and business functions. Viewed together, these dimensions clarify where value is emerging today and how adoption is likely to scale next.

Mapping current progress and readiness to adopt agentic AI across various enterprise segments

Enterprises pursue diverse paths in their agentic AI adoption, from sequential progression to targeted leapfrogging and parallel, multi-track approaches. Across these routes, firms typically manifest as explorers, pilots, or scalers, each with distinct needs and conditions. Exhibit 1 shows current AI adoption levels among enterprises across all three stages.

Exhibit 1: Current state of AI adoption

Source: Everest Group (2026)

	 Explorers	 Pilots	 Scalers
Enterprise characteristic	Enterprises in discovery mode, validating feasibility and fit. Work is limited to low-risk, reversible use cases (for example, prototypes and sandboxes) with humans firmly in the loop, basic guardrails, and light integration	Enterprises running controlled, production-adjacent trials to prove value, reliability, and safety on live data and systems	Enterprises operationalizing agents across functions, achieving task-level autonomy in well-governed workflows and selective process autonomy where audit and reversibility are strong
	29% of the enterprises stand at this stage, running sandboxes and assist-mode agents with light integrations	Majority of the enterprises, about **57%**, are at this stage, making it the center of gravity for adoption	A minority, about **15%,** of enterprises have moved beyond pilots to operational task autonomy in well-bounded workflows

Survey results indicate that agentic AI is beginning to shift from evaluation into early operational use. Most enterprises are experimenting with autonomy in areas where rules are well-defined and outcomes are easy to measure, while scaled deployment remains limited to a few functions.

Across revenue bands, agentic AI is largely still in the pilot phase. In the mid-market, approximately 57% of enterprises are running controlled pilots, positioning this segment as the primary near-term opportunity to convert pilots into scaled rollouts. Only about 15% of enterprises have reached scale, indicating that key enablers such as policy frameworks, integration mechanisms, and clearly defined outcomes are only now starting to take shape. At the smaller end of the market, nearly 30% of enterprises remain in the early explorer stage.

By industry, adoption maps closely to digital intensity. Technology and telecom sectors are in the scaling cohort; internet companies are driving the bulk of pilots; BFSI is advancing cautiously and systematically; and healthcare remains concentrated in the early exploration stage.

Exhibit 2 shows enterprise AI adoption journey.

Exhibit 2: Enterprise AI adoption journey

Source: Everest Group (2026)

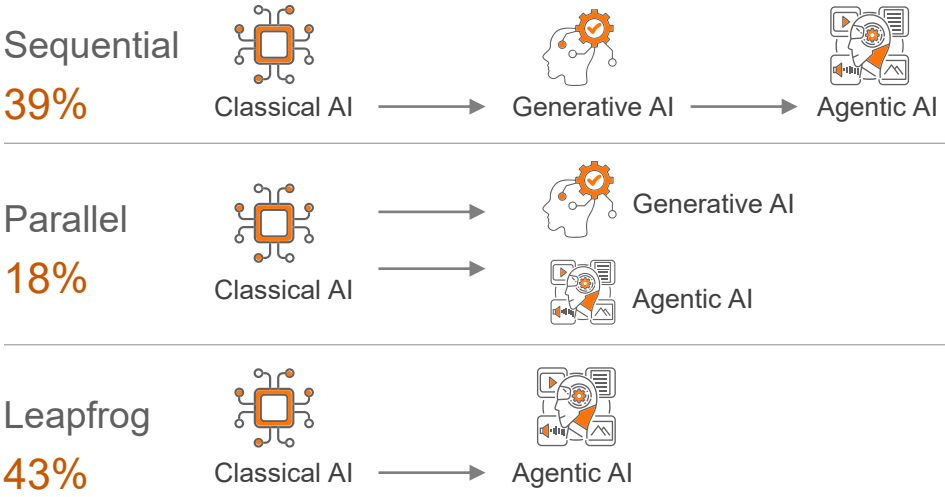

Enterprise adoption journeys typically follow three progression patterns: a cohort that moves directly to autonomous, multi-step agents (Leapfrog: 43%); a cohort that adopts multiple forms of AI simultaneously (Parallel: 18%); and a cohort that advances step-by-step from classical automation and generative AI-assisted workflows toward higher autonomy (Sequential: 39%).

The finding that 43% of enterprises report leapfrogging to agentic AI suggests two interpretations: competitive pressure may be driving organizations to signal rapid progress, or terminology slippage may be at play, with robotic process automation workflows or chatbot-based automation labeled as agentic systems.

Assessing agentic AI feasibility and adoption hurdles for mid-market enterprises

With experimentation giving way to execution, the center of gravity has shifted to business operations and enterprise leaders are now assessing which functions are most feasible for near-term agentic AI deployment under clear guardrails. Across business functions, feasibility concentrates where work is well-instrumented, high-volume, and time sensitive.

Exhibit 3 highlights the enterprise feasibility and adoption state of agentic AI in different business functions.

Exhibit 3: Agentic AI feasibility and adoption across different business functions

Source: Everest Group (2026)

