HFS

In partnership with

genpact

MARKET IMPACT REPORT

Autonomy requires trust in AI

Four leadership decisions that determine whether AI scales

Authors:

Dana Daher, Executive Research Leader
Saurabh Gupta, President, Research and Advisory Services
Hridika Biswas, Associate Practice Leader
Niti Jhunjhunwala, Senior Analyst

Foreword

Every major technology shift produces the same pattern: early enthusiasm, uneven adoption, and then a moment of reckoning where the gap between ambition and execution becomes impossible to ignore.

We're at that moment with agentic AI. The enthusiasm is real, but so is the uncertainty. These are systems that both act on and assist human judgment. This shift demands technical investment and organizational honesty about process maturity, governance, and whether the structures we've built for human work are fit for a world where AI executes it.

This research underscores the urgency of this shift. While 92% of executives believe agentic AI will fundamentally change how work is executed, nearly 80% still operate these systems in supervised or assisted modes. This gap between belief and readiness reflects a broader challenge: trust. Enterprises recognize the potential of agentic AI but hesitate to relinquish control.

At Genpact, we've embedded agentic AI into our own operations and demonstrated how autonomy can drive measurable outcomes. For our clients, we take the same approach: integrating AI into workflows to ensure execution-driven transformation. I hope this report is a useful guide as you navigate your own journey.

Ajay Vasal
Global Lead for AI &
Agentic Services

The hardest part of any technology transition is not the technology. It is the moment an organization has to decide how much authority it will hand over. That moment has arrived for agentic AI, and most enterprises are not ready.

HFS Research has spent two years tracking the shift from generative to agentic AI. The pattern is consistent and concerning. Ninety-two percent (92%) of executives we surveyed believe agentic AI will fundamentally change how work is executed, yet nearly 80% still run these systems under close supervision, with humans retaining sign-off on most consequential decisions. The ambition is enterprise grade; the authorization structures are not.

This is not a technology readiness problem. The models work. What remains unresolved is whether organizations have redesigned the accountability, measurement, workforce, and process logic that autonomous execution requires. Most have not and wonder why autonomy stalls at the edges.

This report, produced with Genpact, makes four constraints visible: accountability, measurement, workforce clarity, and process design. Strip out any one, and agentic AI fragments rather than scales.

The 6% pioneering agentic AI enterprise-wide are not moving fastest. They settled ownership, redesigned workflows, and built governance into the operating model before extending autonomy. That is the model worth studying.

Phil Fersht
CEO and Chief Analyst

Contents

Executive summary

The hard part of agentic AI was never going to be the technology alone. Enterprises have already proven that AI can make employees faster, processes leaner, and decisions more informed. The models work. The pilots deliver. What remains unresolved is whether organizations are ready to allow AI systems to execute work autonomously and whether their people, roles, and oversight structures are ready to evolve with them.

For enterprise leaders, this marks a different kind of transformation challenge. Agentic AI changes the nature of enterprise responsibility. Instead of simply generating outputs that humans review, agentic systems can coordinate tasks, trigger actions, and make operational decisions across workflows. This shifts the bottleneck from technical capability to organizational readiness.

To understand how prepared enterprises are for this shift, HFS Research, in partnership with Genpact, surveyed 545 senior executives across 11 industries and conducted interviews with leaders from Fortune 2000 companies. The results reveal a market aligned on direction but divided on readiness. Ninety-two percent (92%) of executives believe that agentic AI will fundamentally change how work is executed. Yet nearly 80% still operate these systems in assisted or supervised modes, with humans retaining approval authority over key decisions.

Our research shows that four organizational decisions determine whether agentic AI becomes real enterprise-scale autonomous execution or remains tightly supervised experimentation.

1

Trust and accountability determine whether agents are allowed to act

Agentic systems are advancing faster than enterprise confidence in them. Only 22% of organizations are comfortable using agents to operate with domain-level or broad autonomy. The barrier is not technical capability but unresolved questions around accountability, explainability, and regulatory exposure.

2

Measurement determines whether autonomy receives sustained investment

Enterprises expect agentic AI to deliver rapid returns, with spending projected to increase 38% over the year. Yet 67% still rely on productivity metrics designed for earlier automation waves. Without agent-native metrics, leaders struggle to prove value and justify scaling autonomous systems, while employees are left without a clear definition of success

3

Workforce clarity determines whether people accept the shift

Resistance to agentic AI is less about culture than about uncertainty. Employees respond positively when organizations clearly define decision rights, oversight responsibilities, and intervention points. As 44% of enterprises expect flatter organizational structures, role clarity becomes essential

4

Process design determines whether autonomy scales

Agentic AI compounds value only when workflows are redesigned end to end. Thirty-three percent (33%) of enterprises cite unprepared business processes as their top barrier to adoption. Automating inside broken workflows produces brittle autonomy rather than sustainable scale.

Enterprises often frame agentic AI as a technology transformation. In practice, the organizations that succeed treat it as an operating model redesign.

The leaders who scale agentic systems first will not necessarily be the ones moving fastest. They will be the ones resolving these four constraints that determine whether autonomous execution is possible: accountability, measurement, people, and process.

Enterprises proved GenAI works; now they must prove agentic AI can execute

GenAI returns were real, but incomplete

To understand the rise of agentic AI, we first need to understand what GenAI changed and what it did not. Two years ago, HFS and Genpact's two-year GenAI countdown study captured a wave of enterprise ambition around GenAI. Looking back now, many of those early gains materialized quickly. Productivity improvements were the most widely realized outcome, achieved by 46% of enterprises, followed by improvements in customer engagement and satisfaction at 36% (see Exhibit 1).

What stands out in the data is where the impact plateaued. The outcomes leaders associated with transformation, including new business models, faster innovation, and sustained competitive advantage, were much harder to achieve. While GenAI has improved how work is performed, it has not fundamentally changed how enterprises execute it. Tasks accelerated inside existing workflows, but the structures governing decisions, accountability, and coordination remained largely unchanged.

