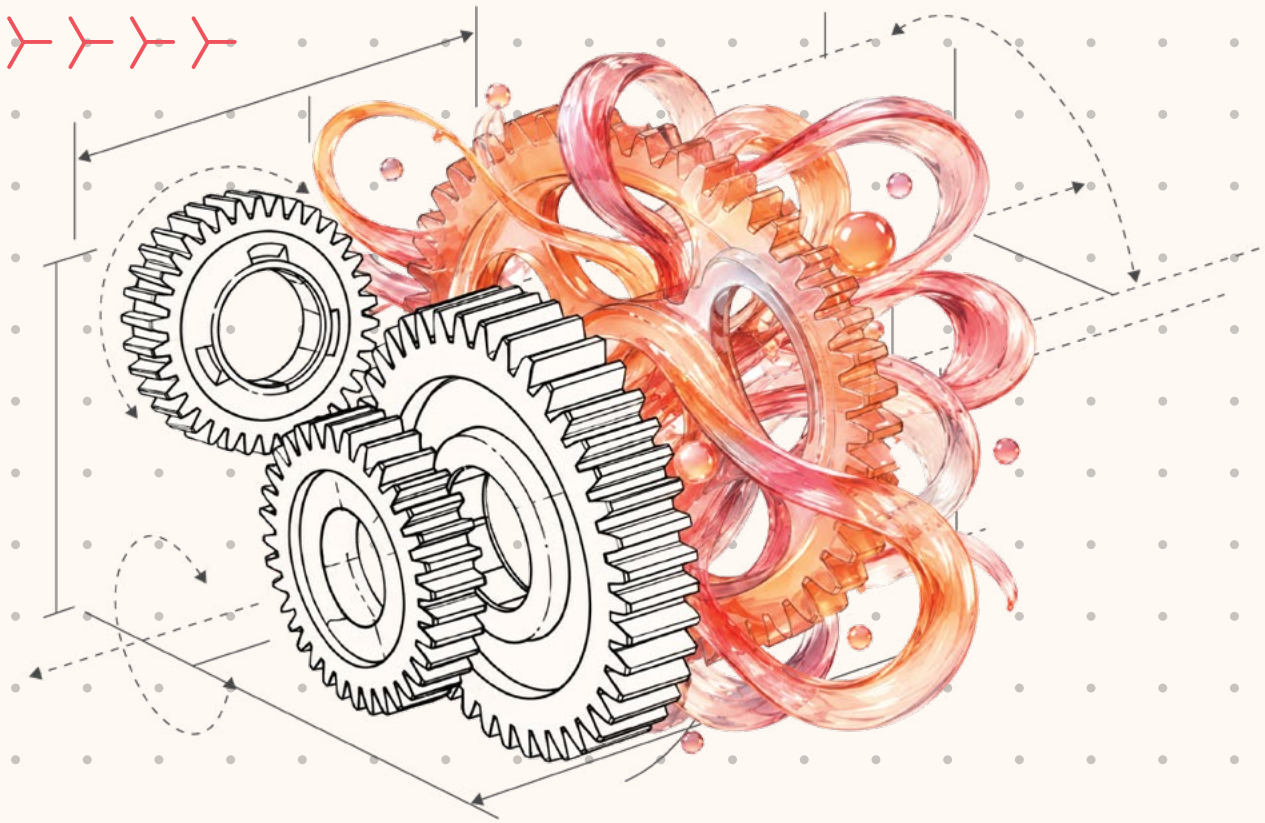
genpact

Autonomy in Practice

How agentic operations can turn AI progress into business value



November 2026

Foreword



Sanjeev Vohra

Chief Technology and Innovation Officer, Genpact

With every wave of technological advancement, people have been the catalyst for change. Agentic AI is no different.

When electricity reached the factory floor, the early gains were modest. Owners swapped steam engines for electric motors and kept everything else the same. The real payoff came years later, when people rethought the factory layout, the flow of work, and the jobs within it. Technology made it possible, but people made it matter.

Many enterprises are at this same moment with AI agents. They have deployed them into operating models designed for human-only work, and much of the value they expected, what we call the autonomy dividend, has yet to arrive. Rethinking the enterprise operating model for people and agents together is what we call agentic operations. This report examines five building blocks critical to making that shift: process reinvention, the human-agent workforce, the data and AI foundation, governance and security, and value capture. Across each, we measure how enterprises are adapting their operating models and what the Leaders do differently.

This is the third report in our thought leadership research series. *Autonomy by Design* set out the vision for the agentic enterprise. *The $18 Trillion Opportunity: How Four Enterprise Debts Will Make or Break Your AI Future*, with HFS Research, showed how process, data, technology, and talent debt trap AI value. This report asks the next question: what separates the enterprises that turn AI into lasting value from the rest?

The gap is stark. Nearly 4 in 5 of the more than 2,000 executives we surveyed report AI agents in production. Fewer than half of their funded initiatives have delivered the expected value. Just 15% have fully tied AI investment to P&L outcomes.

The Leaders behave differently, starting with a CEO-level mandate, and their behaviors are ones any leadership team can learn from. We know because we learned many of them first-hand by redesigning our own operations in Client Zero, our program to transform how Genpact itself works, alongside our work with clients.

As you read on, ask not only what AI can do for your enterprise, but what your people can do with it.

Table of contents

- Foreword 02
- Executive Summary 04
- Enterprises have crossed the AI production threshold 06
- Enterprises are failing to capture the autonomy dividend 12
- Agentic operations is how Leaders capture the autonomy dividend 16
- Make autonomy work in practice 29
- Research methodology 32
- About the authors 36

Executive summary

The enterprise AI conversation has changed in the space of a year. We've moved from pondering whether AI can reach production to whether AI in production delivers value.

In 2025, *Autonomy by Design* explored what it would take for enterprises to move from AI that generates to AI that executes.

In 2026, *The $18 Trillion Opportunity: How Four Enterprise Debts Will Make or Break Your AI Future* with HFS Research showed what kept agentic AI stuck in pilots: broken processes, untrustworthy data, aging systems, and unprepared talent.

For 2027, a new reality is emerging faster than many expected. AI agents are running in production across enterprises. The focus is on achieving autonomy in practice. This does not mean full autonomy. It means deliberately deciding where agents act, where people decide and validate, and what business outcomes the human-agent workforce can produce.

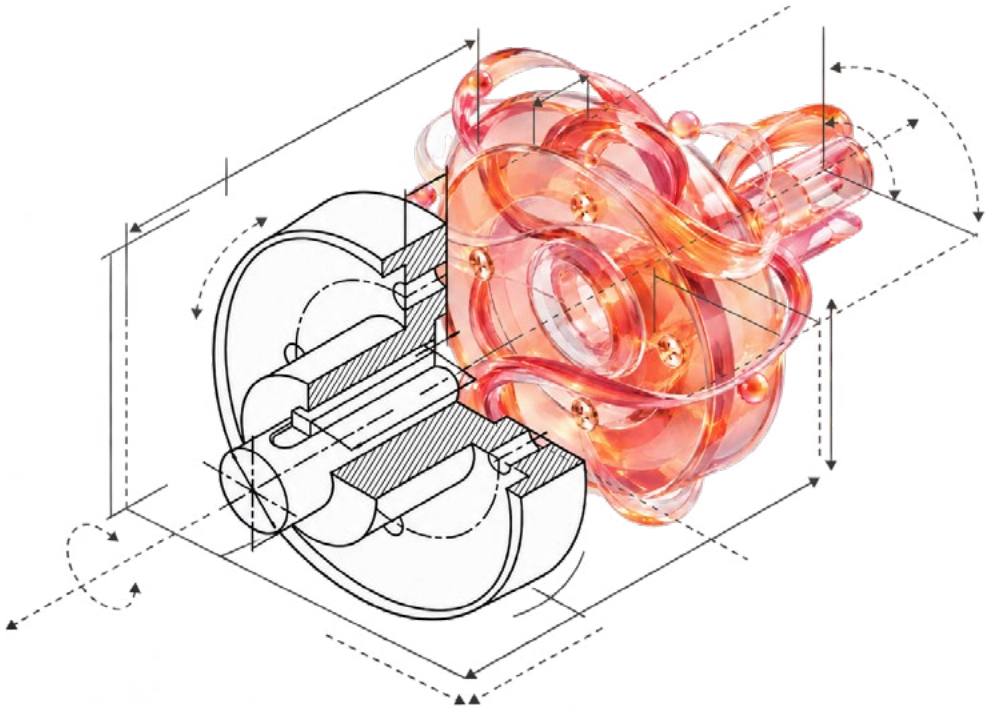
Our survey of 2,005 global executives across 17 industries, supplemented by in-depth interviews, reveals a paradox. Among respondents, 79% report they have agentic AI live in production. Across the eight decision stages we measured, AI now leads or fully runs 40% of decisions. Yet on average, just 49% of AI initiatives funded in the past 24 months have delivered the value expected, while only 15% have fully integrated AI investment to P&L outcomes.

What we found is that the problem is fundamentally operational. AI capability is advancing faster than the building blocks required to capture its value. The story plays out across three critical findings:

1. Enterprises have crossed the AI production threshold

Agents are running in production and are increasingly included in enterprise decisioning. Genpact's annual AI Maturity Index shows the shift: the share of enterprises rated Advanced or Leaders has risen from 27% to 52%, driven by the Advanced tier more than doubling from 15% to 35%. Advanced is the new middle: most enterprises have moved from building foundations to scaling AI. The top tier has moved modestly: Leaders have increased from 12% to 17%. Reaching Advanced takes investment and deployment. Becoming a Leader takes a change in how the enterprise operates and generates business outcomes. AI progress is real. Now enterprises need to convert that progress into measurable value.

AI progress is real. Now enterprises need to convert that progress into measurable value.



2. Enterprises are not capturing the autonomy dividend

AI progress creates the potential for an autonomy dividend: the enterprise value captured when AI-enabled gains translate into better business performance, not simply greater productivity. Our research shows that enterprises are capturing only part of that potential today.

The shortfall shows up in three places.

- **Execution:** on average, only 49% of AI initiatives funded in the past 24 months have delivered the value expected
- **Business ownership:** just 15% of enterprises have fully connected AI investments to specific P&L outcomes
- **Capacity utilization:** agents are scaling alongside people, not simply replacing them. Among enterprises with agents in production, 69% increased their agent population while 92% maintained or increased human headcount. And while 77% say they have a defined approach for redeploying capacity released by AI, 46% identify an inability to redeploy it as the leading barrier to realizing measurable value. Freed time is not freed capacity until someone owns where it goes

Freed time is not freed capacity until someone owns where it goes.

Capturing the autonomy dividend requires the shift from deploying AI to operating with it across the enterprise. What's missing are the building blocks to make the shift.

3. Agentic operations is how Leaders capture the autonomy dividend

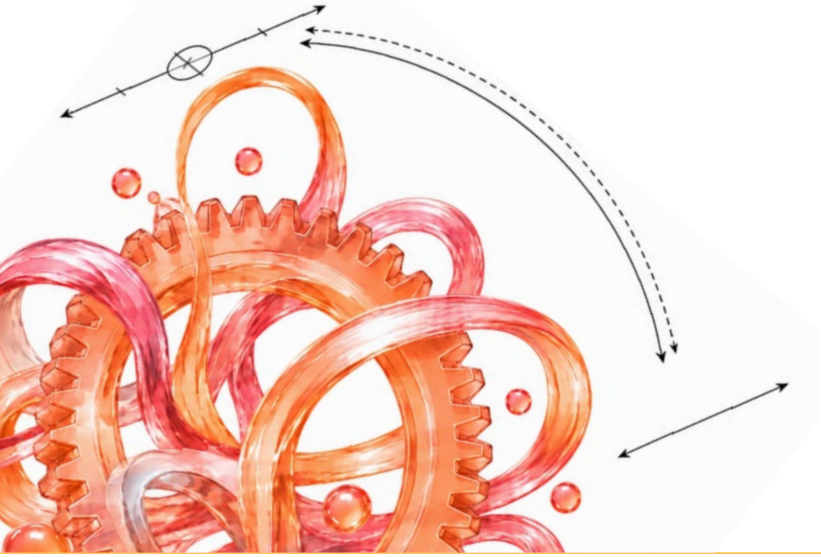
We measured five building blocks critical to agentic operations: process reinvention, a human-agent workforce, a data and AI foundation, governance and security, and value capture. Most enterprises have not connected them: only 9% have completed a significant operating model transformation, 12% have full-stack AI capability at scale, and 17% embed governance across the agent lifecycle.