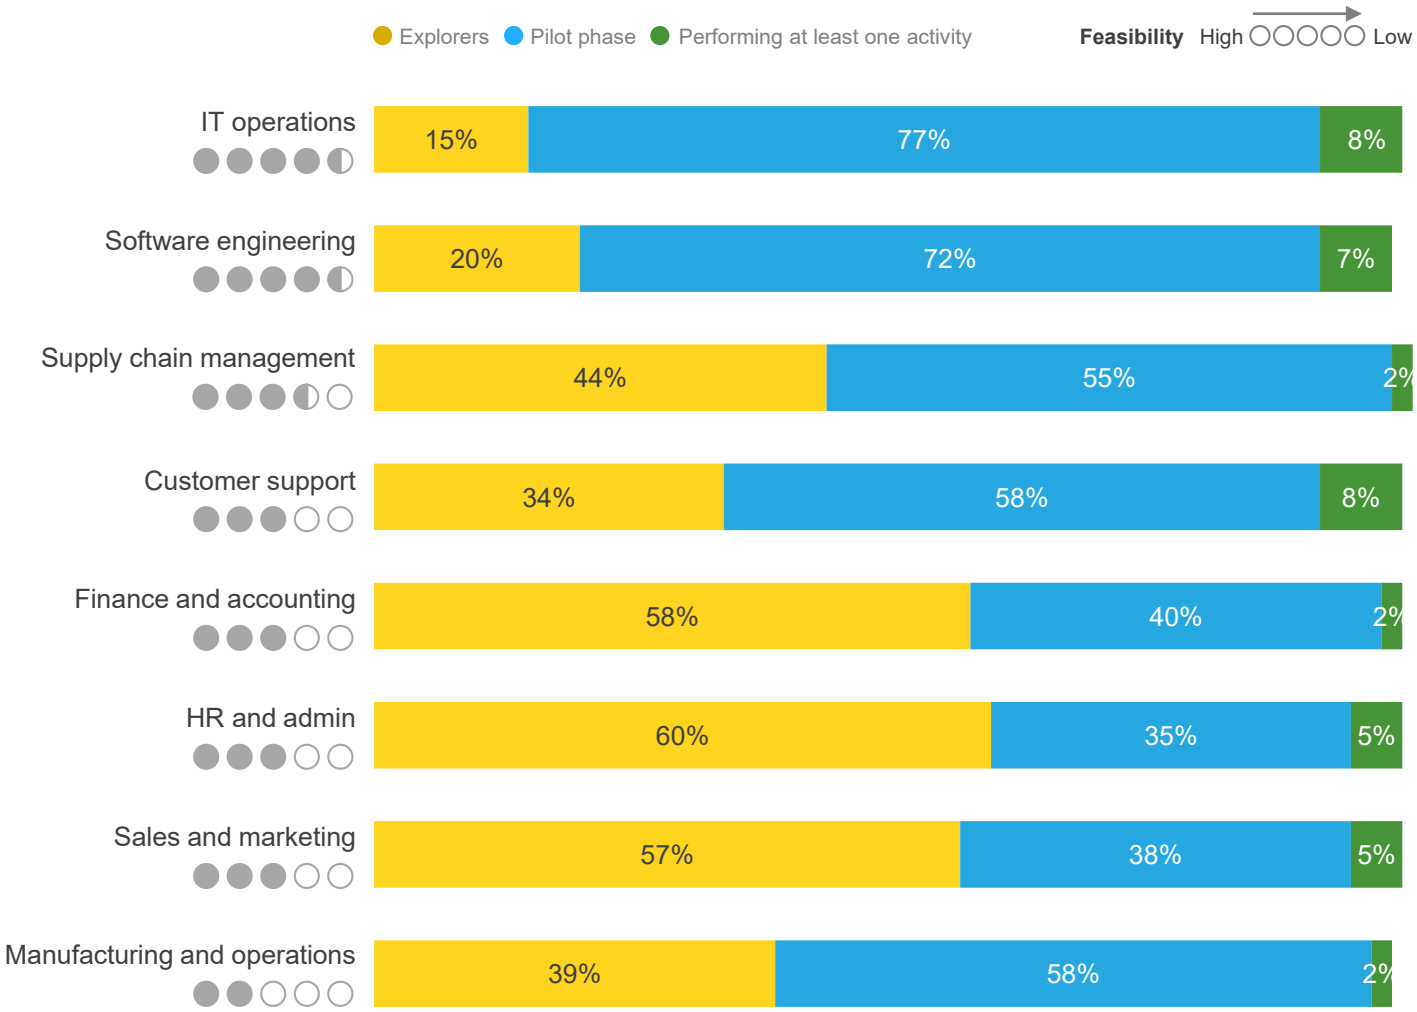

Enterprise leaders see agentic AI as most feasible in technology-heavy areas such as IT operations and software engineering, with most organizations currently operating in the pilot phase.

Customer support is also progressing, as codified knowledge and clear business rules let agents handle multi-turn requests and propose compliant resolutions, improving response consistency and speed.

Commercial and people functions are advancing more selectively. Finance and accounting see traction in reconciliations and close activities where dual control and auditability are built in, while HR applies agents to onboarding, case routing, and policy guidance under well-defined privacy constraints. In operational domains, effective adoption depends on cross-system coordination.

Regardless of where enterprises are on their agentic AI adoption journey, they encounter a consistent set of hurdles that make progress to scaled adoption difficult. Exhibit 4 illustrates the key challenges enterprises face in their agentic AI adoption journey.

Exhibit 4: Hurdles to agentic AI adoption

Source: Everest Group (2026)

43%

Clarity on business case / Return on Investment (RoI)

Many pilots demonstrate effort saved, faster turnaround, or higher first-pass quality, yet those gains are not tied to baselined KPIs or to a defined payback window

40%

Integration complexities

Challenging to integrate multiple disparate systems, preventing agents from operating seamlessly across end-to-end processes

38%

Lack of matured ecosystem

Limited availability of proven partner solutions, standards, and best practices makes decision-makers cautious about committing to large-scale investments

35%

Change resistance

Enterprise leaders are uncertain about new ways of working, which slows organizational adoption

33%

Security and privacy concerns

Security and privacy concerns persist until runtime guardrails such as limits, dual controls, auditability, and tested rollback mechanisms are demonstrably enforced

Evolving from exploration to expansion, guided by realistic RoI windows

Across enterprises, the shift from pilot to scale is being paced by defined RoI windows rather than open-ended experimentation. Leaders are sequencing investments into two broad buckets: near-term and longer-term value.

Near-term efforts typically emphasize cost optimization, customer-experience improvements, and engineering productivity, where outcomes are measurable and guardrails are already well understood.

Longer-term efforts focus on innovation, new revenue models, and end-to-end operational redesign, where payback is meaningful but depends on tighter integration, policy enforcement, and data readiness.

With this baseline in place, we next look at which business functions can unlock the most near-term value and what it will take to cross from pilots to scale.

Crossing the implementation chasm

Uncovering business functions that deliver most value and efficiency improvement

Across industries, enterprises breaking out of pilot purgatory and reaching scale converge on a common pattern: start where value is concentrated and measurability is high. Leaders are prioritizing functions with digital exhaust, mature guardrails, and tight ties to systems of record. Exhibit 5 highlights the business functions poised to gain the most value from agentic AI.