Exhibit 1: GenAI benefits show up in efficiency first; growth outcomes lag behind

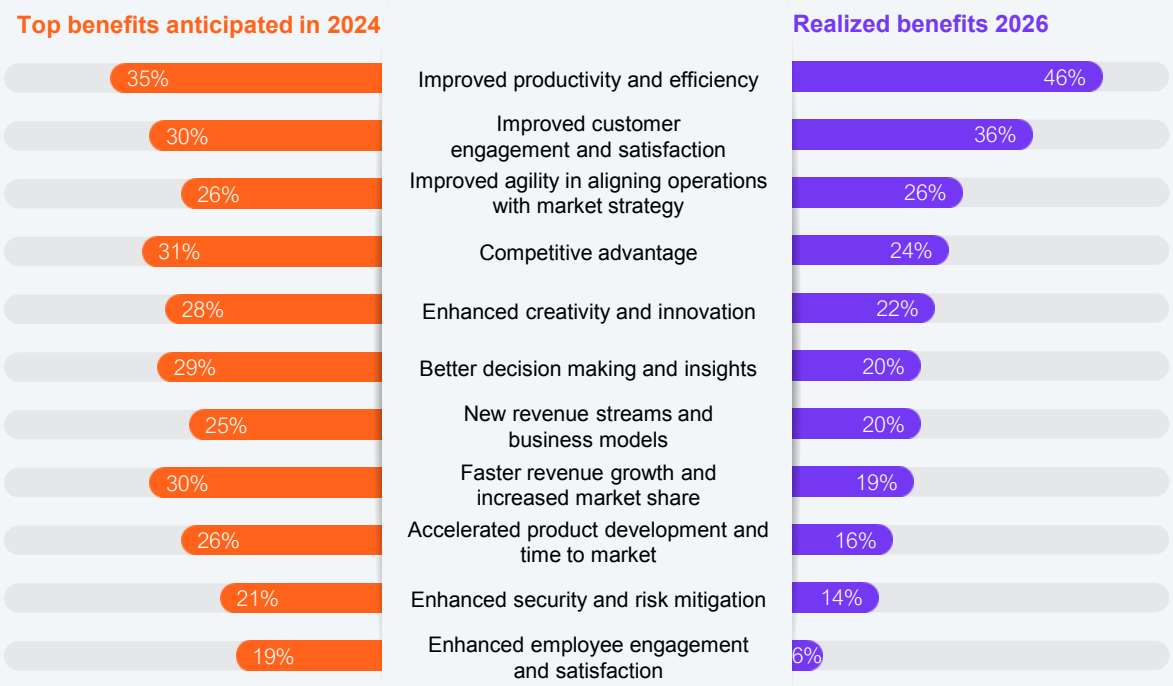

Sample: 550 senior executives (2024) and 545 senior executives (2026) from Fortune 2000 firms
Source: HFS Research, 2026

Agentic AI is expected to close the execution gap

The expectations now placed on agentic AI reflect precisely what GenAI could not deliver. Rather than simply assisting employees, agentic systems are designed to coordinate work across applications, trigger actions across processes, and make decisions without continuous human direction. The ambition is no longer incremental productivity but a shift in how work itself is executed.

Enterprise leaders recognize the significance of that shift. Ninety-two percent (92%) of executives believe agentic AI will fundamentally change how work is executed across their organizations (see Exhibit 2). The expectation is that AI systems will move beyond generating insights to managing workflows, resolving exceptions, and coordinating activity across functions, freeing humans for higher-value judgment and oversight.

As a VP at a Fortune 2000 financial services firm told us, “We are not looking at agents as another layer of productivity tooling. We are looking at them as the next operating model for how work gets done.”

Exhibit 2: Most executives expect agentic AI to change how work gets executed

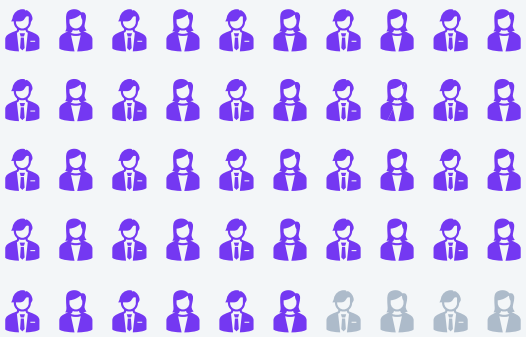

92%

of executives believe agentic AI will fundamentally change how work is executed

Sample: 545 senior executives from Fortune 2000 firms
Source: HFS Research, 2026

The next phase is expected to be faster than the last

The pace at which organizations expect this transition to happen is notable. On average, enterprises expect to scale agentic AI across the business within a 17-month horizon, which means moving from pilots to production where agents are embedded in live workflows and delivering measurable impact on execution, outcomes, and process performance. Thirteen percent (13%) report that agentic AI is already integrated into their organizations, and another 29% expect to reach scale within a year (see Exhibit 3).

These timelines show that enterprises are not treating agentic AI as a distant transformation. They see it as the next stage of AI adoption after their GenAI investments. Yet the organizational conditions required for autonomous systems to operate safely at scale are still emerging. In practice, most enterprises continue to run agentic systems under close supervision, with humans retaining authority over approval for most actions.

That is the contradiction now shaping the agentic AI market. Enterprises expect AI systems to take on more execution, but most have not yet resolved what it means to let them act. The next challenge is not a technical possibility. It is organizational permission.

Exhibit 3: Enterprises expect to scale agentic AI quickly, even as readiness lags

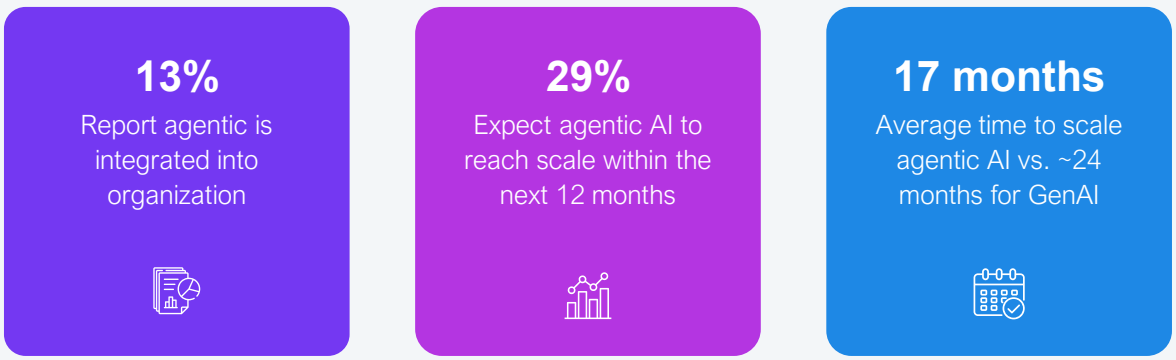

Sample: 545 senior executives from Fortune 2000 firms
Source: HFS Research, 2026

Trust and accountability determine whether agents are allowed to act

Enterprises are already proving that agentic systems can summarize, route, prioritize, draft, and coordinate. What they are unable to prove at scale is that those systems can be trusted to act independently within real workflows. The issue is no longer whether agents can generate value, but whether organizations are prepared to let them own outcomes.

That is why trust matters so much here. Agentic AI asks them to trust not only what the system recommends, but also what it decides and does. That is a different threshold entirely.

As a VP of enterprise automation at a global financial services firm put it, “The opportunity is obvious. The hesitation starts the moment the system is allowed to do something that can’t be quietly reversed.” That is where most enterprises still stop.

Most enterprises are deploying assistive agents rather than autonomous ones

Many enterprises are implementing agentic systems, yet most deployments remain concentrated on recommendation, support, and task assistance rather than autonomous execution. Copilots, task bots, and recommendation agents account for most current activity, while action agents, autonomous coordination, and multi-agent systems remain far less common (see Exhibit 4). This pattern shows where enterprise confidence currently sits. Organizations are far more hesitant to allow autonomous execution, even when they are comfortable with autonomous decision support.