Leaders are further ahead on each building block, advancing them as a system. Backed by a CEO-level mandate, they:

- Reinvent work across the enterprise
- Build the human-agent workforce deliberately
- Strengthen the data and AI foundation for repeatable scale
- Govern and secure agents throughout their lifecycle
- Make value capture a P&L discipline

Deploying agentic AI creates potential. Agentic operations can help turn that potential into business value. That is autonomy in practice.

Deploying agentic AI creates potential. Agentic operations can help turn that potential into business value. That is autonomy in practice.



Enterprises have crossed the **AI production threshold**

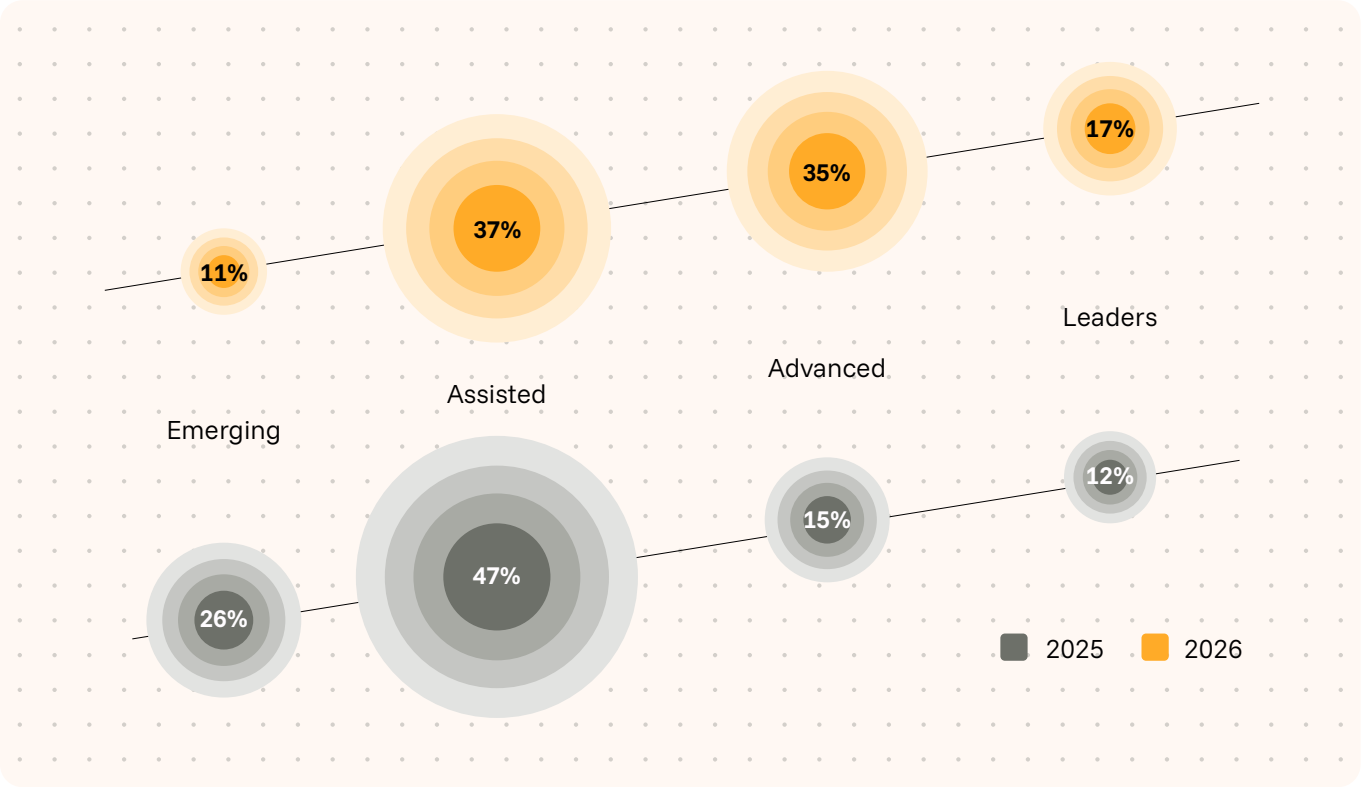
In 2025, we launched the AI Maturity Index to support our research for Autonomy by Design, using it to understand how enterprises were putting the foundations for autonomy in place. That research identified three forces constraining progress: a deficit in realized value, technological complexity, and organizational inertia.

As we approach 2027, we have updated the index to track what has changed as enterprises move from designing for autonomy to running agentic AI in production. The results show significant progress but also reinforce a central premise of our earlier research: deployment alone does not create autonomy or value. Autonomy in practice requires operational change that converts advancing AI capability into measurable business outcomes.

The number of AI Leaders has risen from 12% to 17%. The bigger move was into Advanced, which has more than doubled from 15% to 35% as enterprises have left the Emerging and Assisted tiers. Advanced and Leaders combined have increased from 27% to 52%, while Emerging has decreased by more than half, from 26% to 11% (Figure 1).

What this means: Advanced is the new middle. Most enterprises have moved from building foundations to scaling AI. The barriers we identified a year ago have not disappeared, but enterprises have made enough progress to shift the maturity curve significantly. Getting to Advanced is a matter of investment and deployment. Becoming a Leader requires changing how the enterprise operates to convert that deployment into value, and few have made that leap yet.

Figure 1. Advanced is the new middle as AI maturity shifts sharply upward



Sample: 2,005 global enterprise executives
Source: Genpact, 2026

More than half of respondents are now Advanced or Leaders. The evidence of what this means is in how agents and decisions now operate. Keith Lapointe, *Vice President of Global Business Services at Olympus*, shares his views on the past 12 months of AI progress:



A year ago, there were dozens of stories about companies not realizing value and whether AI is really all it's cracked up to be. In the last 12 months, we have seen it accelerate dramatically, and people are starting to see the value pull through. This is real, and we need to keep going."

Keith Lapointe
Vice President of Global Business Services, Olympus

What defines an **AI Leader**

Leaders are not defined by AI adoption alone. They are the enterprises where AI deployment, decision authority, ways of working, technical capability, secure governance, and value realization advance together. They redesign how work gets done, shifting from predominantly human-led processes to new ways of working in which agents execute defined work, with people providing judgment, guidance, and validation. Autonomy is not a collection of technology initiatives; it is a fundamental change in how the enterprise works.

Our AI Maturity Index assesses respondents' capabilities across seven dimensions. The full methodology is provided in the research methodology section.

Agentic AI adoption is progressing from pilots to production

While AI adoption is only one facet of maturity, it's a critical element to understand because it reflects how enterprises are scaling the use of AI. When we look specifically at the growth in agentic AI adoption over the past year, 79% of respondents report AI agents live in production (Figure 2). While the distribution shows different levels of in-production deployment scale, with more than a quarter identified as isolated use cases and just 20% reported in broad and consistent use across functions, it reflects meaningful progress from pilot purgatory to production use. Process penetration also tells us that agents are no longer confined to peripheral experiments. When we asked respondents to indicate what percentage of their processes are now primarily run by AI agents, two out of five enterprises (40%) say agents now primarily execute or manage at least a quarter of their operational processes.

The movement reflects how agentic AI is now operating inside real business workflows, giving enterprises practical experience of what it takes to run, govern, and scale it. Marina F. Bellini, *President of Global Services and Digital Technologies at Mars*, comments on the nature and pace of the company's AI progression:



We started with a very back-office and consumer-marketing-oriented focus. Now we're going more directly into value chain processes like sales, supply chain, and the whole employee experience. The progression isn't at the same pace in every area, but it's going broader, covering more business processes."

Marina F. Bellini
President of Global Services and Digital Technologies, Mars

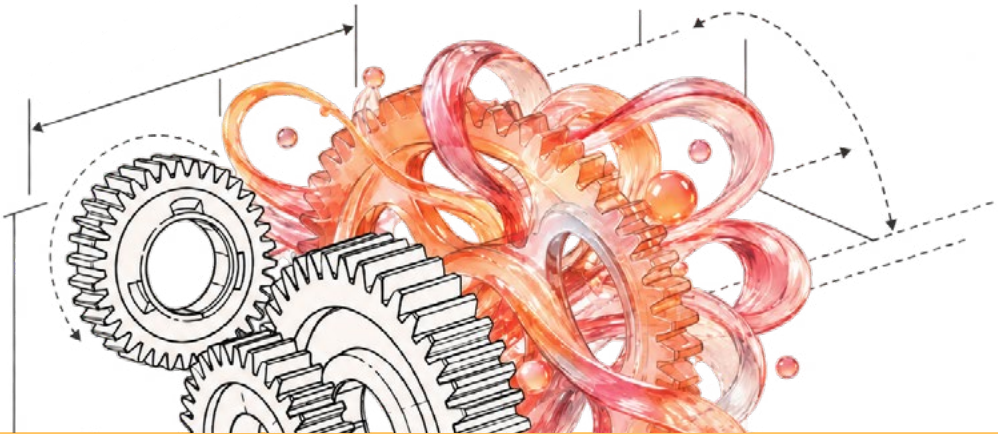
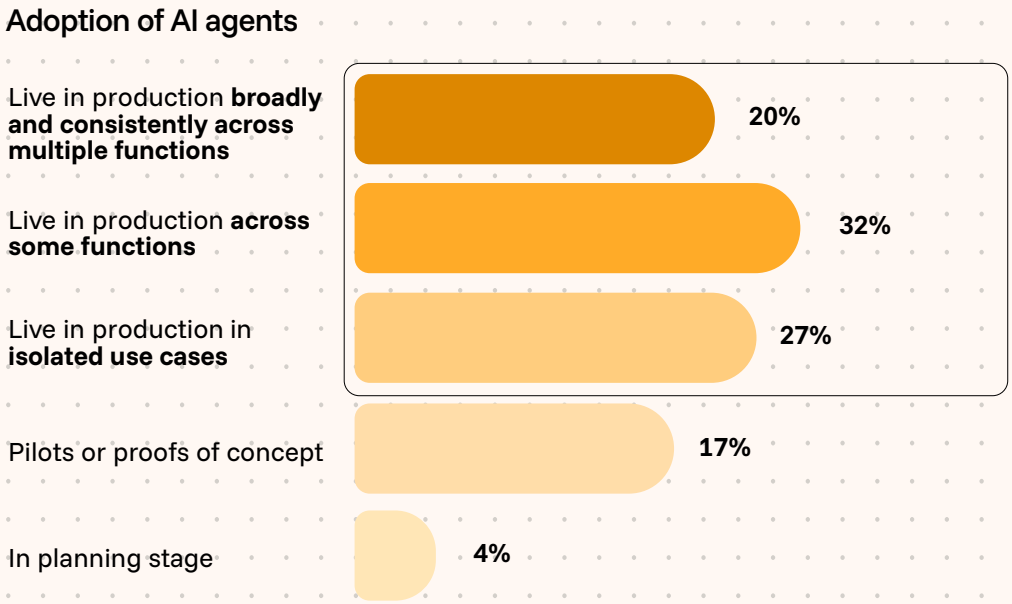


Figure 2. Agentic AI adoption is progressing from isolated production to broader enterprise use across functions

Q. Please characterize the current state of adoption of AI agents across your organization. (% of respondents)



Q: Approximately what percentage of your organization's operational processes are currently executed or managed primarily by AI agents?

Enterprises where agents run 25%+ of operational processes.



Two in five enterprises report that agents execute a quarter or more of their operational processes.