Exhibit 5: Business functions poised to gain most value from agentic AI adoption

Source: Everest Group (2026)

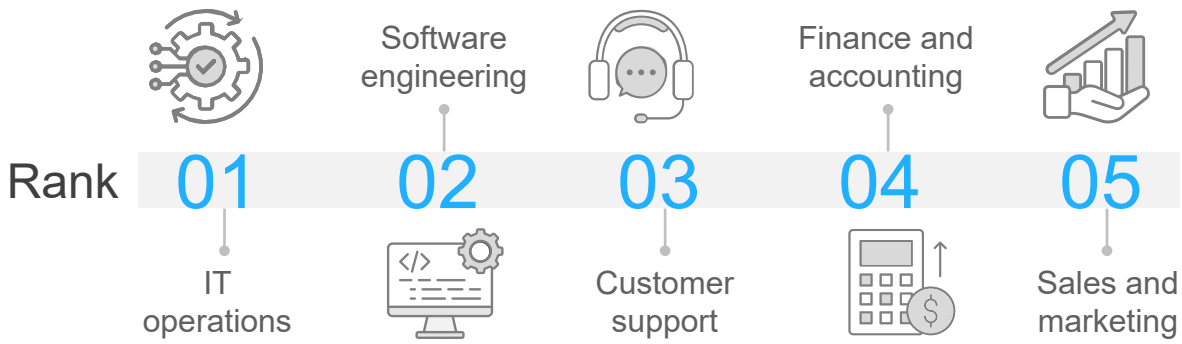

Survey results indicate a clear pecking order in where agentic AI is already delivering measurable efficiency at scale. The strongest impact appears in high-volume, rules-led operations that still require light contextual interpretation – for example, classifying inputs, routing work, handling exceptions, and selecting from a bounded set of actions. Five business areas stand out for the fastest, most reliable returns, each combining high work volume, strong digital exhaust, and well-governed action surfaces into core systems.

- **IT operations** is the most scale-ready domain for measurable value and efficiency. Incident triage, change management, and detect-diagnose-remediate loops are becoming semi-autonomous. Agents read alerts, correlate signals across observability stacks, draft root-cause analyses, and trigger runbook actions through ITSM and configuration tools with humans approving high-risk steps, reducing toil and escalations
- **Software engineering** serves as a potent engine for productivity when governed by strong quality gates. Teams are moving beyond code completion to agentic patterns across the SDLC – story decomposition, project scaffolding, test generation and repair, dependency upgrades, and secure refactoring – wired into CI/CD pipelines for verifiable impact
- **Customer support** has emerged to be the most immediate front-office beneficiary as organizations scale from deflection to resolution. Retrieval-augmented agents answer with grounded citations, assemble context from CRM/ERP/entitlement systems, and execute actions – refunds, replacements, entitlement changes – under policy, improving first-contact resolution and reducing handoffs
- **Finance and accounting** provides a natural next wave for safe autonomy. Structured, dual-control workflows allow agents to classify and code invoices, reconcile exceptions, draft close packages, and orchestrate collections – posting only with controller approval – compressing cycle times and manual effort while strengthening working-capital discipline
- **Sales and marketing** follows with more variable realized impact. Near-term gains concentrate in operational tasks – pipeline hygiene, enrichment, lead routing, forecast commentary, campaign QA, and compliant content assembly – where strong brand guardrails and approval flows convert pilots into durable, scalable outcomes

Focusing on SDLC value hotspots for near-term impact and scale

Software engineering is the safest starting point for agentic AI. Given the work is already codified, instrumented, and governed by CI/CD and test gates, teams can raise autonomy safely and verify impact quickly. Our survey shows software engineering delivers the highest efficiency uplift (nearly 30%), making it the natural launchpad for scale.

The next step is to examine SDLC stages to identify where agentic AI can drive the fastest near-term impact and expansion. Exhibit 6 highlights the SDLC stages that leaders rank among their top three to realize near-term value.

Exhibit 6: SDLC value hotspots for agentic AI adoption

Source: Everest Group (2026)

% of respondents selecting the SDLC stage as a top near-term value driver

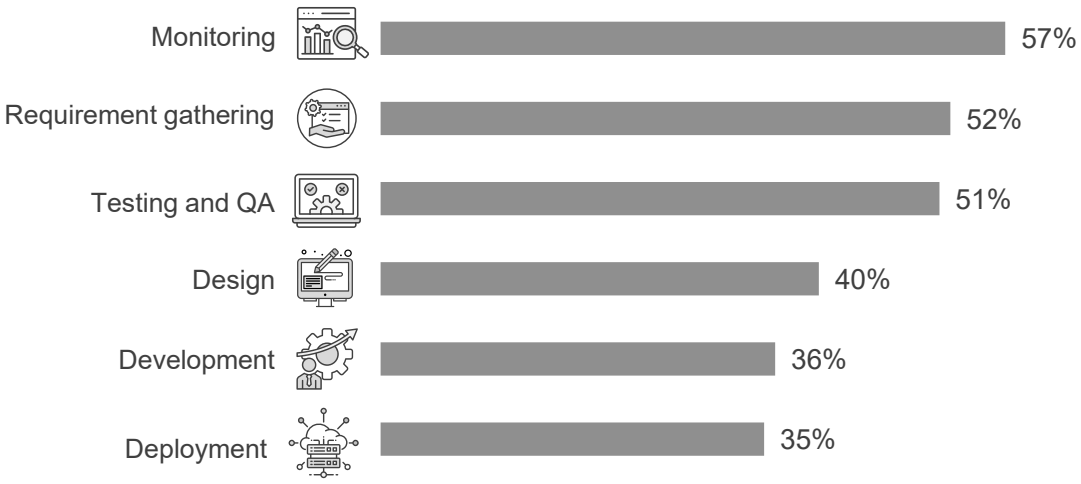

Key barriers to agentic AI adoption in SDLC (ranked as per severity level)

1. **Security and compliance**
2. **Skill gaps**
3. **Integration complexities**
4. **Change resistance**
5. **Unclear RoI**
6. **Lack of trusted partners**

Enterprise leaders see monitoring (57%) as the lowest friction starting point, where agents can spot issues, summarize root causes, and trigger safe fixes for quick wins with minimal process change.