Exhibit 4: Most deployments stay assistive; few enterprises push into autonomy

What is the most advanced form of agentic AI your organization has deployed or is developing?

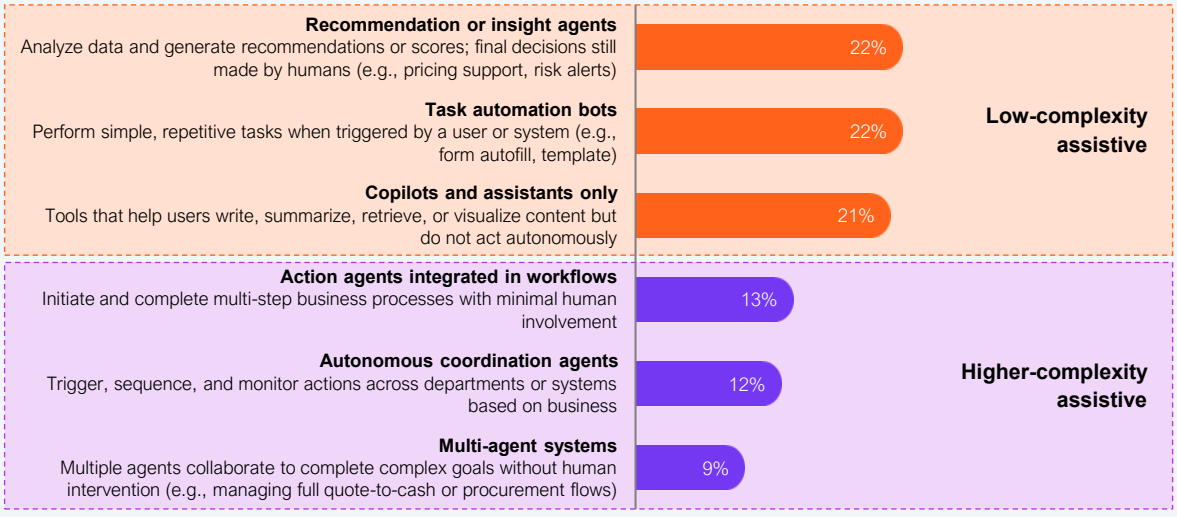

Sample: 545 senior executives from Fortune 2000 firms
Source: HFS Research, 2026

As an AI leader in a large retail bank noted, “We keep saying we want agents, but what we actually approve are better assistants.” The gap between ambition and authorization reflects a trust boundary that many enterprises are still navigating.

Most enterprises operate agents under tightly managed oversight; those moving aren’t making a single leap to full autonomy

The trust boundary becomes even clearer when examining how enterprises authorize agent behavior. Most continue to operate agents in assisted or supervised modes, with humans retaining approval authority over key actions (see Exhibit 5). This can make autonomy seem binary when, in reality, it is conditional. Enterprises grant autonomy incrementally. They adjust permissions based on context and risk, not through a single flip from supervision to full independence.

This creates the central trust gap in the market. Enterprises increasingly believe that agentic AI will transform execution, yet most still govern these systems as if they were tightly bound automation tools. Human checkpoints remain embedded across workflows. Agents may analyze, draft, and recommend, but when responsibility becomes material, the decision, the action, or both still return to a person.

A senior executive at a global bank summarized the issue directly: “We do not have a technology problem. We have a risk appetite problem.” The ambition for autonomy is real, but the organizational confidence to authorize it remains uneven.

Exhibit 5: Most enterprises are not comfortable authorizing high autonomy

Percentage of enterprises by highest autonomy level they are comfortable operating

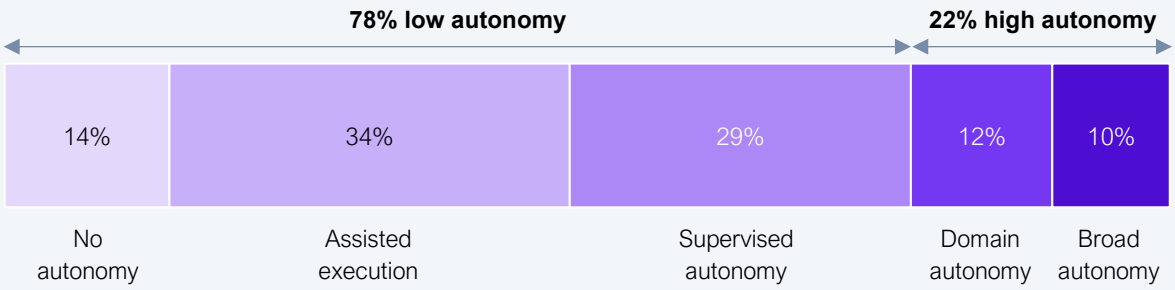

Sample: 545 senior executives from Fortune 2000 firms
Source: HFS Research, 2026

The trust gap is really an accountability gap

As enterprises move deeper into agentic deployments, the conversation shifts. The debate becomes less about whether models perform well and more about who owns the consequences, both operationally and legally, and whether those decisions can be mapped to existing regulatory and liability frameworks when autonomous systems act.

This is reflected in the concerns executives often raise around agentic AI. Compliance exposure, reputational risk, lack of explainability, and unclear accountability are among the top barriers to expanding autonomy, especially when decisions can't be easily mapped to existing legal or regulatory frameworks (see Exhibit 6).

The shift from the GenAI era to the agentic one becomes clear here. With GenAI, the primary question was whether the system produced useful outputs. With agentic AI, it is whether the organization can stand behind what happens when the system takes action. A healthcare executive captured this insight clearly: "The challenge is not getting the model to respond, but being able to stand behind what it just did."

Enterprises moving first are designing for governable autonomy

The organizations moving ahead are not waiting for trust to emerge on its own. They are designing systems, so autonomy remains governable. Instead of relying on humans to approve every action, these enterprises define where agents can act independently, where they must escalate, and what evidence must be captured when they do. Controls are built into the operating model itself rather than layered on through manual approvals.

In these environments, autonomy is not treated as a leap of faith. It is treated as something observable, controllable, containable, or reversible when needed, and defensible under pressure.

Exhibit 6: Trust breaks on accountability, explainability, and regulatory exposure

What concerns are making it hard to trust agentic systems with autonomous action?