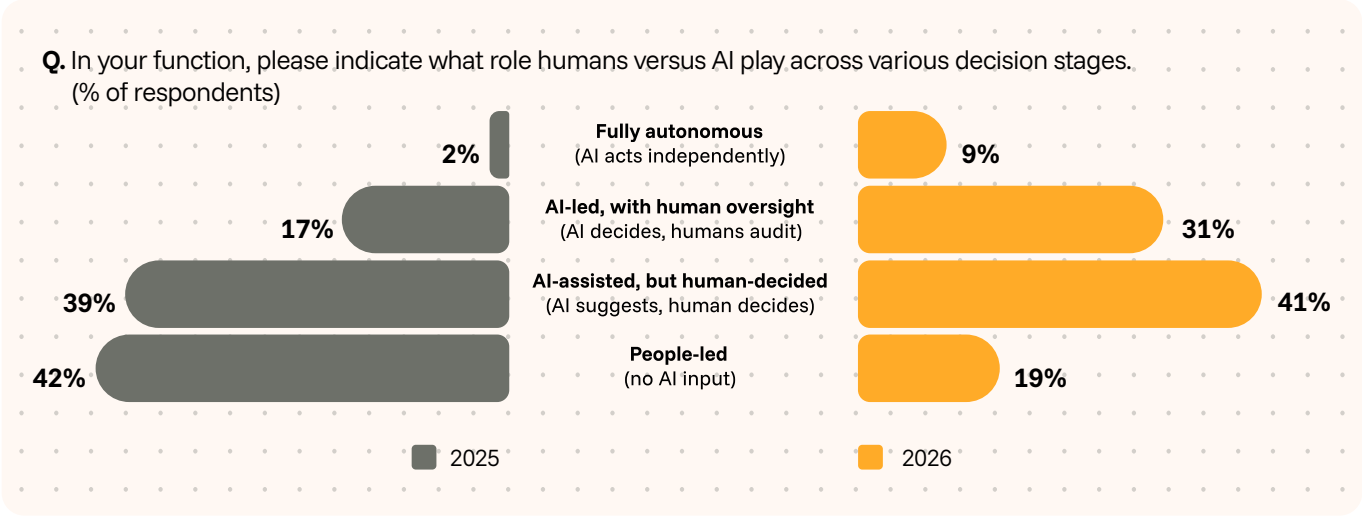
Sample: 2,005 global enterprise executives
Source: Genpact, 2026

AI is moving into the **decision loop**

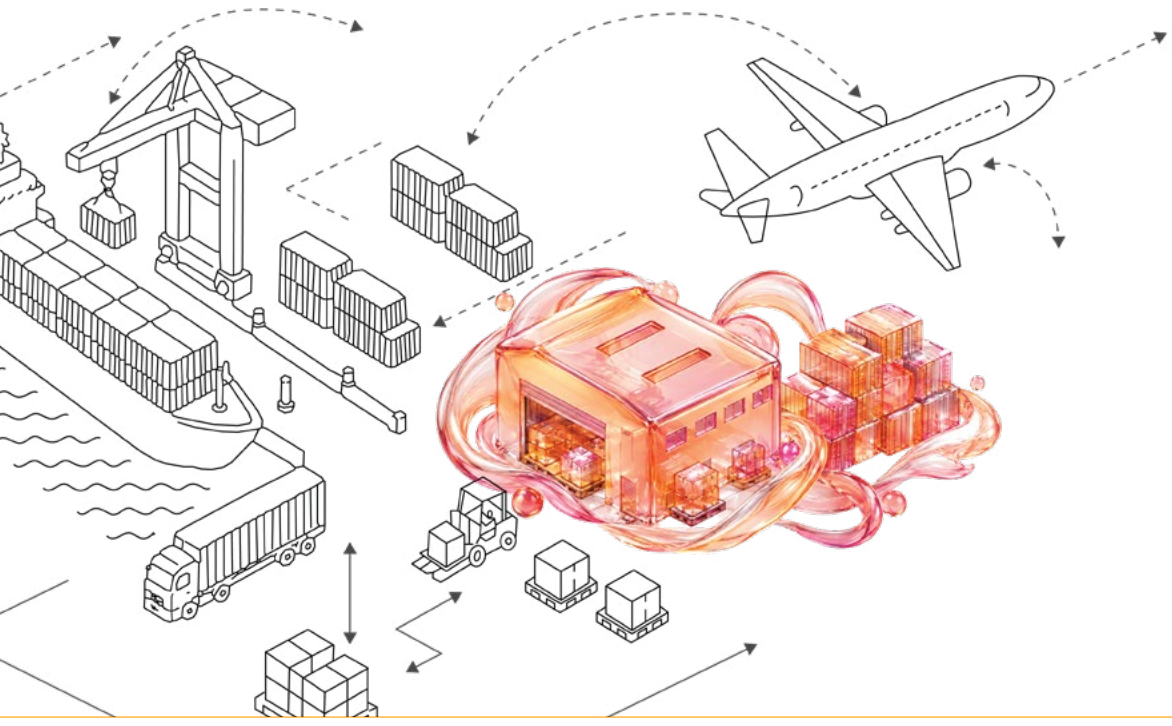
The year-on-year movement in AI decision authority is significant. Averaged across the eight decision stages we measured, people-led execution with no AI input has fallen from 42% in 2025 to 19% in 2026. AI-led decisions (where AI decides and humans audit) and fully autonomous execution (where AI acts independently) together have more than doubled, from 19% to 40%, as shown in Figure 3.

AI-assisted, human-decided execution has remained broadly stable at around 40%, underscoring that the shift is not simply from human work toward greater AI assistance. Enterprises are allowing AI to assist and increasingly lead decisions.

Figure 3. Enterprises are allowing AI to assist and increasingly lead decisions



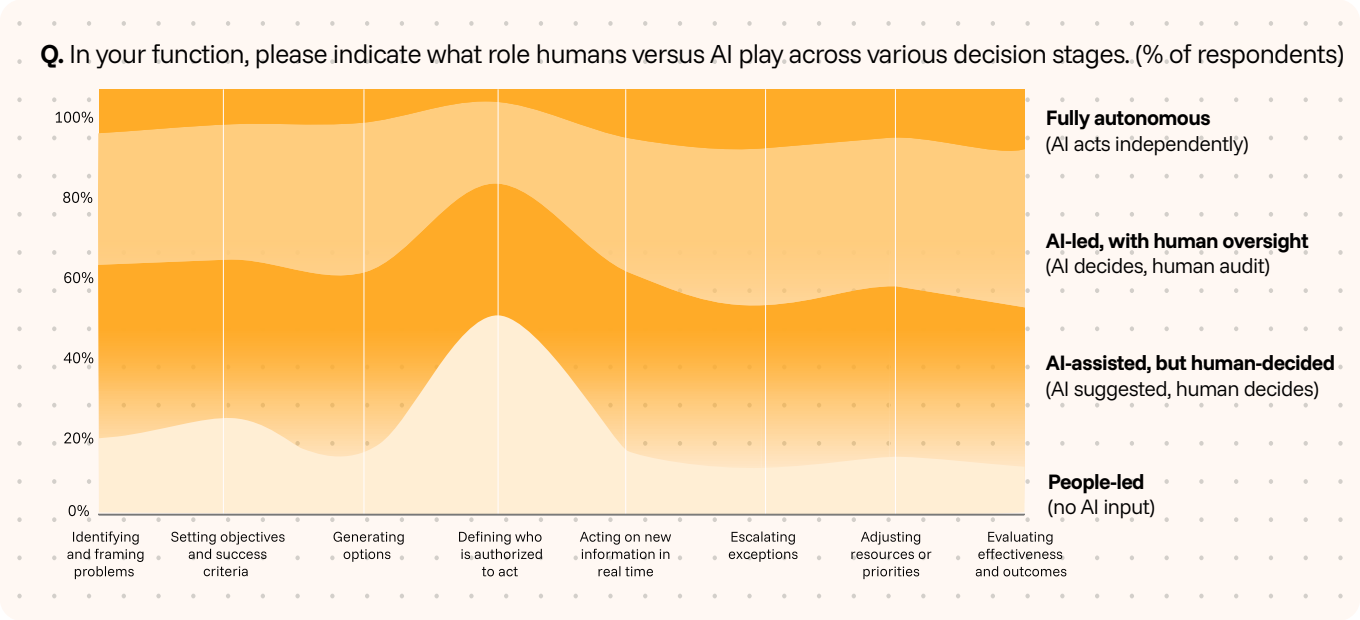
Sample: 2,005 global enterprise executives
Source: Genpact, 2026



This is not full autonomy. Fully autonomous execution averages 9% and ranges from 2% to 13% across the eight decision stages we measured (Figure 4). Humans still decide or audit in 87%–98% of cases. For some enterprises, this reflects limited trust and immature governance. For others, it is a deliberate choice when, for regulatory, legal, or risk reasons, full autonomy is not the ultimate end goal. A bank executive describes the distinction:

> **We distinguish between AI exercising delegated authority and AI determining its own authority. An organization can let an agent make decisions, but humans should always be responsible for establishing the boundaries, particularly in banking. The target isn't one autonomy level across the operating model. AI can have greater autonomy in generating options, responding to information, and evaluating outcomes, while humans retain stronger control of objectives and authority."**
>
> *Strategy & Transformation Leader at a large bank*

Figure 4. Greater AI decision authority does not mean full autonomy



Sample: 2,005 global enterprise executives
Source: Genpact, 2026

What this looks like in **practice**

To make AI decision authority concrete, we examine how the following four categories play out across a single function within finance and accounting: accounts payable (AP). We chose AP because virtually every enterprise runs it, regardless of industry. The examples below draw on Genpact's experience running agentic AP operations.

People-led. Human teams validate nearly every invoice because legacy extraction tools still require substantial checking and correction across varied document formats and line-item data.

AI-assisted. AI takes on invoice data extraction at header and line-item level, while humans shift from reviewing every invoice to handling only exceptions outside defined tolerances.

AI-led. Matching and exception-identification agents resolve extraction errors, mismatched remit-to data, and unit-of-measure discrepancies. Humans audit the exception queue.

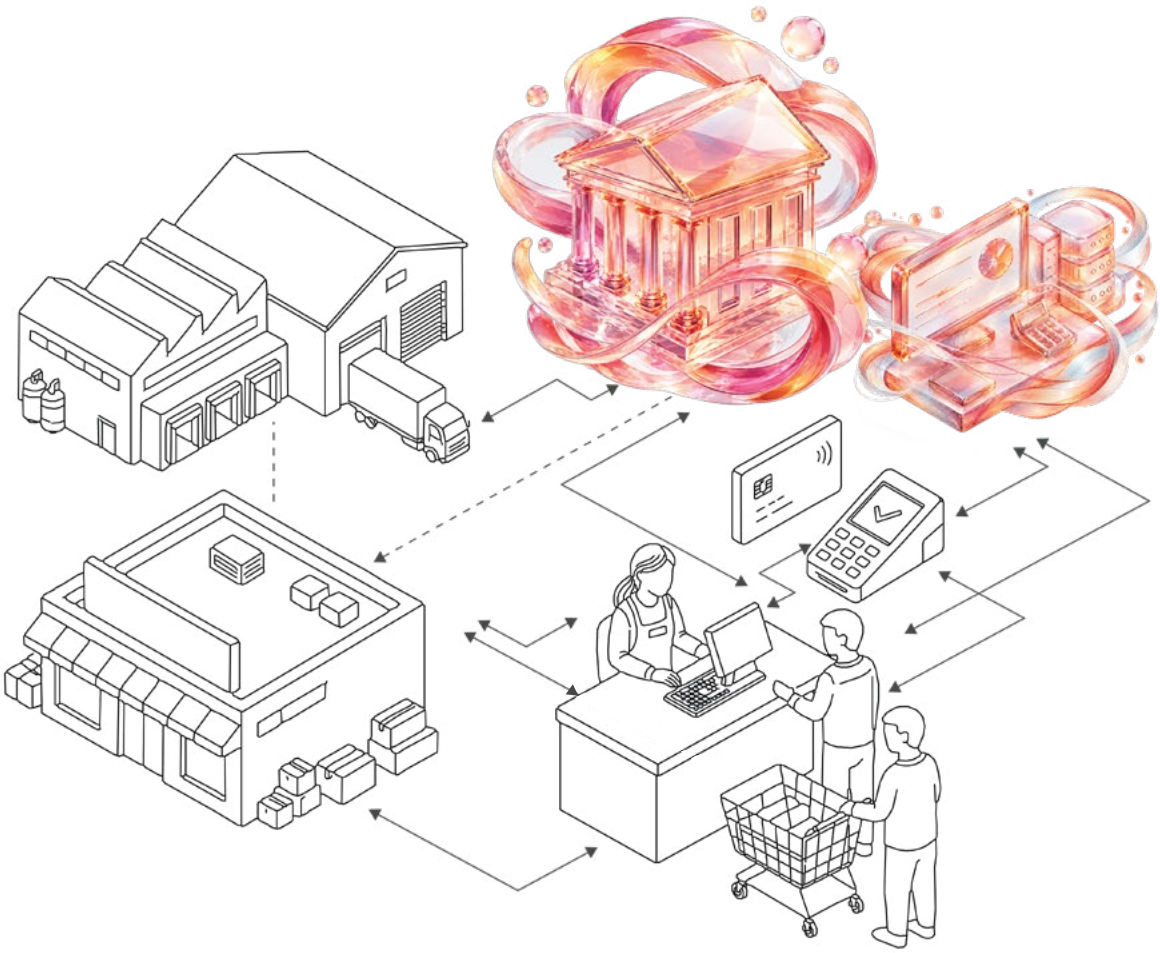
Highest autonomy. Agents process standard, in-tolerance invoices end to end within defined business rules. Across the AP operations Genpact runs, straight-through processing ranges from 45%–90% of volume, depending on the client. Higher-risk exceptions are managed by people.