Following monitoring, enterprises rank requirements gathering (52%) and testing and QA (51%) as the next priorities. Both functions share a similar value profile: high-throughput, well-structured work where agents can accelerate execution.

Agents can synthesize inputs into usable artifacts and strengthen test creation and prioritization, improving release readiness while operating within existing quality gates.

Over time, value compounds as agents move deeper into design, development, and deployment, supporting design reviews, safe refactoring, and controlled rollouts once guardrails are in place.

Although leaders have identified SDLC hotspots for near-term value, they point to six cross-cutting hurdles that can slow adoption:

- **Security and compliance:** Sensitive code and data raise concerns about provenance, secret exposure, and auditability, slowing approvals and production use
- **Skills gaps:** Teams lack hands-on patterns for moving from assistance to execution, which limits safe autonomy and confidence
- **Integration complexities:** Fragmented toolchains and brittle connectors make it hard to plug agents into systems of record and automate end-to-end flows
- **Change resistance:** Fears about quality, accountability, and job impact create hesitation, slowing policy sign-offs and day-to-day adoption
- **Unclear RoI:** Benefits diffuse without baselines, making impact hard to attribute and funding decisions harder to justify
- **Partner gaps:** An uneven ecosystem (assurance, references, deployment options) increases perceived risk in procurement and scale

Charting the current and expected autonomy levels for agentic AI across various business functions

Across the enterprise, agents are mostly operating in assist and pre-approval modes today, with two-year road maps pointing toward task autonomy wherever work is digital, instrumented, and governed. Functions with rich telemetry and repeatable workflows (ITOps, customer support, finance) are set to move fastest; those with heavier judgment, physical constraints, or partner dependencies (manufacturing, supply chain) will advance more deliberately. Exhibit 7 outlines the current state of agentic AI autonomy across core business functions.

Early agentic AI wins will come from tightly scoped, metrics-driven workflows – building trust through controlled automation before expanding into higher-stakes, multi-party processes.

Exhibit 7: Autonomy level of agentic AI across business functions

Source: Everest Group (2026)

Autonomy level High → Low

Business function	Current autonomy level	Tasks agents perform currently	Autonomy expected in two years
Customer support	●●●○○	Query resolution, policy-bounded requests, post-resolution surveys	●●●●○
HR and admin	●●◐○○	Policy query resolution, schedule interviews, maintain HRIS records	●●●◐○
IT operations	●●◐○○	Execute safe runbook steps, correlate alerts, suppress noises, scale infrastructure within bounds	●●●◐○
Software engineering	●●◐○○	Generate unit test scripts, run static scans, code refactoring	●●●○○
Finance and accounting	●●○○○	Invoice auto-reconciliation, classification and coding	●●●◐○
Sales and marketing	●●○○○	CRM records enrichment, A/B test emails, pop-ups	●●●●○
Supply chain management	●●○○○	Auto-reconciliations, identification of goods received note mismatches	●●●○○
Manufacturing and operations	●◐○○○	Work order creation and assignment	●●◐○○

Across core business functions, agentic AI is steadily moving beyond simple assistance into controlled execution of everyday work.

- Functions such as customer support, HR, and admin are moving from simple assistance to policy-bounded resolution of routine requests with strong audit and Personal Identifiable Information (PII) controls
- Finance and accounting are evolving from invoice tagging and low-risk reconciliations to task-level posting and collections orchestration under dual control, compressing close cycles
- IT operations and software engineering are front-runners, with agents already driving alert triage, safe runbook execution, test generation, refactoring, and gated CI/CD changes
- Manufacturing and operations see assistive autonomy in work orders, quality records, and digital instructions, with longer-term potential for supervised physical actuation under strict safety interlock

- Sales and marketing use agents to clean CRM data, run experiments, and streamline operational tasks such as lead routing and campaign QA, while brand-sensitive content stays under human review
- Supply chain agents are expanding from auto-reconciliation and mismatch detection to exception triage, ETA updates, and policy-gated reorder and reprioritization recommendations as data quality improves

Over the next two years, enterprises expect agentic AI autonomy to step up by a full notch across functions, with customer support, IT operations, software engineering, and sales and marketing approaching relatively high levels of controlled autonomy.

Human-AI collaboration and oversight

Scaling safely and quickly starts with formal, agentic-specific policies that define boundaries, permissions, auditing, and escalation. Equally important are equipped, skilled teams and clear ownership models, including named product owners, agent owners, and risk partners with defined decision rights, which turn policy into practice and lift autonomy from assistance to execution. As enterprises move from single-agent pilots to multi-agent systems, orchestration becomes essential to manage handoffs and dependencies. Agent orchestration serves as the coordination layer that assigns roles, routes work between agents, enforces guardrails, and ensures multiple agents operate as one coherent, governable system instead of disconnected bots.

Benchmarking current trust levels and formal policies in place for agentic systems

Exhibit 8 highlights the current enterprise trust levels for agentic AI.

Exhibit 8: Current trust levels for agentic AI

Source: Everest Group (2026)

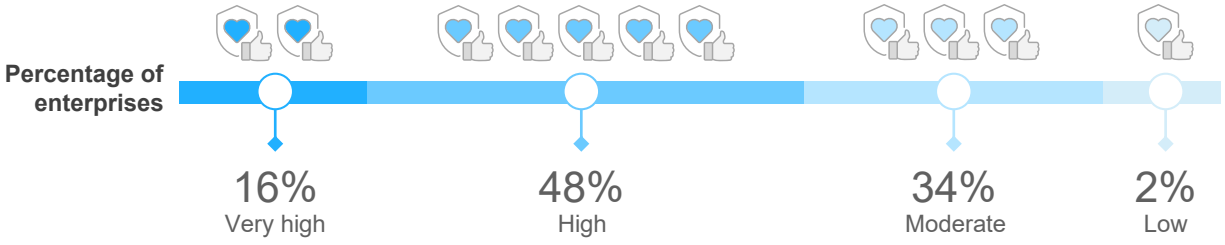

Most enterprises express they are comfortable moving beyond trials as long as humans remain in the loop and actions are auditable.