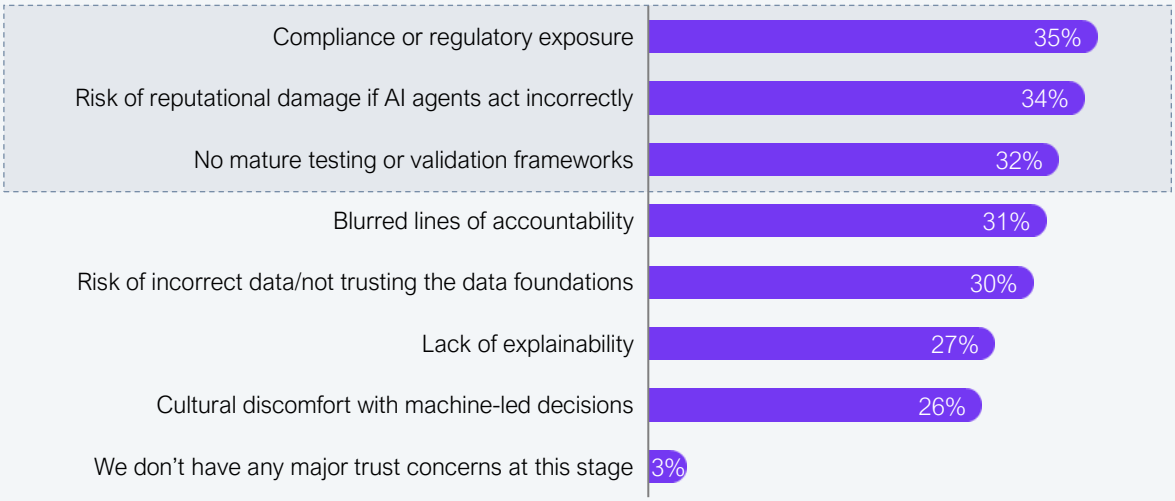

Sample: 545 senior executives from Fortune 2000 firms
Source: HFS Research, 2026

Measurement determines whether autonomy receives sustained investment

Executive expectations for agentic AI are unusually high. Seventy-one percent (71%) of senior executives believe it will deliver ROI faster than any previous wave of technology, including cloud and robotic process automation (RPA), and the first generation of enterprise AI.

What makes this striking is that this confidence exists alongside weak measurement readiness. More than half of enterprises report they don't yet have KPIs to reflect the performance of autonomous systems accurately, and 67% say they still rely on productivity-based metrics that can't capture the value of adaptive, decision-driven systems (see Exhibit 7).

This creates an immediate tension. Leaders expect agentic AI to transform execution, but most organizations still lack the measurement framework required to demonstrate whether it is actually delivering that impact.

Exhibit 7: Leaders expect fast ROI, but most lack agent-ready measurement

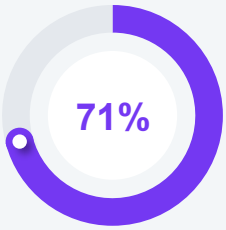

Expect faster ROI

Believe agentic AI will deliver ROI faster than any previous wave of technology

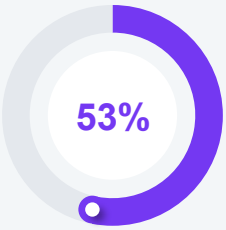

Lack accurate KPIs

Don't have KPIs in place to reflect the performance of autonomous systems

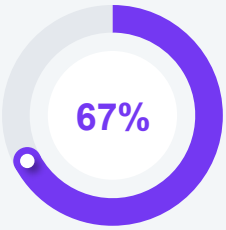

Using wrong metrics

Still depend on productivity-based metrics that cannot capture the value of adaptive, decision-driven systems

Sample: 545 senior executives from Fortune 2000 firms
Source: HFS Research, 2026

Investment is accelerating faster than the ability to account for it

That confidence is already shaping capital allocation. Agentic initiatives account for an average of 7.2% of enterprise technology budgets today, with planned spending increases of 38% over the next twelve months (see Exhibit 8).

Enterprises are not hedging on agentic AI. Budgets are moving, timelines are compressing, and leadership teams are increasingly positioning autonomous systems as part of the next operating model. But commitment and accountability are different things. Investment is accelerating faster than the infrastructure needed to evaluate it.

If organizations can't measure what these systems are doing, they will struggle to defend the investment, identify which deployments are working, or determine when autonomy should expand.

Exhibit 8: Agentic budgets are rising faster than accountability for value

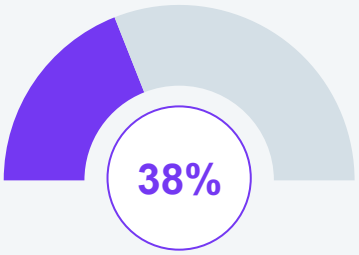

Planned spending increase

Average increase in agentic AI investment over the next 12 months

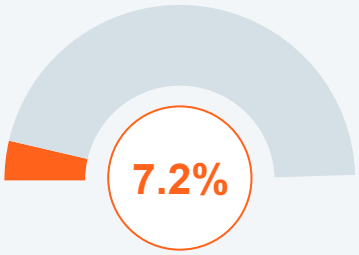

Of technology budgets

Average share of technology budget currently allocated to agentic initiatives

Sample: 545 senior executives from Fortune 2000 firms
Source: HFS Research, 2026

Most enterprises are still measuring autonomy through a productivity lens

The measurement gap becomes clearer when examining how enterprises currently evaluate agentic systems. Cost savings, productivity gains, and cycle-time reduction dominate the metrics leaders are using today (see Exhibit 9).

Those measures made sense for earlier automation waves such as RPA and assistive AI tools. They are far less suited to systems designed to coordinate work across processes, resolve exceptions, and execute decisions with limited human intervention.

Only 10% of enterprises report using agent-native measures that reflect the outcomes autonomous systems are intended to produce. These include resolving decisions without escalation, completing workflows end to end without human intervention, and handling operational exceptions autonomously.

As a technology executive at a healthcare enterprise explained, productivity metrics show whether employees are working faster. Agentic systems force organizations to ask a different question: whether the system itself is beginning to carry part of the work.

Enterprises that scale autonomy will redefine what success looks like

Organizations that move ahead will not simply invest more aggressively. They will redefine how success is measured.

Agentic systems require metrics that capture execution rather than activity, including autonomous workflow completion, reduced escalation, and independent exception handling. Without that shift, enterprises risk judging agentic AI through a productivity lens designed for assistive tools.

And once organizations begin measuring execution rather than activity, another constraint becomes visible. As systems take on more work, enterprises must also redefine how human roles, oversight responsibilities, and workforce structures evolve alongside them.

Exhibit 9: Enterprises still measure agentic AI like automation, not execution

Which metrics does your organization prioritize when defining agentic AI ROI?