Source: Genpact operations data, 2026. Results vary by client and depend on client systems, process maturity, data quality, implementation scope, and configuration. AI agents operate within client-defined business rules, with exceptions and higher-risk items subject to human review.

Context matters greatly here. A retailer can reverse errors that a healthcare provider, bank, insurer, or industrial operator cannot. A warehouse agent that recommends where inventory should move is different from one that physically triggers the move. The Chief Operating Officer at a large US-based healthcare provider weighed in on this point from the context of the healthcare industry:

> **One of our guiding principles is that no AI platform, agent, clinical companion, or clinical decision support will ever be human out of the loop. Our second guiding principle is that human judgment remains essential, so you can't fully automate a function."**
>
> *Chief Operating Officer, large US-based healthcare provider*

The movement toward AI-led decision-making should not be interpreted as a linear march toward 100% autonomy. Autonomy will vary by task, risk, and context. Within the same process, agents may execute some tasks autonomously while others require human review or approval. The question is where to set the guardrails: a routine $100 invoice may be handled autonomously, while a $1 million invoice may require human judgment. Autonomy in practice means deliberately defining these boundaries, then building the agentic operations that allow people and agents to work within them.



Enterprises are failing to capture the **autonomy dividend**

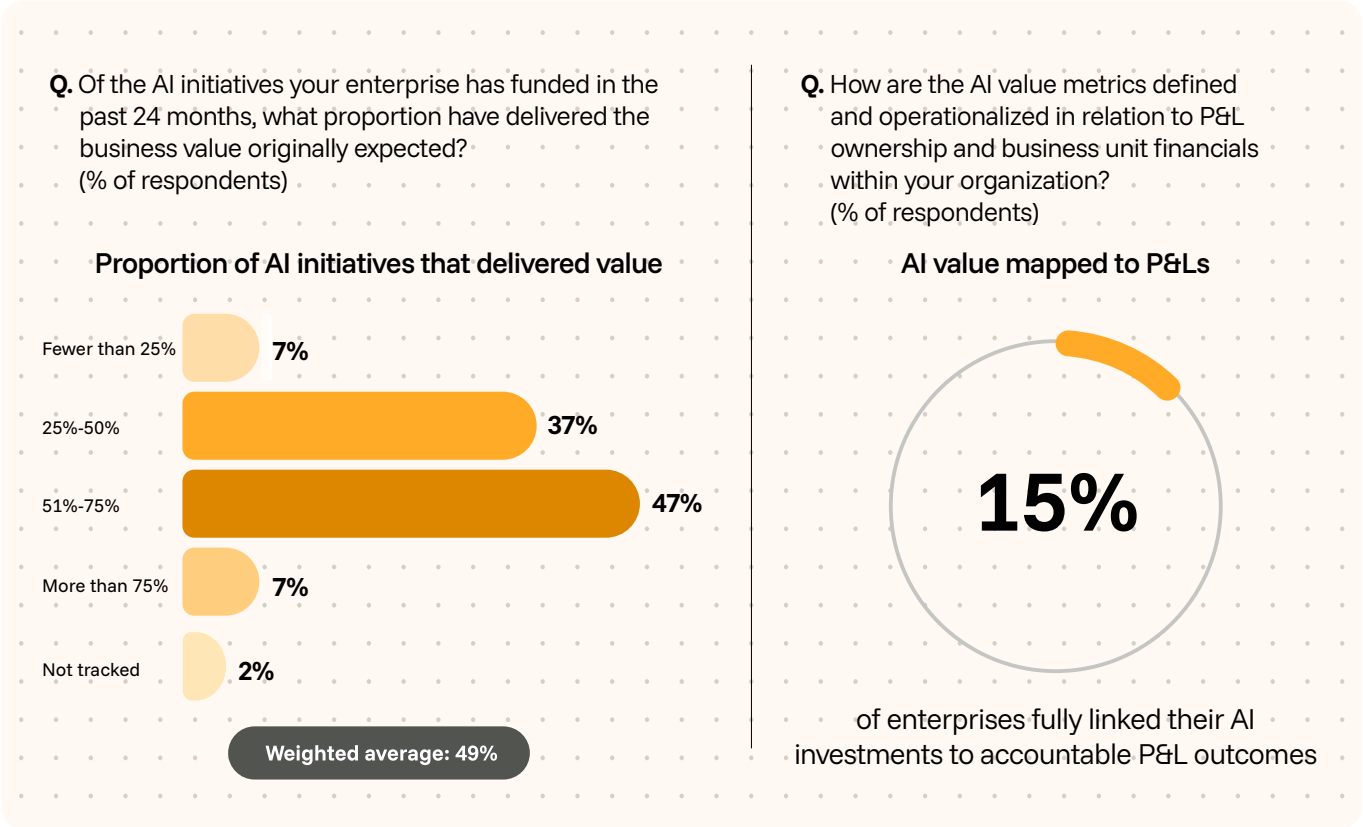
AI progress is accelerating, but enterprise value realization is not keeping pace. The autonomy dividend is the enterprise value captured when AI-enabled gains translate into better business performance, not simply greater productivity. Faster work, higher quality, greater capacity, and better decisions create the potential for value. Enterprises capture the dividend when those gains help them grow faster, operate leaner, convert more cash, or deliver other measurable business outcomes.

Our research shows that enterprises are capturing only part of this potential value. The shortfall shows up in three areas: execution, business ownership, and capacity utilization.

It starts with execution. Looking back at AI initiatives funded over the past 24 months, respondents report that a weighted average of just 49% delivered the business value originally expected (Figure 5). Much of the execution remains fragmented at the task and functional level, making it difficult for enterprise leaders to see how individual gains add up to meaningful business value.

The second is business ownership. Although enterprises are tracking AI value, few connect it directly to financial ownership. Just 15% of respondents have fully integrated their AI investments into specific P&Ls. Without a named owner and financial target, AI-enabled gains can remain visible within a workflow without ever reaching enterprise performance.

Figure 5. Only half of AI initiatives deliver expected value, and just 15% are fully connected to P&L outcomes



Sample: 2,005 global enterprise executives
Source: Genpact, 2026

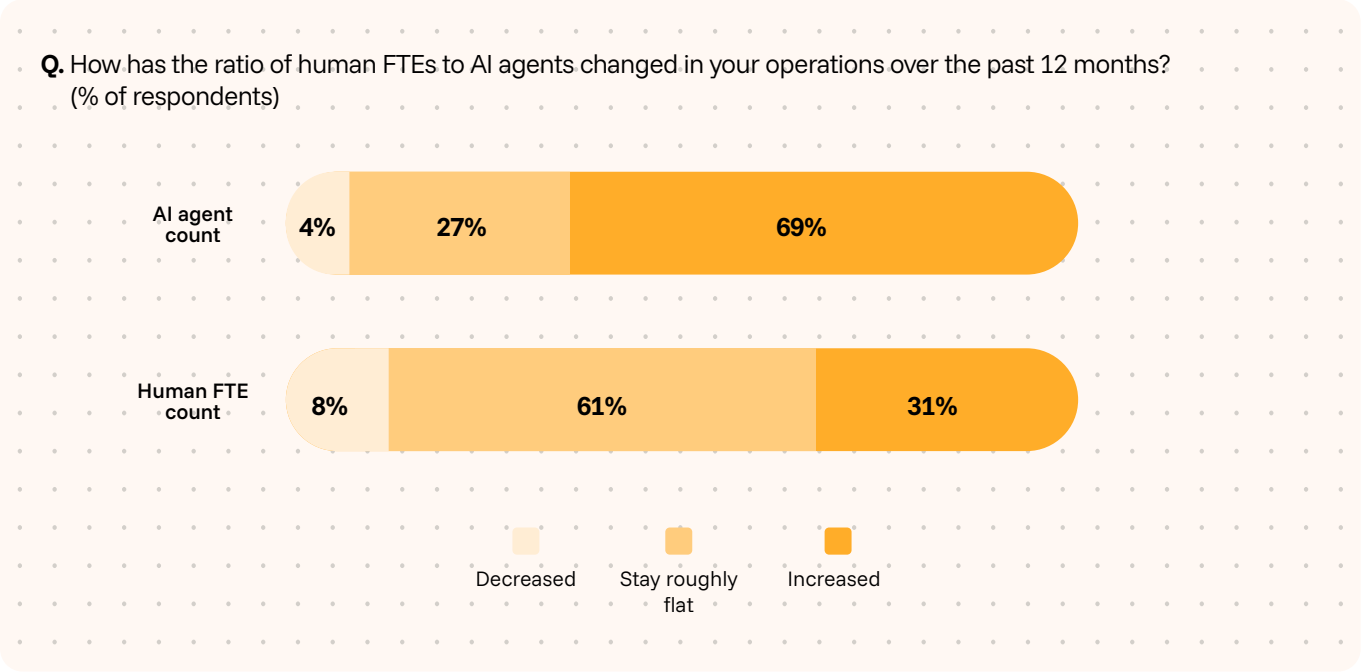
AI agents are **scaling alongside people**

The third shortfall is capacity utilization. The autonomy dividend cannot be captured through a business case based solely on replacing people and counting labor savings. That reduces the transformation of work to a one-time cost calculation and overlooks the larger opportunity: redesigning work so people and agents together increase the enterprise's capacity to grow, operate more effectively, make better decisions, and respond faster.

Our data reinforces this point. Among respondents with agents in live production, 69% report that their agent population has increased over the past year. Human headcount has remained roughly flat or increased at 92% of these enterprises; just 8% report a decrease (Figure 6).

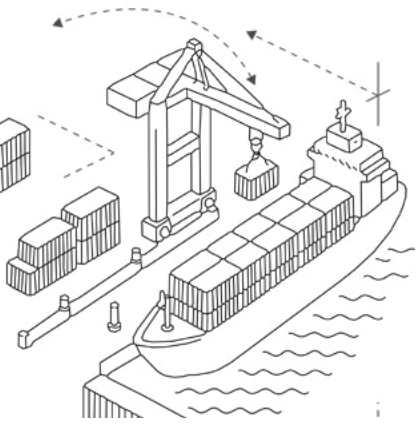
Agents are being added, but people are not simply being removed. The economic question is not how many jobs AI replaces but whether enterprises can convert their expanding human-agent capacity into better utilization and business performance.