Across enterprises, trust in agentic AI is no longer tentative; it is broadly confident yet governed. Viewed through an industry lens, trust in agentic AI mirrors governance maturity, strongest in tightly controlled sectors and more uneven in faster-moving arenas:

- Regulated and process-intensive segments such as healthcare and BFSI show high trust backed by codified oversight and audit; leaders allow task-level autonomy under dual control
- Engineering-led sectors such as technology show high trust due to mature CI/CD and testing discipline; comfortable moving from assistance to execution in bounded flows
- Telecom sector shows moderate-to-high trust; reliability and regulatory obligations keep autonomy tightly scoped, with momentum in NOC monitoring and playbook execution
- Digital-native sectors such as internet and media and entertainment depict mixed confidence. Rapid release cycles encourage trials, but brand safety / user-generated content risks keep many teams in assistance or pre-approval modes

To deepen this trust, enterprises need clear agentic AI policies that guide oversight, permissions, and data use at every adoption journey stage. The policy landscape for agentic AI is still emerging. Only 7% of enterprises report having agentic-specific policies, while the majority are either developing formal policies (33%) or have very nascent ones. Around 30% remain exposed, with only generic AI frameworks (9%) or no framework at all (20%), underscoring the gap between intent and robust governance.

Calibrating enterprise and workforce readiness for effective collaboration with agents

Readiness splits into two questions: which skills matter most for scaling agents, and how prepared teams are today.

Exhibit 9 outlines enterprise workforce readiness for agentic AI.

Exhibit 9: Enterprise workforce readiness

Source: Everest Group (2026)

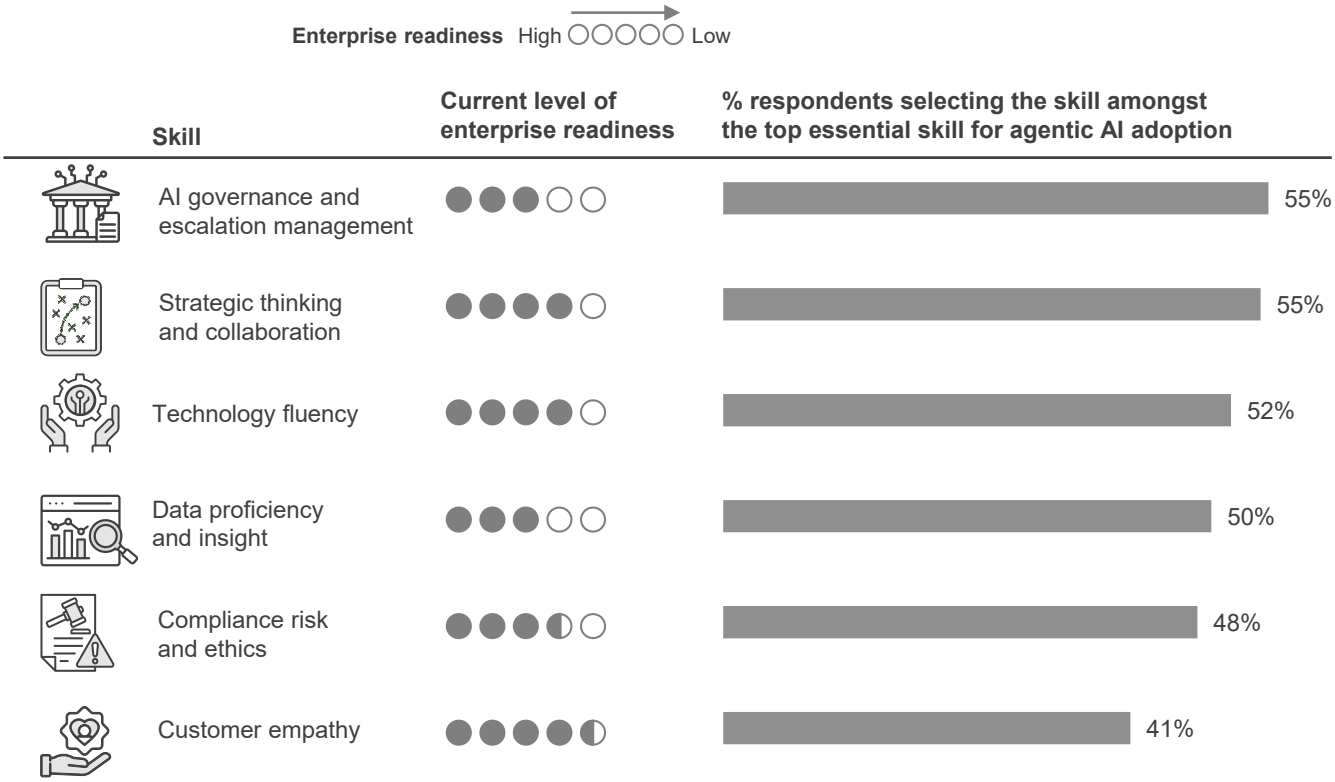

Leaders put AI governance and escalation management, strategic thinking, and collaboration at the top of the essential skills list, clearly indicating that safe autonomy hinges on decision rights, approval thresholds, and fast handoffs across functions. Technology fluency and data proficiency and insights sit close behind – practical operators needed to interpret signals, judge agent outputs, and keep pipelines healthy. Compliance, risk, and ethics are widely recognized but still need to be embedded in day-to-day workflows rather than living only in policy, while customer empathy is a relative strength that supports effective human-in-the-loop review.

When we compare importance versus current readiness, the biggest capability gaps appear in AI governance/escalation and data proficiency: both are viewed as essential, yet many organizations rate their maturity as only mid-tier. Strategic collaboration and technology fluency present smaller gaps, signaling generally higher readiness; these areas can act as accelerators for near-term scale. Compliance and ethics are established in principle but uneven in practice.

Viewed through an industry lens, readiness clusters along familiar lines: regulated sectors show stronger governance muscles.

Strengthening trust through formal measures and defined ownership models

Enterprises are strengthening trust in agentic AI along two axes:

- **Ownership and accountability:** who is accountable for agentic AI
- **Formal trust measures:** which formal measures are in place to govern day-to-day use

Ownership models

Enterprises can choose from a range of ownership models for AI depending on their appetite for control, speed, and assurance. Exhibit 10 outlines the most common approaches and where each model works best.