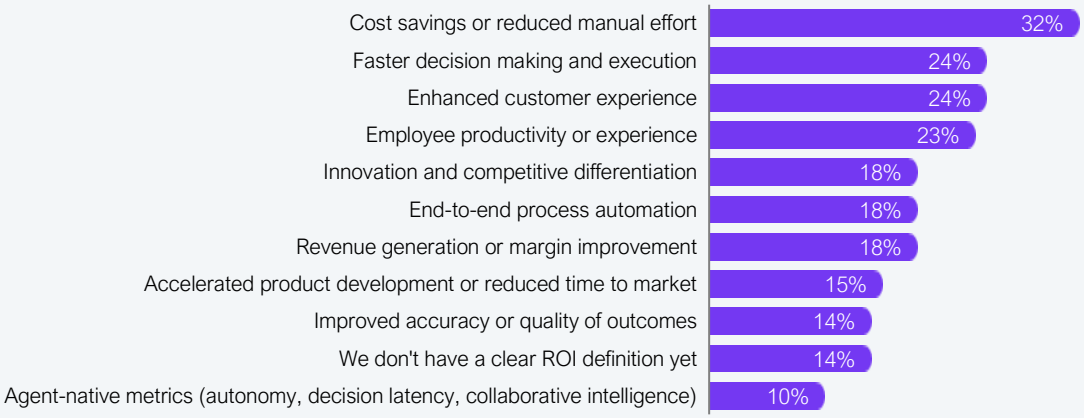

Sample: 545 senior executives from Fortune 2000 firms
Source: HFS Research, 2026

Enterprises are restructuring for agentic AI before their people can govern it

Agentic AI is not only changing what work gets done. It is also changing what the organization expects people to oversee, authorize, and own. Structural redesign is already underway, but the workforce expected to operate within that redesign is still emerging. The result is a widening gap between how enterprises intend to run work and whether their people are equipped and clear about their roles to govern it.

Enterprises are redesigning organizations faster than they are redefining the human role

The structural changes are already visible. Forty-four percent (44%) of enterprises expect agentic AI to flatten traditional hierarchies by reducing management layers, while 36% expect to eliminate specific roles entirely rather than augment them (see Exhibit 10). These are not incremental adjustments. They reflect growing expectations that agentic systems will absorb coordination and oversight tasks historically performed by middle management. This shifts not only work but also decision authority, raising new governance, ethical, and leadership questions.

The challenge is that these structural bets imply different futures for the workforce. Flattening hierarchies assumes humans shift toward higher-order judgment, exception handling, and oversight. Eliminating roles assumes the loop no longer requires human involvement. Many enterprises seem to be pursuing both simultaneously, often without a clear framework for which roles evolve, which ones disappear, and which new responsibilities emerge.

As an operations leader at a global telecommunications firm explained, “We are reorganizing around speed and autonomy, but the human role in that model is still being invented.” In many organizations, structural redesign is moving faster than role design.

Exhibit 10: Agentic AI is expected to flatten organizational structures and remove roles

Percentage of enterprises expecting each structural change

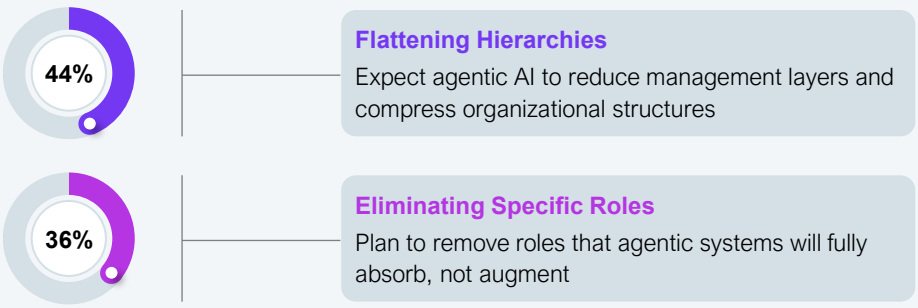

Sample: 545 senior executives from Fortune 2000 firms
Source: HFS Research, 2026

The skills gap is shifting from building AI to operating alongside it

The skills enterprises now prioritize reinforce this shift. Workflow orchestration and integration top the list at 42%, followed by data engineering at 39% and monitoring and observability at 36% (see Exhibit 11). These are not capabilities required to build AI models. They are the capabilities required once systems are embedded in workflows and must be connected, monitored, and governed in real time.

Exhibit 11: Scaling agentic AI shifts demand to orchestration, data, and observability skills

What are the top three technical skills you believe are most essential to drive successful agentic AI initiatives?

Sample: 545 senior executives from Fortune 2000 firms
Source: HFS Research, 2026

The business skills story reinforces the same shift. Enterprises increasingly prioritize business and domain knowledge combined with technical fluency, followed by critical thinking, strategic planning, and accountability and judgment (see Exhibit 12). The emphasis is moving away from simply using AI tools toward supervising how autonomous systems operate.

What really stands out in the skills data is the drop in understanding ethical AI, from 41% to 21% in two years, the steepest decline on the list. This suggests that many organizations believe that governance mechanisms and guardrails are already in place. Earlier findings on trust indicate that most enterprises are still building them.

Exhibit 12: Enterprises prioritize domain fluency and judgment over ethics know-how

What are the top three business or domain skills most critical to agentic AI success in your organization? Also, what are the top three business skills you believe were most crucial to drive the success of your GenAI initiatives in 2024?

Sample: 545 senior executives from Fortune 2000 firms
Source: HFS Research, 2026

Workforce sentiment reflects whether accountability has been explained

Employee sentiment toward agentic AI is neither overwhelmingly positive nor overwhelmingly negative. Instead, it is fragmented in ways that reveal how unevenly enterprises are managing the transition. Only 16% of organizations report very positive sentiment toward agentic AI, while 60% feel neutral or negative (see Exhibit 13).

Where sentiment becomes negative, the drivers are consistent. Employees cite concerns about job security, lack of clarity around how AI decisions will be governed, and discomfort with the pace of change. These reactions reflect uncertainty about responsibility rather than skepticism about the technology itself.

Organizations reporting stronger workforce acceptance emphasize transparency around how agents operate, structured training programs, and clearer communication about how human roles will evolve. Confidence increases when employees understand the boundaries between human accountability and system autonomy.

Enterprises that move ahead will redesign roles as deliberately as they redesign systems

The workforce challenge is not simply one of training. It is one of the roles of design, accountability, and organizational clarity. Enterprises can't scale agentic systems if the people expected to oversee them don't understand where their responsibility begins and ends, and how their work changes alongside autonomous execution.

As agentic systems assume more responsibilities, the workforce challenge becomes inseparable from a process challenge. Enterprises must not only redefine human oversight, but also redesign the workflows those systems are expected to govern.