Figure 6. AI agents are scaling alongside people



Sample: 1,584 global enterprise executives with AI agents in live production
Source: Genpact, 2026

According to the Vice President at a US-based life insurance firm, scaling agents alongside people is a deliberate choice being used as a means of growth:



> **The message from our most senior leadership has been from the beginning that we're investing in AI not to reduce headcount, but instead to grow our business."**
>
> *Vice President, a US-based life insurance firm*

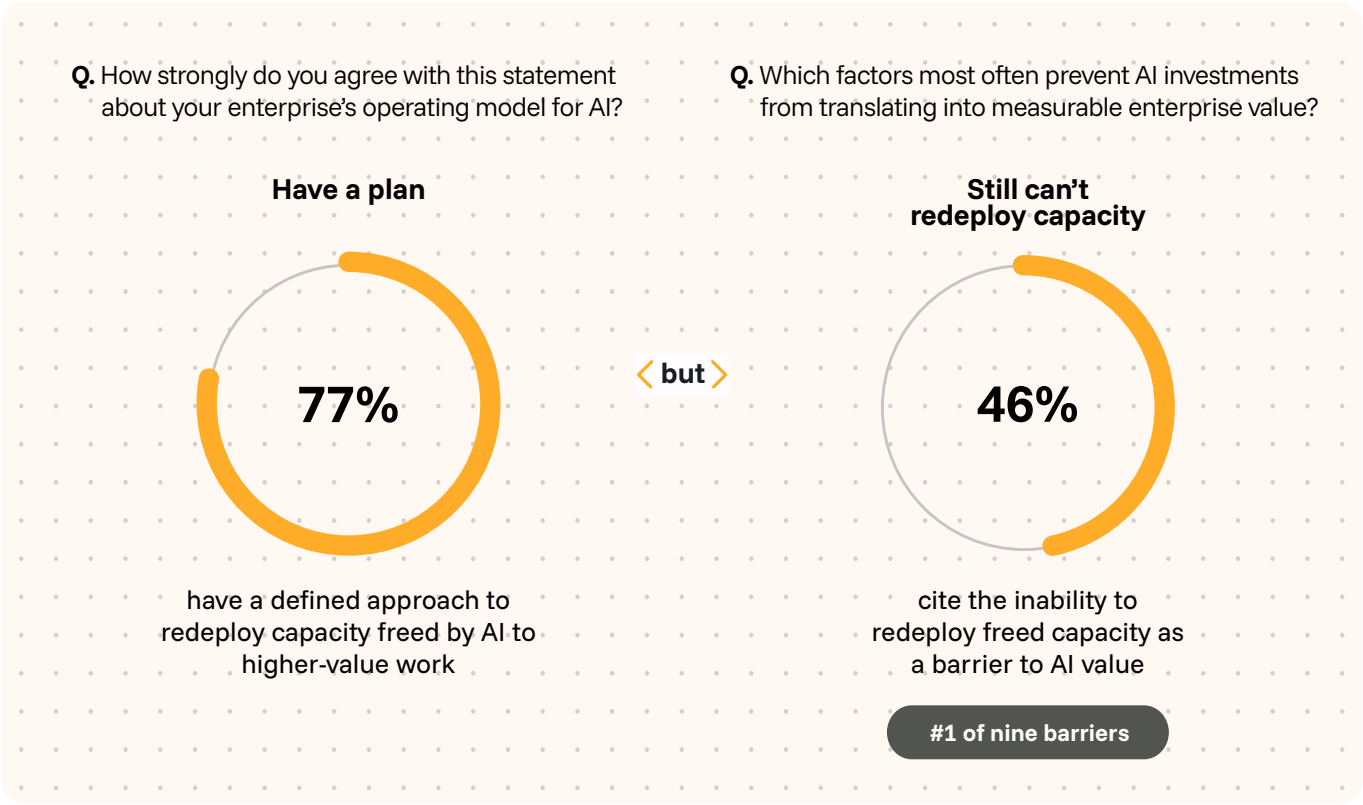
Enterprises **struggle to redeploy the capacity** freed by agentic AI

As AI agents scale alongside people, enterprises are creating new capacity. Capturing the autonomy dividend depends on directing that capacity toward higher-value or entirely new work.

Our data shows a revealing contradiction. Seventy-seven percent of enterprises say they have a defined approach to redeploying capacity freed by AI to higher-value work (Figure 7). Yet the inability of the workforce or operating-model to redeploy that capacity is the most-cited barrier to turning AI investment into measurable value (46%).

The disconnect is between capacity identified on paper and capacity redeployed in practice. Freed time is not freed capacity until someone owns where it goes. And freed capacity does not contribute to the autonomy dividend until it changes a business outcome.

Figure 7. Planned capacity isn't redeployed capacity



Sample: 2,005 global enterprise executives
Source: Genpact, 2026

Capturing the autonomy dividend

The models work, and agents are scaling. But the foundations around them have not kept pace. Most enterprises are inserting agents into structures built for human-only work. Tasks get faster and hours are saved, but these gains do not automatically translate into enterprise performance.

Capturing the autonomy dividend requires a deliberate way of operating that converts these gains. Freed capacity must be deliberately redeployed. Better decisions and execution must improve growth, cost, cash, quality, resilience, or risk. And those outcomes need named owners and financial accountability.

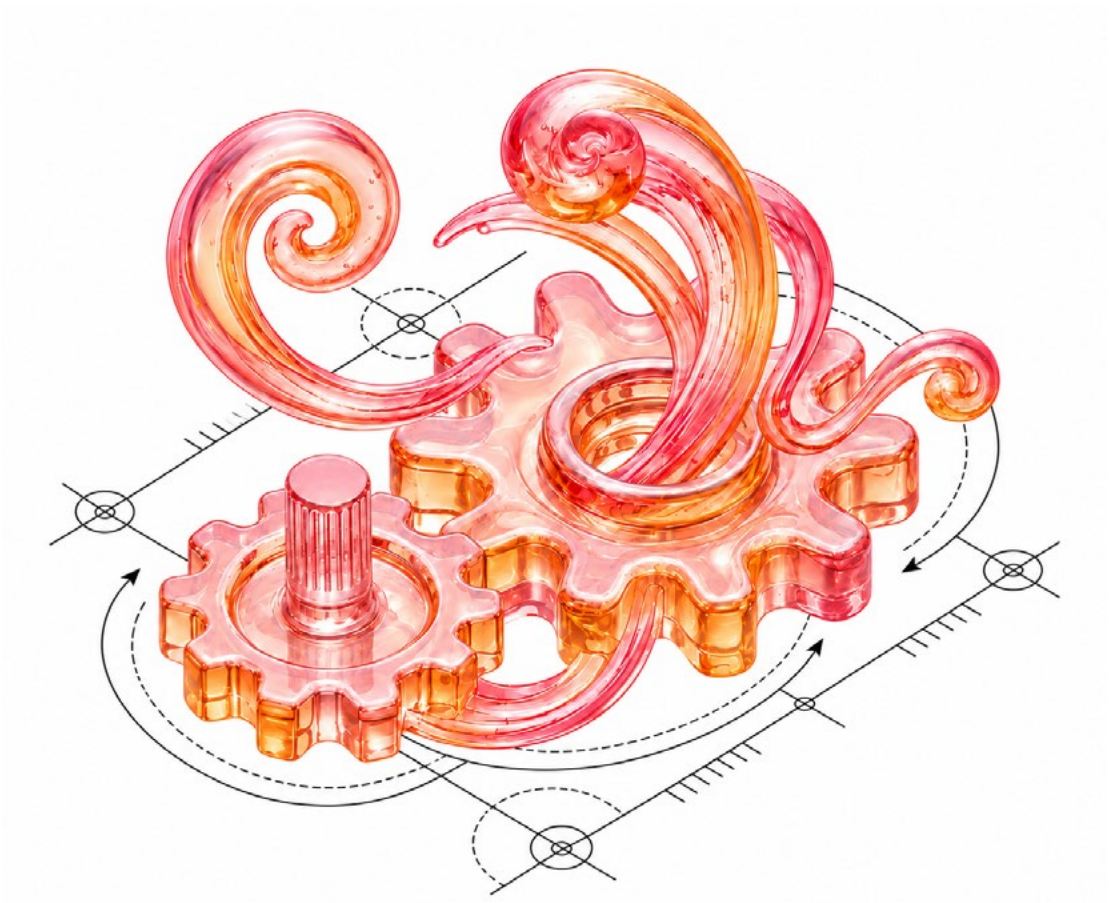
This requires more than deploying agentic AI. It means changing how work is designed, how people and agents contribute, how data and technology enable them, how their actions are secured and governed, and how the resulting value is captured. As Marina F. Bellini, *President of Global Services and Digital Technologies at Mars*, explains:



We can only get true value from AI in our processes if we rethink the operating model and if we care for data. The technical challenges are "easy" to overcome. But when it comes to rethinking how operating models are put in place, and which data you need to make AI work, if that's not addressed well, you end up with a series of AI use cases, but not AI business cases."

Marina F. Bellini, *President of Global Services and Digital Technologies, Mars*

Capturing the autonomy dividend requires the shift from deploying AI to operating with it. Agentic operations is how enterprises can put autonomy into practice, orchestrating people and agents to help deliver faster growth, leaner operations, and greater conversion to cash.



Agentic operations is how Leaders capture the autonomy dividend

Agentic operations is the operating model that runs people and agents as one workforce. We measured enterprise progress across five building blocks critical to the operating model: process reinvention, a human-agent workforce, a data and AI foundation, governance and security, and value capture. Leaders are ahead on all five. Enterprises capture the autonomy dividend when these building blocks work as a system, connecting redesigned processes and human-agent work with the technology, controls, and financial accountability needed to turn AI progress into economic outcomes.

Jo Thorn, *Global Finance Operational Excellence Lead at News Corp*, explains that the challenge is coordinating these changes rather than advancing any one element in isolation:



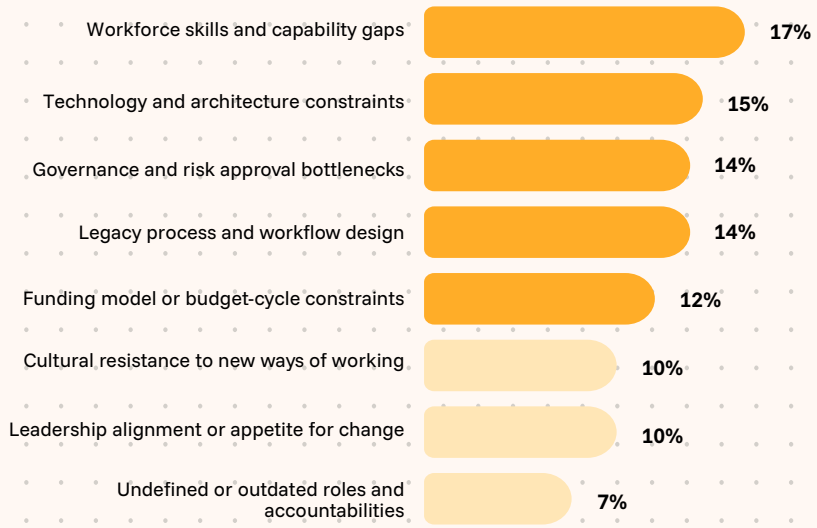
Realizing the full value of AI requires more than investing in technology. It demands coordinated change in how we work, how we organize ourselves and how we build capability. People, culture, and the operating model must evolve together. Without that alignment, organizations will struggle to achieve the outcomes they expect from AI."

Jo Thorn
Global Finance Operational Excellence Lead, News Corp

These building blocks are also where enterprises face the greatest strain. All five appear among the leading constraints on operating model transformation, and the underlying data reveals distinct weaknesses within each one (Figure 8). Workforce skills and capability are the most-cited constraint, while technology and architecture, governance and risk approval, process and workflow design, and funding constraints all feature prominently.

Figure 8. The barriers to agentic operations span people, technology, governance, process, and value

Q. What is the single biggest roadblock to reaching your target operating model — i.e., what, if removed, would most accelerate the change?
(% of respondents)



Sample: 2,005 global enterprise executives
Source: Genpact, 2026

Leaders show what it looks like when the building blocks work as a system. They are the 17% of respondents at the top of our AI Maturity Index, and their advantage is not simply greater adoption of agentic AI. They are further ahead in bringing the building blocks together for running people and agents across the enterprise. **The rest of this section takes each building block in turn: where most enterprises are today, what Leaders do differently, and what it looks like in practice.**

It starts with the **CEO**

Building agentic operations requires changes that cut across functions, technologies, processes, workforce decisions, governance and security, and financial ownership. No single function can make these changes alone. Function owners may drive execution day to day, but Leaders are more likely to make agentic operations a CEO-level mandate: 44% report that the CEO leads the AI change mandate, compared with 29% of the rest of the study.

44%
Leaders

compared to

29%
Rest of the study

report that the CEO leads the AI change mandate.