Exhibit 10: Ownership models

Source: Everest Group (2026)

Autonomy level High → Low

Ownership models	Enterprise preferences	Description	Best suited for
CIO/CTO-led	●●●●○	Technology leadership owns platforms, guardrails, and integrations; business consumes within defined limits	Early-scale programs needing speed, unified tooling, and clear accountability
Cross-functional governance council	●●●●○	Product, engineering, risk, legal, and security share decisions on standards, approvals, and escalations	Enterprises ready to scale use cases across functions; balances speed with assurance
Enterprise AI/trust board	●●●●○	Executive body sets policy, approves high-risk exceptions, and arbitrates enterprise risks	Large/Regulated firms needing top-down alignment and auditability
Federated hub-and-spoke	●●●○○	Central hub provides standards and shared services; BUs (spokes) implement locally	Multi-BU portfolios seeking scale with local autonomy once foundations are stable
Dual ownership (business and technical)	●●●○○	Named business owner for value pairs with a technical agent owner for safety/operations	Ideal once agents move from assistance to execution in production
Product/BU-owned (decentralized)	●●◐○○	Each product/BU runs its agents within enterprise policies and audits	Mature organizations with strong BU ops and consistent enterprise guardrails
Risk/Compliance-led	●●◐○○	Risk function leads approvals/monitoring; speed trades off for assurance	Highly regulated domains or early stages with heightened risk tolerance needs

Along their AI adoption journey, organizations often start with CIO/CTO-led ownership, introduce a cross-functional council as use cases spread, then layer dual ownership and policy-in-path to make accountability and controls visible in day-to-day workflows.

Formal measures

As enterprises scale AI, they steadily introduce trust measures that strengthen control, visibility, and reliability. Exhibit 11 highlights the actions most commonly taken along that journey.

Exhibit 11: Trust measures enterprises take along their AI adoption journey

Source: Everest Group (2026)

48% **Data governance**
Define data sources, PII handling, retention, and lineage for agent use

42% **Transparency and explainability**
Provide citations, reasoning summaries, and decision artifacts

41% **Pilot first before scale**
Stage rollouts to reduce risk and learn safely

41% **Policy-in-path enforcement**
Embed approvals, checks, and logging directly in CI/CD and workflow tools

37% **Human-in-the-loop**
Define pre-approval levels, auto-escalations, and dual control

36% **Audit-by-default**
Log prompts, context, actions, and outcomes for review and forensics

To make agentic AI dependable, organizations pair explicit ownership with a progressive stack of controls. This stack typically starts with data governance and transparency, then matures into human-in-the-loop oversight and policy-in-path enforcement as deployments scale. As this control stack matures, multi-agent orchestration becomes the operating layer that coordinates agents, manages handoffs, and enforces controls where work is executed.

Enterprise use cases increasingly move beyond one agent, one task toward multi-agent systems, where specialized agents collaborate to deliver end-to-end outcomes. In practice, orchestration becomes the control plane that coordinates:

- **Task decomposition and routing:** Breaks an outcome into discrete tasks and assigns each to the appropriate specialist agent in the right sequence
- **Tool access and permissions:** Controls which systems/tools an agent can access and what actions it can take under defined conditions
- **Quality and risk checks:** Validates outputs and actions against quality thresholds, security requirements, and policy guardrails before execution
- **Human escalation:** Routes work to a human reviewer when uncertainty, impact, or policy thresholds require judgment or approval

Done well, multi-agent orchestration improves scale and safety by standardizing how agents interact with tools, data, and humans, rather than relying on ad hoc prompt logic embedded within individual agents.

Planning for scale – governance, ecosystem roles, and commercial models

As enterprises move agentic AI from pilots into production, governance becomes the foundation for sustainable scale. An equally important enabler is the partner ecosystem, which helps translate experimentation into enterprise-wide impact. Providers play differentiated roles – some deliver control planes and guardrails, while others offer domain accelerators. Managed service and assurance partners support day-to-day operations and independent validation, so success depends on assembling the right mix of providers.

This shift brings a defining set of choices to the forefront: Should we build, partner, or buy? And if we buy, which commercial model best aligns with value so RoI remains defensible at scale?

Governance-led scaling of agentic AI from experimentation to enterprise rollout

Governance provides the foundation that moves agentic AI from promising pilots to scalable, trusted operations. By making agent behavior predictable, explainable, and accountable, governance strengthens enterprise trust in agentic AI and builds the confidence needed to move beyond controlled trials.

Without these mechanisms, enterprises accumulate agentic AI debt as they scale. This debt shows up as rapid agent sprawl across teams; duplicated or conflicting agent logic; brittle integrations and hidden dependencies across tools and workflows; inconsistent permissions; and zombie agents that persist without clear owners. The result is operational drag, harder troubleshooting, and rising compliance and incident risks.

Done well, governance prevents this debt and accelerates replication across functions. It also enables autonomy to progress from assistance to execution with measurable outcomes.

Leveraging the ecosystem to scale through build, partner, or buy decisions

Build versus partner versus buy

Exhibit 12 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. With the build-partner-buy lens set, the following view maps the ecosystem players and how each helps translate strategy into scale.

Exhibit 12: Key players in the agentic AI ecosystem

Source: Everest Group (2026)

Partner category	Differentiator	Best suited for
Hyperscalers	Secure control plane, elastic scale, native AI services	Establishing a stable platform and integrating with the existing estate
Foundation model providers	High-performing models, safety tooling, fine-tuning options	Fast experimentation and multi-model routing; task specialization
Agentic orchestration platforms/frameworks	Coordinate tool use and workflow/state with policy hooks, memory, and multi-app action routing	Building and scaling agents that act across many systems
Enterprise platforms	Ship pre-built, compliant actions/connectors inside core apps to unlock quick wins with auditability	Standardized workflows with clear SLAs and audit
Providers	Integrate agents with systems of record and pipelines; industrialize patterns for scale	Complex integration, platform build-out, enterprise rollout
Niche AI firms	Offer focused agentic capabilities and rapid innovation tailored to specific tasks/domains	Targeted use cases, fast pilots, differentiated features
Academic institutions	Contribute cutting-edge research, benchmarks, and talent pipelines to de-risk frontier methods	Co-development, evaluations, and skills development for advanced projects

This landscape is a concise snapshot of the agentic AI ecosystem, outlining the major provider categories, what differentiates them, and where each is best applied. It helps leaders align provider strengths with enterprise priorities, scalability, speed to deploy, customization, controls, and cost efficiency and make informed build-partner-buy choices. In practice, a hybrid mix is common, with teams selecting different categories based on use case maturity, governance needs, and time-to-value.