Exhibit 13: Workforce sentiment turns negative when governance and role clarity are unclear

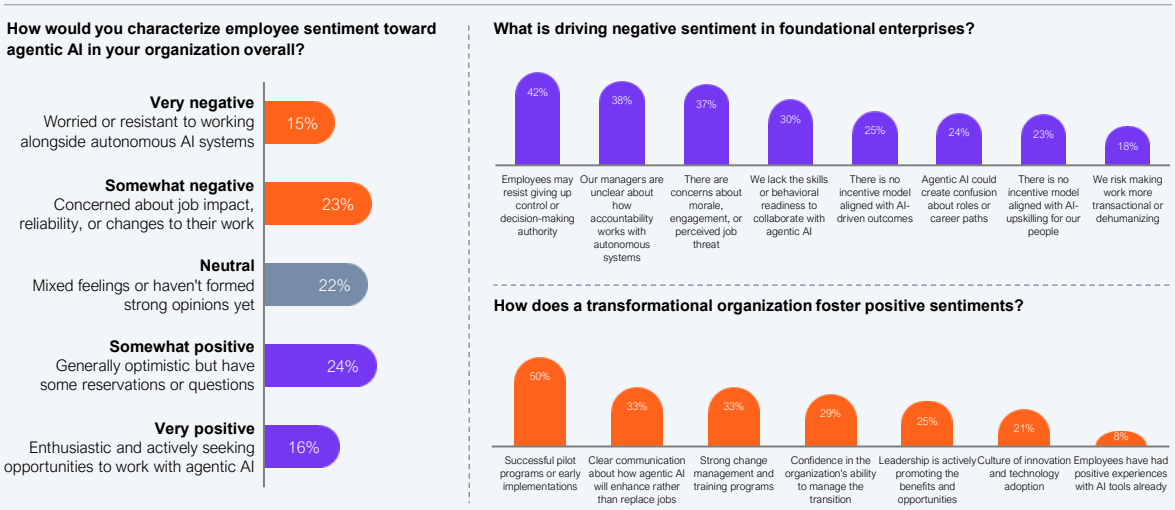

Sample: 545 senior executives from Fortune 2000 firms
Source: HFS Research, 2026

Agents don't fix broken processes; they expose them

Most enterprises are not automating work. They are automating the symptoms of broken work. The process logic underneath remains intact, which is why autonomy often becomes brittle, ROI remains difficult to demonstrate, and accountability stays unclear.

Workflow redesign is emerging where process failure is most visible

Agentic AI is no longer confined to isolated pilots. Workflow redesign is underway in customer service, IT infrastructure, and finance, functions where decision volume is high, exceptions are frequent, and the cost of manual coordination is most visible. Nearly half of enterprises report reworking customer service workflows so agents can manage decision chains rather than individual tasks. IT and finance follow closely behind (see Exhibit 14).

The pattern is not accidental. These functions sit at the intersection of data, decisions, and coordination across organizational boundaries. They are also where traditional process logic breaks first when systems attempt to act autonomously. In many cases, redesign is not a planned transformation. It is the process forcing itself to change because the existing structure can't support autonomous execution.

Exhibit 14: Workflow redesign is concentrating on where decisions and exceptions are the highest

In which business functions has agentic AI triggered or is expected to trigger workflow redesign?

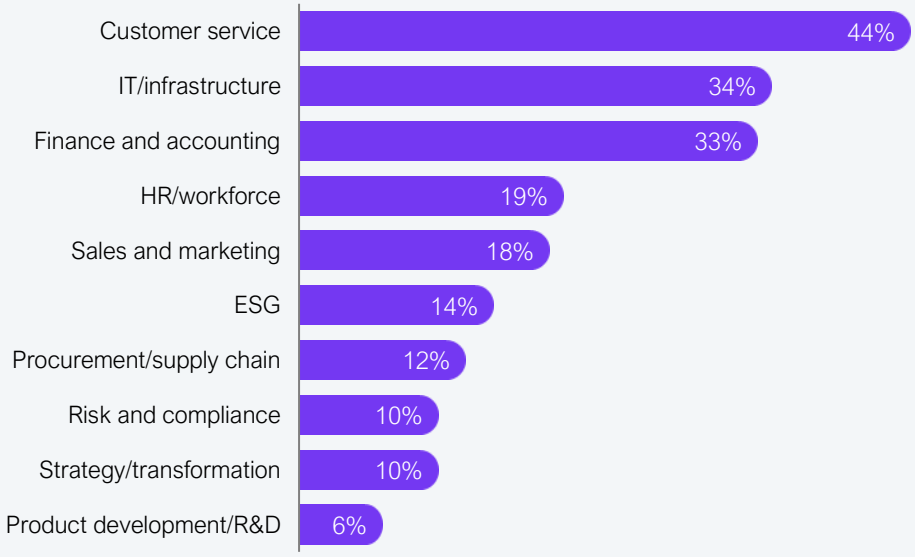

Sample: 545 senior executives from Fortune 2000 firms
Source: HFS Research, 2026

Process readiness, not ambition, determines whether autonomy scales

Process readiness has emerged as the top barrier to agentic AI adoption, ahead of governance concerns, expertise gaps, regulatory risk, and budget constraints (see Exhibit 15). This finding connects many of the challenges observed across the market.

Trust is difficult to establish in a process that does not make accountability visible. ROI is difficult to measure when workflows rely on human judgment at every step rather than explicitly surfacing decision logic. Governance is difficult to embed in processes designed for sequential handoffs rather than distributed decision making.

The barriers listed below the top line reinforce the same structural problem. Lack of governance structures, shortage of internal expertise, and difficulty scaling beyond pilots often reflect workflows that were never redesigned for autonomous execution. Enterprises attempting to layer autonomy onto legacy process logic are not scaling agentic AI. They are discovering where the process itself is not ready for it.

Autonomy scales only when processes are redesigned for system execution

Enterprises making progress share a common approach. They stop optimizing individual tasks and begin redesigning workflows end to end. Sequential approvals that assume a human will authorize each step are removed. Decision ownership is clarified so accountability remains visible when agents act across functions. Manual handoffs are replaced with system triggers designed for autonomous execution.

However, as workflows become more compressed, risk can also become more concentrated. Fewer handoffs may mean fewer control points, increasing the blast radius when failures occur. In practice, every workflow left undesigned becomes a ceiling on how far autonomy can expand.

Exhibit 15: Process readiness is the leading blocker to scaling agentic AI

What are the top three organizational barriers hindering your agentic AI adoption?

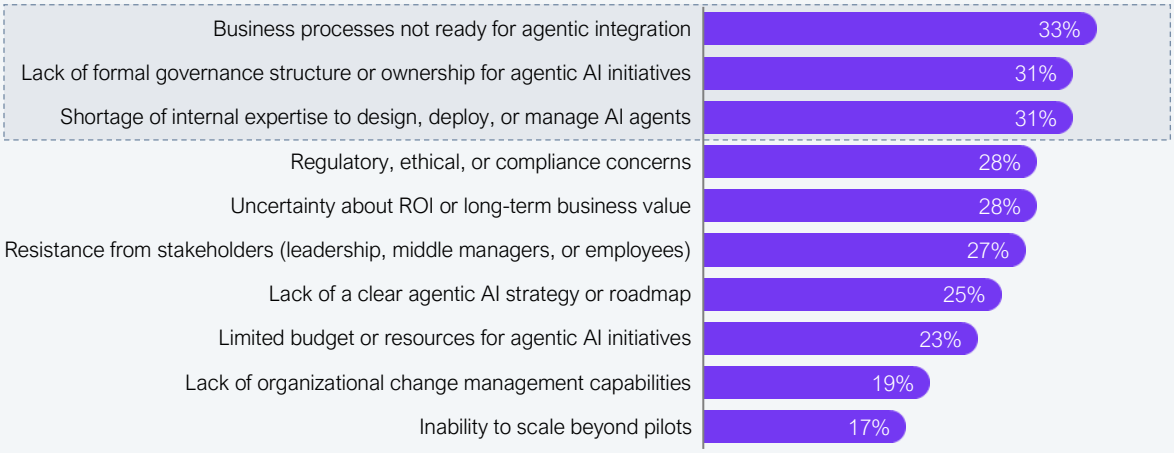

Sample: 545 senior executives from Fortune 2000 firms
Source: HFS Research, 2026

Four actions for leaders

Four decisions separate enterprises compounding value from those running expensive pilots that never scale.