The CEO's role is not to run every initiative. It is to establish the enterprise ambition, align business and technology leadership, resolve competing priorities, and maintain accountability for economic outcomes. Execution then runs across the leaders responsible for the five building blocks.

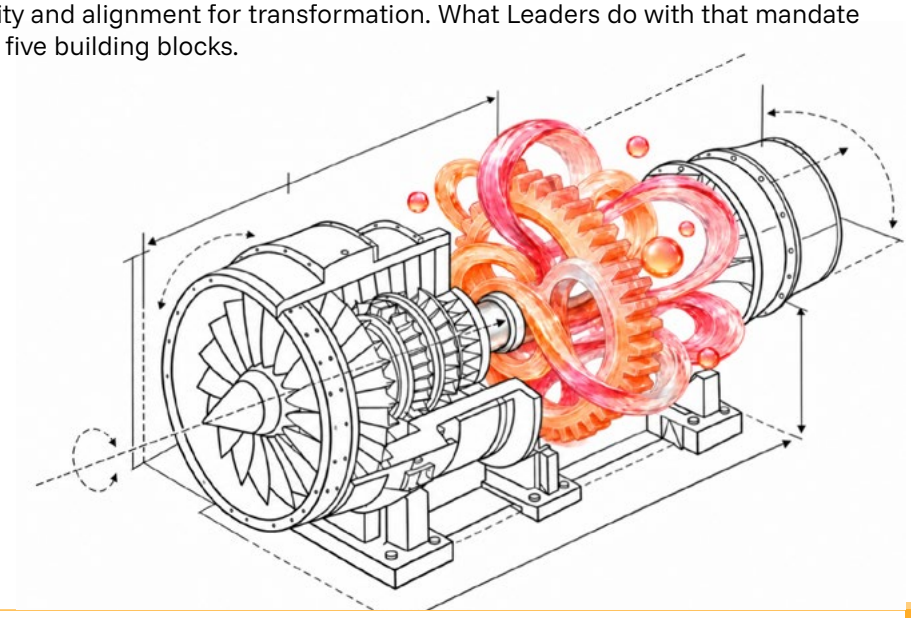
Caitlyn Carr, *President and CEO of Coca-Cola Bottlers' Sales and Services,* describes why:



I am far more involved in AI-driven operating model transformation than I've been in any kind of other program. This is one of those transformations that requires leadership to not only set the stage and define an end state of vision, but to go through it themselves and be open about their own learning experiences, their own trials and tribulations. It's required because of the amount of change and the speed of the change that you're asking the whole organization to go through.

Caitlyn Carr
President and CEO, Coca-Cola Bottlers' Sales and Services

A CEO mandate creates the authority and alignment for transformation. What Leaders do with that mandate becomes visible across each of the five building blocks.

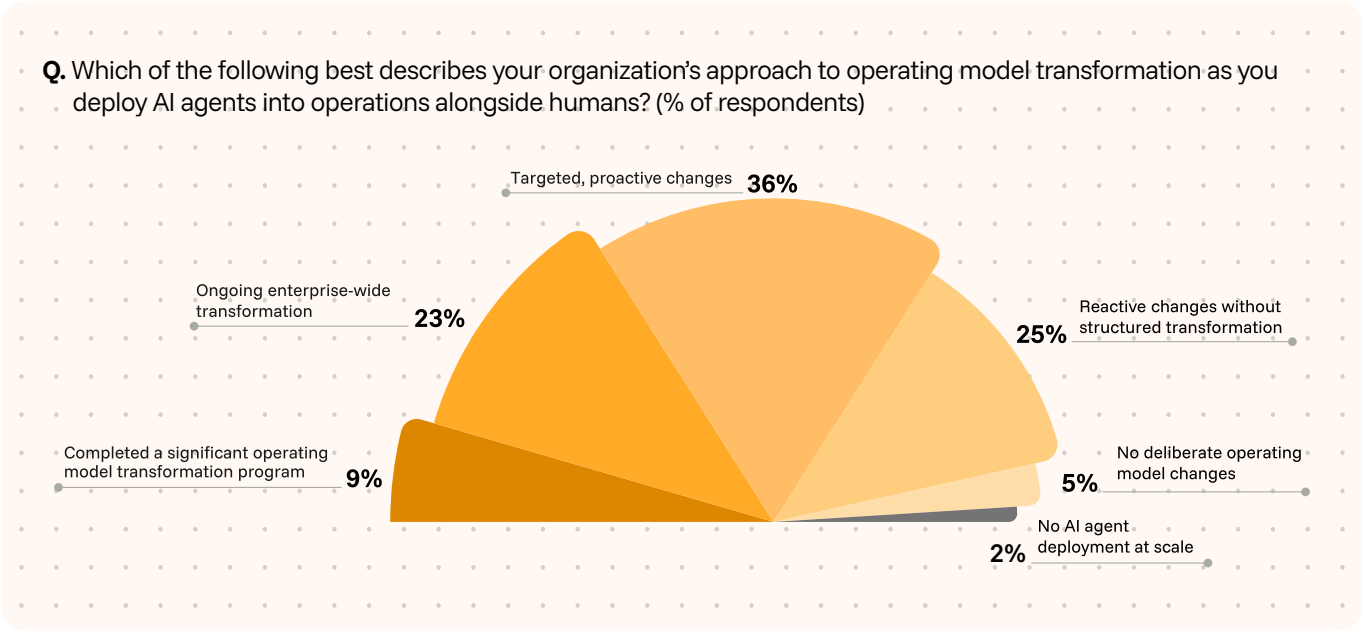


1. Process reinvention: most enterprises change function by function. Leaders redesign work end to end

Process reinvention starts with understanding where work should change. It means determining, process by process, where agents can take on work, where human judgment remains essential, and where workflows should be redesigned.

Most enterprises have not yet turned that understanding into broad organizational change. Only 9% have completed a significant operating model transformation, while 36% are making targeted, proactive changes, and another 25% are changing reactively without a structured transformation (Figure 9).

Figure 9. Most enterprises are changing work incrementally, not end to end



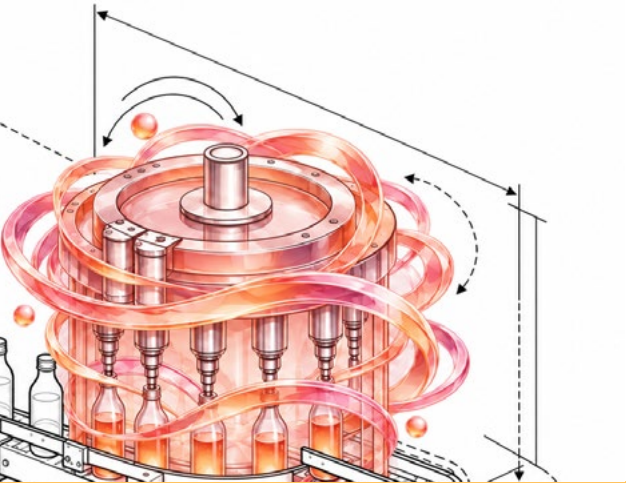
Sample: 2,005 global enterprise executives
Source: Genpact, 2026

Process-by-process change is not the problem. Disconnected change is. Process reinvention is what connects changes across processes, functions, systems, and decision points. Kevin Ichhpurani, *President, Global Partner Ecosystem at Google Cloud,* shares an ecosystem perspective on grounding that change in how work actually gets done:



Process intelligence earned over decades informs how AI is applied to solve real-world problems."

Kevin Ichhpurani
President, Global Partner Ecosystem, Google Cloud



Leaders treat process reinvention as an enterprise program. A majority (51%) are running proactive, enterprise-wide operating model transformation specifically to enable agentic AI compared to just 17% for the rest of the study. They redesign beyond individual tasks, defining what the agent owns, what the human owns, how authority passes between them, and how workflows connect across functions and systems. That end-to-end view keeps an improvement in one area from simply moving the bottleneck somewhere else.

51% compared to **17%**

Leaders **Rest of the study**

are running proactive, enterprise-wide operating model transformation

This coordination challenge is as much about leadership as method. Krishna Cheriath, *Chief Digital and AI Officer at Thermo Fisher Scientific, Clinical Research,* describes the kind of leader it takes:



AI transformation requires people who are boundary spanners. If you want to affect change, you need to be part technologist, part strategist, part storyteller, part visionary, part shrink. It's less about the organization and more about whether you have leaders who understand the challenges and needs across roles and functions."

Krishna Cheriath
Chief Digital and AI Officer, Thermo Fisher Scientific, Clinical Research

Case study

From 27% to 71% straight-through processing, by redesigning AP for humans and agents

A global automotive components supplier was processing nearly 50,000 invoices a month across two legacy ERPs that didn't talk, with 38% of invoices stuck on exception hold. Genpact had run its AP operation for several years, which is how both sides knew incremental improvement had hit its ceiling.

They redesigned the work rather than layering agents onto the old process. Twelve agents support invoice capture, matching, validation, and ERP posting, with postings subject to defined approval controls and human oversight. An expert finance team owns exceptions, system tuning, and the judgment calls. Decision rights and escalation paths were defined before go-live; the AP team moved into seven newly defined agentic roles, and a shared command center gives both organizations one view of every exception.

From January to June 2026, straight-through processing rose from 27% to 71%, capture accuracy reached about 98%, and exception holds fell from 38% to 20%. A new duplicate-rejection control now stops a quarter of duplicates upfront. When the company closed an acquisition, the same agents, governance, and platform absorbed the new entities without a rebuild.

**Client-specific results. See the disclaimer section for more details.*

2. Human-agent workforce: most enterprises adapt roles. Leaders redesign them for people and agents

Redesigning operations changes the work people do and how they do it. As agents take on more execution, people increasingly handle judgment, exceptions, and oversight. They also build, train, and improve the agents they work alongside. These capabilities may sit across internal teams and partners, changing the skills enterprises need, the roles they design, and how work is organized across organizational boundaries.

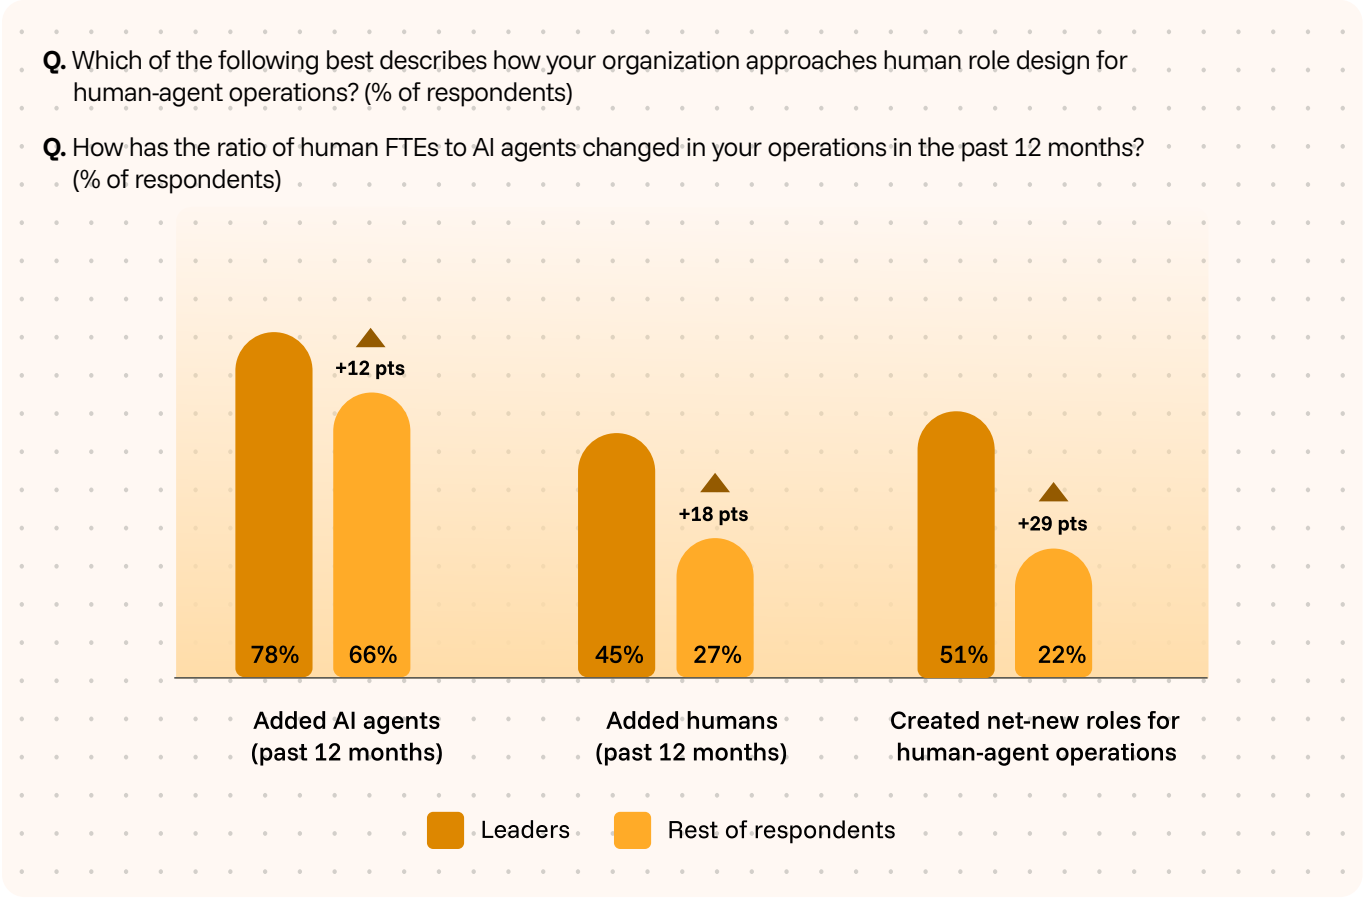
When asked about the skills required for work alongside AI agents, 47% of respondents prioritize techno-functional talent, people who combine technical and business expertise, ahead of purely technology-first (34%) or domain-first (18%) profiles. Because the profile is a hybrid, it has to be built through training: domain experts learn to evaluate and improve agents, and technologists learn the business.