Aligning pricing to business value to ensure RoI-backed sustainable scaling

Once the enterprises have chosen the right mix of partners for its agentic AI journey, the next move is deciding how to structure the engagement.

Enterprises sentiment point to a clear order of preference across commercial models:

- Platform subscriptions emerge as the most favored construct for scale because they bundle shared guardrails, support, and tools under one predictable fee – ideal once multiple teams are using a common control plane
- Consumption-based pricing is the next preference, valued early for transparency and budget control (per user/API/workflow) and as a way to prove unit economics before broader rollout
- License-based models follow for standardized capabilities with stable usage where procurement simplicity matters
- Interest in tier-based pricing is steady, especially where functionality and autonomy rise in stages (for example, assist → pre-approval → execute) and governance maturity is still ramping
- Outcome-based pricing is increasingly viewed as an evolutionary endpoint for commercial constructs in agentic AI. It becomes more feasible as enterprises strengthen baseline definitions, attribution logic, and measurement discipline. As these capabilities mature, organizations can transition from input- or effort-linked arrangements to structures tied directly to demonstrable business impact

Exhibit 13 presents a pricing model selection framework mapped to two axes – measurement maturity (clear baselines, KPIs, and auditable impact) and usage predictability (consumption stability) – and overlays them across adoption journey stages.

Exhibit 13: Pricing model selection framework

Source: Everest Group (2026)

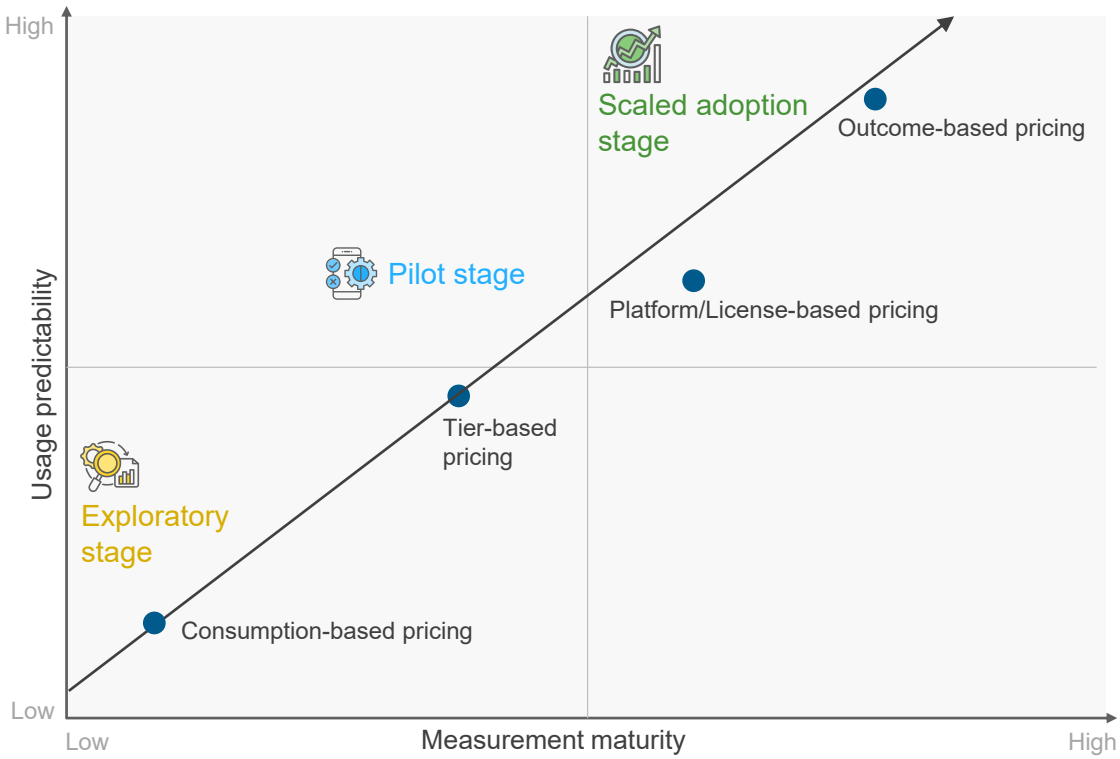

Different pricing models fit best at different stages of the agentic AI journey; the progression below shows how enterprises typically match commercial constructs to their measurement maturity and usage predictability:

- **Exploratory (low measurement × low predictability):** Start with consumption-based pricing to cap risk and learn unit economics
- **Pilot (mid measurement × mid predictability):** Move to tier-based pricing as capability and governance step up (assist → pre-approval → execute)
- **Scaled (mid-high measurement × high predictability):** Consolidate on platform/license models for predictable spend, shared guardrails, and SLAs
- **Mature, well-bounded processes (high measurement × high predictability):** Layer in outcome-based elements (bonus/malus bands) once KPIs, baselines, and attribution are auditable

The way ahead

Roadmap to operationalize agentic AI with outcomes, guardrails, and repeatable value

The next phase for enterprises along their agentic AI journey is about moving from promising pilots to a reliable, repeatable way of working. To unlock enduring value, enterprises should advance through a deliberate adoption sequence. Exhibit 14 depicts a clear road map for agentic AI adoption.