1

Define accountability before expanding autonomy

Before any agent enters production, the organization must answer three questions explicitly: who owns the agent, who is responsible when it fails, and how escalation works. These are not legal questions. They are design decisions.

Enterprises that have not defined these responsibilities are not ready to expand autonomy. They are simply deferring the accountability problem until the consequences become more expensive.

2

Replace the ROI framework before scaling investment

Productivity metrics systematically undervalue what agentic AI delivers. Before the next budget cycle, organizations must define what agent-native success looks like. That includes decisions removed from the human queue, workflows completed end to end without escalation, and organizational capacity created through autonomous execution. It also requires isolating the agent's contribution from surrounding changes such as process redesign, data improvements, or workflow simplification, to avoid misattributing value.

If the measurement framework can't capture what the system is doing, investment decisions will continue to misdirect capital.

3

Make the human transition a design constraint

Workforce anxiety is not simply a change management issue to fix after deployment. It signals that accountability and oversight structures have not been made visible to the people expected to operate within them.

Enterprises that successfully scale agentic systems address that by building role clarity, escalation paths, and oversight responsibilities directly into deployment design rather than relying on communication programs later.

4

Redesign the process before deploying the agent

Every workflow left unchanged becomes a ceiling on how far autonomy can scale. Enterprises must identify sequential approvals, manual handoffs, and unclear ownership structures built for human coordination and redesign them before agents are introduced.

Autonomy compounds only when the process itself is rebuilt for system execution.

The Bottom Line:
Enterprises that will lead in agentic AI are not those moving fastest. They are the ones resolving the four constraints that determine whether autonomy compounds: accountability, measurement, people, and process.

As organizations address those constraints, they are doing more than preparing AI to execute work. They are redefining how decisions are governed, how responsibility is assigned, and where control sits when systems begin to act.

HFS Research authors

Dana Daher
Executive Research Leader

Dana Daher is an Executive Research Leader at HFS Research, spearheading research initiatives in emerging technologies and employee experience. With a unique blend of expertise in anthropology and IT, Dana leads cutting-edge research that shapes industry landscapes across various domains, including employee experience, Agentic AI, generative AI, diversity, equity, and inclusion (DEI), and sustainability.

Saurabh Gupta
President, Research and Advisory Services

Saurabh Gupta is president, Research and Advisory Services for HFS Research. He sets the strategic research focus and agenda for HFS Research, understanding and predicting the needs of the industry and ensuring that HFS maintains its position as the strongest impact thought leader for business operations and services research. Saurabh oversees HFS' global research function, managing the team of analysts and operations across the US, Europe, and Asia.

Hridika Biswas
Associate Practice Leader

Hridika Biswas is an Associate Practice Leader at HFS Research. She joined the team in 2021 and her key coverage areas include business services (F&A, S2P), process intelligence, intelligent document processing, and automation. Her interests include analyzing how emerging tech enables enterprises to reach their transformation goals.

Niti Jhunjhunwala
Senior Analyst

Niti Jhunjhunwala is a senior analyst for HFS Research. Her coverage areas include banking and financial services and GenAI. She also regularly contributes to competitive intelligence across IT and business process services and the HFS Market Index, a quarterly report that analyses the performance and major developments of top service providers over the past quarter.

Genpact contributors

Dr. Pramukh Nanjundaswamy Vasist
Research and Thought Leadership

Dr. Pramukh Nanjundaswamy Vasist drives research and thought leadership at Genpact. A Data and AI practitioner with over 17 years of global leadership experience across consulting and professional services firms, he bridges senior industry practice with rigorous academic scholarship. He holds a Ph.D. in Information Systems and has published extensively in leading international journals and presented at international conferences across information systems and management disciplines. His thought leadership spans AI-enabled enterprises, disinformation, responsible AI, future of data-driven organizations, and the societal impact of digital technologies.

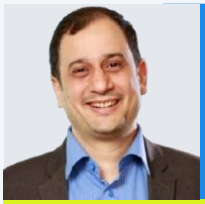

Ojas Save
Technology Strategy Lead

Ojas Save leads strategy for the Data, Tech, and AI business at Genpact, where he focuses on driving market positioning, shaping growth, and scaling next-generation AI-led offerings for enterprise clients. With a background spanning corporate strategy, business planning, marketing, and client engagement across global enterprises and startups, he brings a distinctive perspective on translating data and AI capabilities into business value. His current focus is on defining market narratives, sharpening competitive positioning, and aligning the business to emerging opportunities in an increasingly AI-driven landscape.

Ajay Vasal
Global Lead for AI & Agentic Services

Ajay is Genpact's global lead for AI & Agentic Services. His focus is to help clients make effective use of AI and related technologies in their business operations. He has been in the industry for more than 20 years and has deep experience in helping client shape their AI and data strategies, build new capabilities, and design complex transformation programs. He has worked across a number of industries and brings in tremendous learnings and points of view on the future of AI and how companies should adopt it.

Zainab Andrabi
AI Programs and Strategy Analyst, Office of the CTIO

Zainab drives AI strategy and programs within the Office of the Chief Technology & Innovation Officer at Genpact. She leads thought leadership and research initiatives, including serving as a contributor to Genpact's first primary research report, *Autonomy by Design*. Zainab manages key AI programs for the CTIO, using performance across Genpact's Data, Tech, and AI portfolio to guide organizational direction and strategic priorities. She also supports Genpact's growing agentic advisory practice, helping organizations assess their readiness for agentic AI, define adoption roadmaps, and build the frameworks needed to deploy autonomous systems at scale.

Demographics

What is your role in your organization?