Enterprises are also more likely to adapt existing roles (39%) than to create new ones (27%), so much of the transformation depends on reskilling the current workforce. Adapting roles means adding new capabilities such as agent oversight to existing responsibilities, while net-new roles, such as exception specialists and agent trainers, are designed around the agents from the start.

Leaders make a different choice. Just over half (51%) create net-new roles built for human-agent operations, such as exception specialists, agent trainers, and output auditors. Only 22% of the rest of the respondents take this approach.

Hiring follows the same pattern. In the past 12 months, 78% of Leaders have added agents (versus 66% of others), and 45% also added humans (versus 27%). Leaders are not swapping people for agents. They are building both sides of the human-agent workforce at once (Figure 10).

Figure 10. Everyone is adding agents. Leaders are adding the people and roles to run them.



Sample: 1,584 - 2,005 global enterprise executives. Human and agent data excludes respondents who have not deployed agents
Source: Genpact, 2026

Shellie Molina, *Lumen Technologies Chief Procurement Officer and SVP of Supply Chain*, describes what intentional role design looks like:



My charter was to re-imagine the procurement organization, so we flipped it from procurement to strategic sourcing. We completely changed the job responsibilities, from purchasing managers to category leaders and category managers, and we were very explicit in the must-haves: domain plus experience building AI agents and managing agents. The change has paid dividends."

Shellie Molina
Chief Procurement Officer and SVP of Supply Chain, Lumen Technologies

Role design is only half of it. The other half is who answers for the agents. Agents carry business logic and act inside live processes, so they straddle the line between business and IT. Ownership on one side alone leaves a gap: IT may lack the process context and the business may lack the controls. Most enterprises have not settled it. When an agent produces an incorrect output or causes a downstream error, 26% hold the technology team accountable and 23% the business leader for the process. Only 20% have a dedicated, named agent owner. Even among Leaders, no approach dominates. Where ownership is unsettled, it needs to be made explicit. In Genpact's experience, that takes a formal business-IT engagement model with shared accountability, rather than leaving it to default to whoever built the agent or whoever feels its effects.

A human-agent workforce is not simply the existing workforce plus AI. It requires deliberate choices about what agents execute, where people apply judgment, who handles exceptions, and who owns the performance of the overall human-agent system. These choices cut across traditional business and technology boundaries.

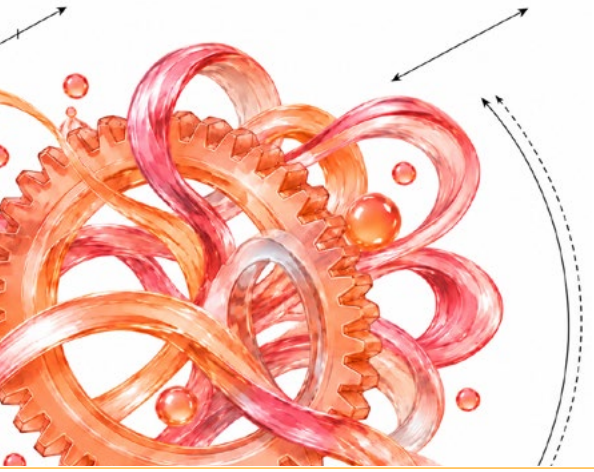
Case study

Designing the roles before scaling the agents

A telecommunications company built a global capability center (GCC) to take ownership of finance and procurement, and decided from day one that it would run on employees and agents working side by side, not automation bolted onto old workflows. The obstacles were real: fragmented tools, manual interventions, and invoices running to hundreds of pages, beyond what standard extraction could read.

Genpact brought its agentic AP Suite solution and process intelligence to engineer a custom extraction capability for the client's invoice complexity. The division of labor was designed, not assumed. Agents own extraction, validation, and workflow orchestration. Employees own oversight, exceptions, decisions, and continuous improvement. The company used a mix of role redesign and upskilling to help ensure its human talent was capable of agent oversight.

The result is a workforce model designed to let capacity scale with volume, with employees focused on oversight, exceptions, and improvements. The GCC now treats AP as the first step in extending human-agent operations into procurement, supply chain, and other functions.



3. Data and AI foundation: most enterprises are rearchitecting. Leaders have made it repeatable

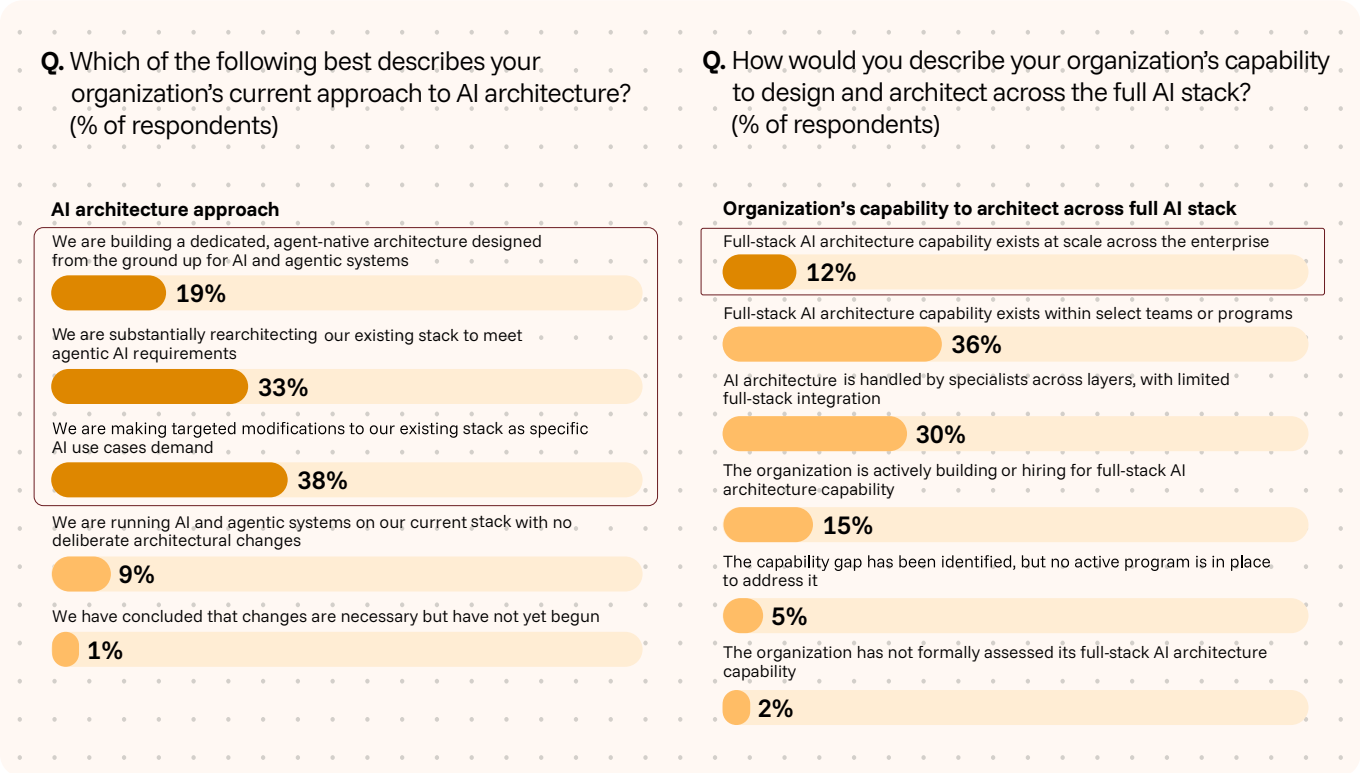
Agentic operations needs more than capable models or teams of agents. Agents need a connected foundation that brings together:

- Trusted, governed data, and semantic context
- Standardized machine interfaces, increasingly built on open standards such as the Model Context Protocol (MCP)
- Agent-to-agent (A2A) protocols for coordination across platforms
- Non-human identity and least-privilege permissions
- Orchestration and observability across applications, clouds, and models

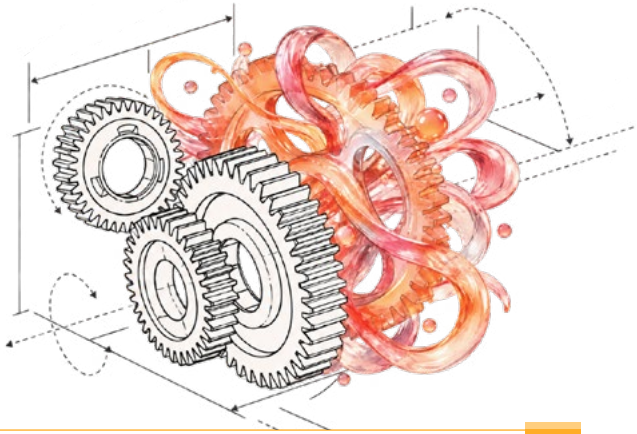
Data sovereignty and residency add further constraints: which data agents can access, where models run, and where outputs are stored.

The capabilities remain fragmented. Around 90% of enterprises report some deliberate architectural change, yet only 12% have full-stack AI architecture capability at enterprise scale. Even among those substantially rearchitecting or building agent-native environments, only 20% have full-stack capability (Figure 11).

Figure 11. Enterprises are rearchitecting, but few have built AI capability at scale



Sample: 2,005 global enterprise executives
Source: Genpact, 2026



An agent is only as effective as the environment it operates in. It may perform well within a function but still depend on data, applications, identity, or systems of record controlled elsewhere. Autonomy breaks at these boundaries: agents can retrieve information but not act on it, and workflows stall between systems. Chris Growe, *Chief Financial Officer at Advantage Solutions,* describes it as follows:



Data integrity has been one of our hardest problems. We have a lot of disparate systems, and many weren't connected to our data lake, which limited how far we could go with agentic AI. Agents are only as good as the foundation underneath them, so we're getting the architecture right first."

Chris Growe, *Chief Financial Officer, Advantage Solutions*

Technology fragmentation is compounded by fragmented ownership. Only 14% of respondents assign end-to-end architectural responsibility to a small number of architects.