Exhibit 14: Agentic AI adoption road map

Source: Everest Group (2026)

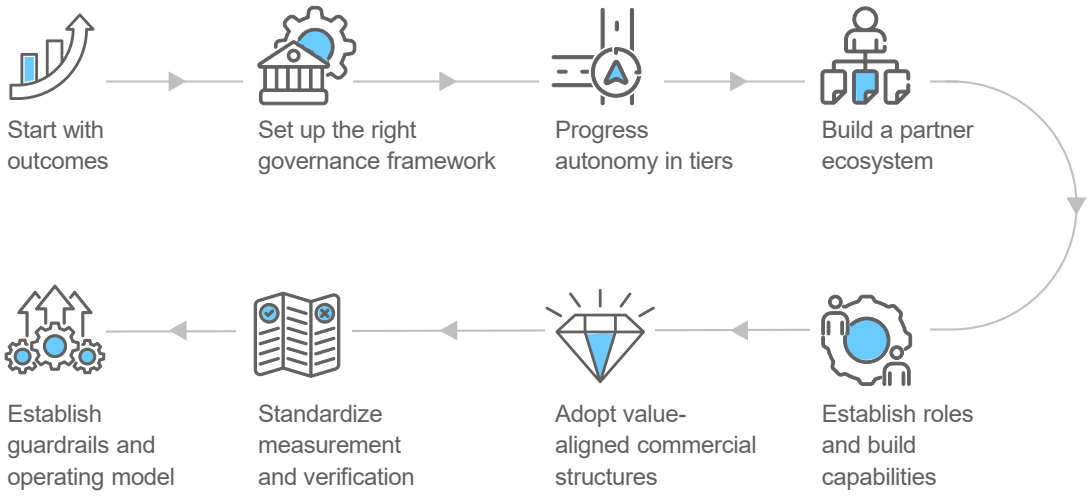

A durable agentic AI operating model rests on a handful of practical disciplines – each reinforcing trust, scale, and measurable value:

- **Start with outcomes:** Prioritize use cases where impact is easy to see and undo, such as SDLC hotspots (monitoring, requirements, testing) and run-heavy domains (ITOps, customer support, finance). Set one or two primary KPIs per case with quality, safety, and compliance guardrails. Run a short shadow period to lock baselines before expanding scope or autonomy
- **Set up the right governance framework:** Move beyond paper policies and embed governance into day-to-day execution through access controls, risk-based approvals, audit-by-default logging, rollback and kill switches, and standardized evaluation. Anchor accountability with CIO/CTO-led ownership and a cross-functional council to keep decisions fast, consistent, and defensible

- **Progress autonomy in tiers:** Treat responsibility as a staircase (assist → pre-approval → task autonomy → process autonomy). Graduate only when KPIs improve and guardrails stay green; revert automatically on breach. Aim broadly for task autonomy and reserve process autonomy for well-bounded, auditable workflows
- **Build a partner ecosystem:** Build what differentiates, partner for domain depth and integration/change, and buy mature components. Design for portability with open interfaces and interchangeable building blocks to avoid lock-in
- **Establish roles and build capabilities:** Upskill in data proficiency, AI governance, strategy, technology fluency, empathy, and compliance. Assign named product, agent, and risk owners with clear decision rights
- **Adopt value-aligned commercial structures:** Pick pricing to match evidence and demand stability – consumption for exploration, tiered as capability/risk step up, platform/license for stable multi-team use, and bounded outcome bands when KPIs are auditable. Add caps/collars, dual-metric triggers, right-to-reprice, and a standing value-realization workstream
- **Standardize measurement and verification:** Publish a KPI dictionary, attribution rules, and data lineage; standardize dashboards for executives and operators. Invite periodic independent validation (internal audit / external assurance) to keep stakeholders confident as autonomy rises
- **Establish guardrails and operating model:** Embed governance into execution through a central control layer that manages permissions, approvals, monitoring, and agent coordination. Define clear ownership, lifecycle standards, and incident response playbooks to prevent agent sprawl and enable a measured progression from assistive support to controlled task execution

The road map can serve as a practical guide to turn pilots into operations, aligning outcomes, guardrails, and economics so value compounds over time.

From manual execution to outcome supervision: the rise of human-AI pods

Rising competitive pressure and the widespread availability of generative AI have reset expectations for execution velocity, with enterprises increasingly assuming faster cycle times and higher automation throughput. As these expectations rise, a range of AI solution offerings promise rapid gains, but enterprises often see performance degrade at scale because production environments demand deep contextualization – domain nuance, process variability, policy constraints, and exception handling. Humans, therefore, remain essential to contextualize workflows, fine-tune models and tools, and calibrate guardrails so agent behavior aligns with business outcomes and risk boundaries.

This need for human context and control is now being codified into new operating models where humans define intent, constraints, and approvals, while agents execute repeatable workflows across functions and systems. As organizations progress from single-agent pilots to multi-agent workflows, orchestration becomes the coordinating

layer that assigns roles, sequences tasks, manages dependencies, and enforces guardrails so agents operate as one coherent operating model.

This evolution culminates in the human-AI pod: a smaller outcome-owned team where humans focus on decisions, risks, and correctness, while specialized agents accelerate delivery under controlled permissions and auditability. Looking ahead, these pods are positioned to become the core execution unit for scaled automation, integrating human accountability with coordinated agent execution to accelerate outcomes while maintaining operating model.

Appendix

Exhibit 15 provides an overview of the 206 respondents who form the basis of the study. It breaks them down by revenue, industry, and designation.

Exhibit 15: Percentages of survey respondents by revenue, industry, and designation

Source: Everest Group (2026)

100% = 206

Split of respondents by **revenue** (100% = 206)	
Small enterprises (< US$500 million)	24%
Midsize enterprises (US$500 million-2 billion)	55%
Large enterprises (> US$2 billion)	21%

Split of respondents by **industry** (100% = 206)	
BFSI	21%
Healthcare and life sciences	21%
Internet	18%
Technology	22%
Telecom	18%

Split of respondents by **designation** (100% = 206)	
C-level executive (CIO, CEO, CTO, CMO, CFO)	20%
Senior Director	16%
Director	28%
President	06%
Senior Vice President	12%
Vice President	18%

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 R Systems

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):

Manukrishnan SR, Practice Director
manukrishnan.sr@everestgrp.com

Nandita Pandey, Senior Analyst
nandita.pandey@everestgrp.com

Sneha Singh, Senior Analyst
sneha.singh@everestgrp.com

Notice and Disclaimers

THANK YOU!

R Systems is a leading digital product engineering company that designs and builds platforms, and digital experiences empowering clients across various industries to overcome digital next-gen products, barriers, put their customers first, and achieve higher revenues as well as operational efficiency. We constantly innovate and bring fresh perspectives to harness the power of the latest technologies like cloud, automation, AI, ML, analytics, Mixed Reality etc.

Contact Us

For more information about our solutions or to discuss how we can help your business, please contact us at:

marketing@rsystems.com