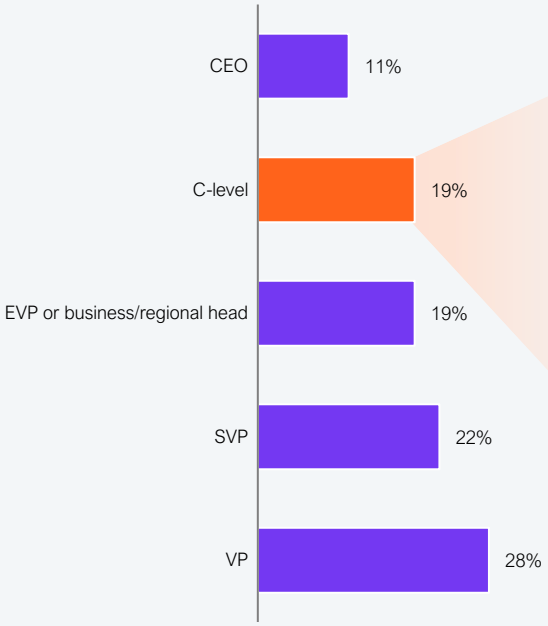

What best describes your specific role?

Which of the following best describes your primary function or department within your organization?

Which of the following best describes your primary function in IT department within your organization?

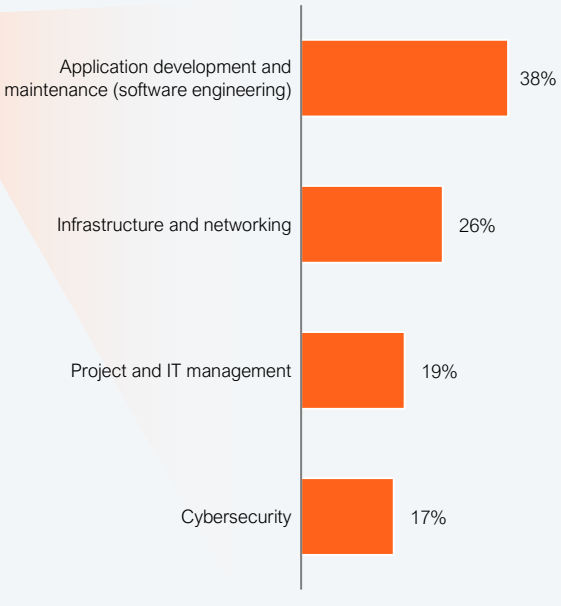

Demographics

What industry does your business operate in?

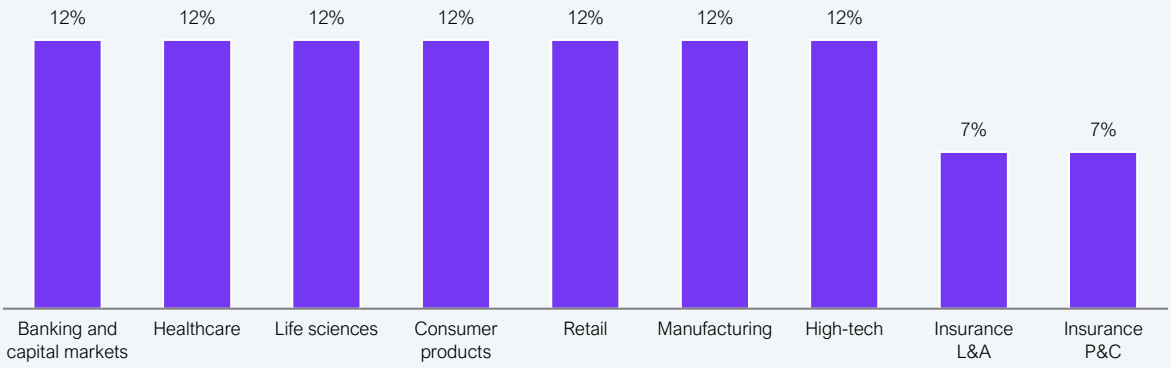

In which country is your organization headquartered?

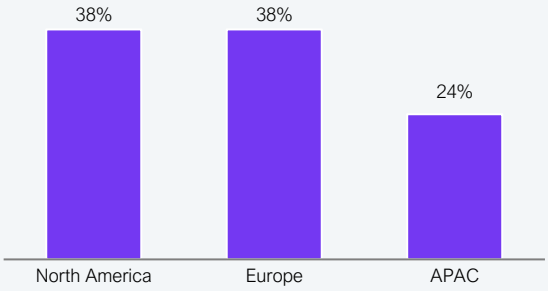

What is your organization's annual revenue range?

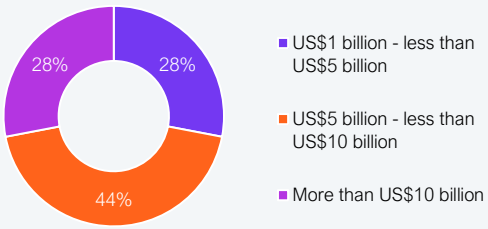

How would you best describe your role in relation to your organization's investments and initiatives around AI?

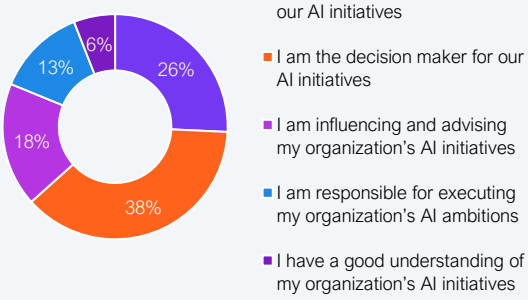

What best describes your organization's current status with agentic AI?

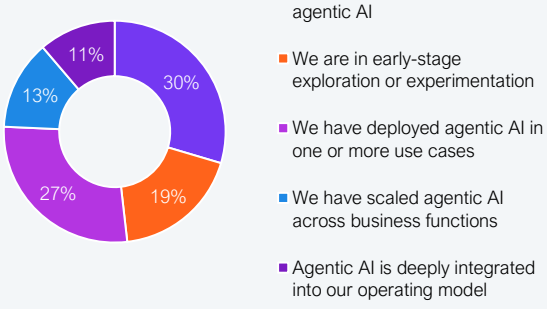

About Genpact

Genpact (NYSE: G) is a global professional services and solutions firm delivering outcomes that shape the future. Our 125,000+ people across 30+ countries are driven by our innate curiosity, entrepreneurial agility, and desire to create lasting value for clients. Powered by our purpose – the relentless pursuit of a world that works better for people – we serve and transform leading enterprises, including the Fortune Global 500, with our deep business and industry knowledge, digital operations services, and expertise in data, technology, and AI.

Powered by our purpose – the relentless pursuit of a world that works better for people – we serve and transform leading enterprises, including the Fortune Global 500, with our deep business and industry knowledge, digital operations services, and expertise in data, technology, and AI.

Get to know us at genpact.com and on LinkedIn, X, YouTube, and Facebook.

HFS

About HFS

- **INNOVATIVE**
- **INTREPID**
- **BOLD**

HFS Research is a leading global research and advisory firm helping Fortune 500 companies through IT and business transformation with bold insights and actionable strategies.

With an unmatched platform to reach, advise, and influence Global 2000 executives, we empower organizations to make decisive technology and service choices. Backed by fearless research and an impartial outside perspective, our insights give you the edge to stay ahead.


www.hfsresearch.com


hfsresearch


www.horsesforsources.com


www.horsesmouthpodcast.com