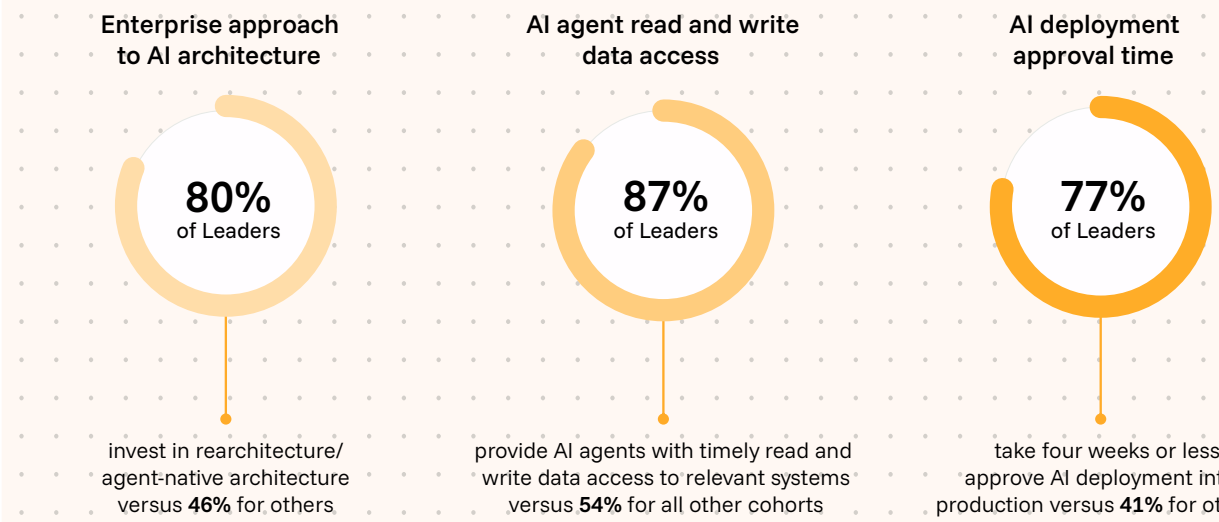
Leaders make that foundation reusable, so each new workflow builds on what already exists. 80% are rearchitecting their stack or building agent-native architectures, versus 46% of the rest of the study. Among Leaders, 87% give agents timely read and write access across the systems their workflows need, versus 54%. That is the difference between agents that retrieve information and agents that act on it. And 77% clear governance and approval for a new agent or workflow in four weeks or less, versus 41%, which suggests security and responsible AI (RAI) controls are built into the platform rather than bolted on (Figure 12).

Figure 12. Leaders are building the data and AI foundation for repeatable scale

Q. Which of the following best describes your organization's current approach to AI architecture? (% of respondents)

Q. Which of the following best describes your organization's current ability to provide AI agents with timely data access and enable them to write back or take action within systems required for their workflows? (% of respondents)

Q. From the point at which a new AI agent or agentic workflow is ready for deployment, how long does your organization's governance and approval process typically take before it can go into production? (% of respondents)



Sample: 2,005 global enterprise executives
Source: Genpact, 2026

In Genpact's experience, repeatability comes from agnostic platforms that bring the estate together: shared connectors and tool interfaces, a common orchestration layer, agent identity and access controls, and observability designed to trace agents' actions. Open standards such as MCP and A2A can help enterprises add applications, models, agents, and partners with less need to re-platform. A dynamic ecosystem of model, data, and software partners plugs into those platforms, so agents can run on the technology best suited for the job. This flexibility matters because the data and AI foundation will typically span an ecosystem of cloud providers, data platforms, model providers, and specialist capabilities.

It is a platform problem as much as a data one and in a multi-vendor estate, the platform has to be agnostic. Priya Vijayarajendran, *Genpact's Chief Product and Platform Officer,* explains what enterprises need to scale agents without engineering each one from scratch:



Agentic AI reframes the architecture challenge. It is no longer enough to give models access to good data. Agents need operational context, connections to applications, the ability to invoke tools and APIs and act across systems, all within enterprise identity and security controls. The role of the platform is to make those capabilities reusable and scalable, so enterprises can scale without engineering every agent from scratch."

Priya Vijayarajendran
Chief Product and Platform Officer, Genpact

Case study

Rearchitecting the finance data foundation for accuracy at real-time speed

A multi-billion-dollar global consumer brand's financial close didn't reconcile cleanly. When its channel mix shifted between wholesale and direct-to-consumer, the legacy data architecture began blocking strategic decisions, not just slowing operations.

Genpact delivered a unified finance data hub through a multi-year engagement, combining Genpact's finance and process expertise with Databricks as a technology ecosystem partner. The team started from pre-defined finance canonical and semantic data models and embedded automated validation throughout the migration.

The outcomes the engagement helped deliver show up on the cash line first: a cash conversion cycle cut by several days, customer payment deductions reduced by tens of millions of dollars, and several hundred million dollars in unlocked cash flow. Revenue forecasts now run at high-80s to high-90s percent accuracy, with real-time visibility across the substantial majority of revenue.

**Client-specific results. See the disclaimer section for more details.*

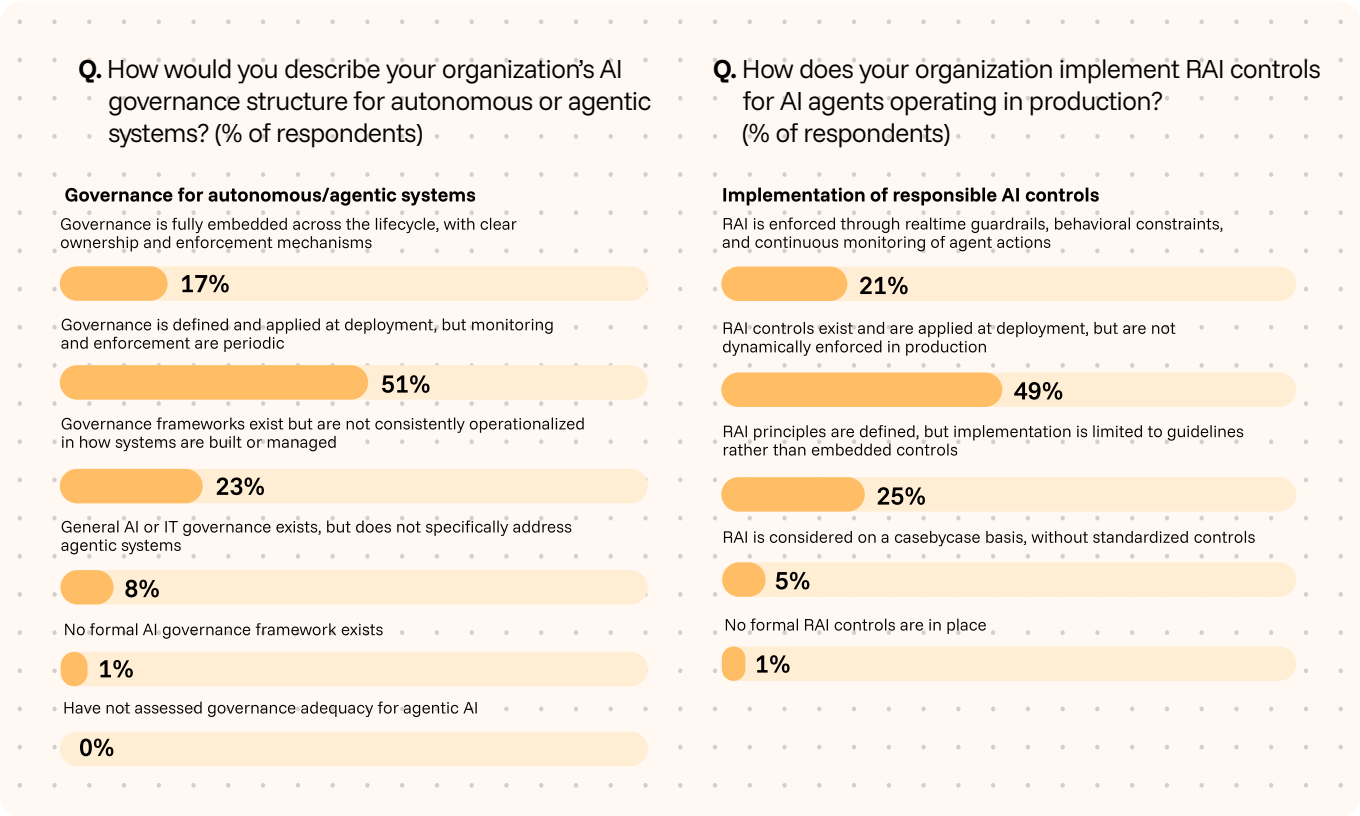
4. Governance and security: most enterprises govern at the gate. Leaders govern at runtime

Agentic operations requires governance that goes beyond deciding whether a model or agent is safe to deploy. It includes the policies, accountability, monitoring, and controls that determine what autonomous systems are allowed to do throughout their lifecycle. Security and RAI are integral to this equation. As agents access enterprise data, systems and tools, security controls around identity, permissions and access must operate alongside RAI principles such as fairness, transparency, privacy and safety. Both must translate into enforceable controls that continue to work once an agent is acting in production. This also includes agents the enterprise did not build itself, such as those supplied by technology and services partners.

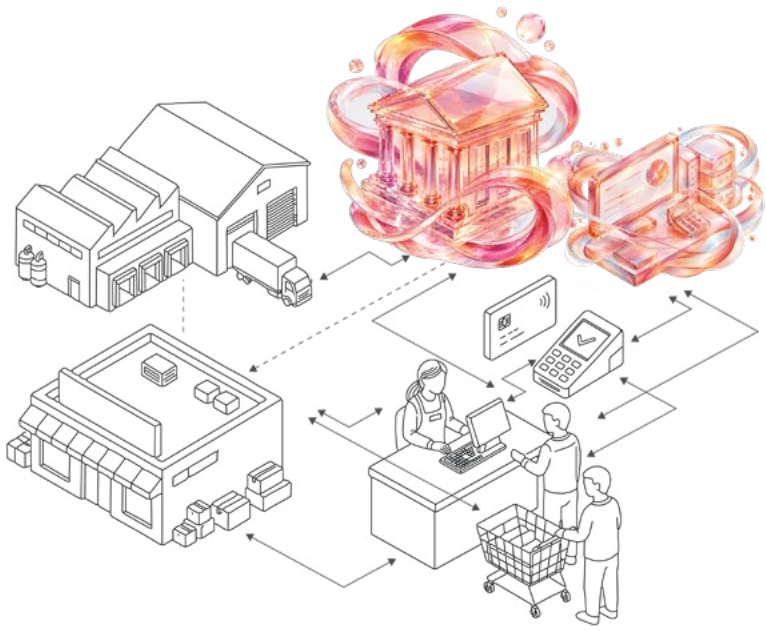
Most enterprises have built the front door to governance more strongly than the runtime controls behind it. Across all respondents, 51% apply governance at deployment with only periodic monitoring, while 17% report governance embedded across the lifecycle. On RAI, 49% apply controls at deployment without dynamic enforcement in production, compared with 21% using real-time guardrails and continuous monitoring (Figure 13).

Most enterprises have built the front door to governance more strongly than the runtime controls behind it.

Figure 13. AI governance is stronger at deployment than at runtime



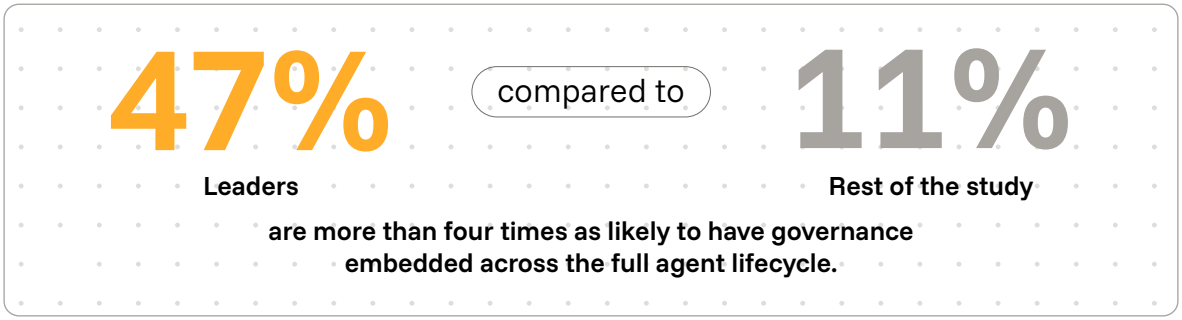
Sample: 2,005 global enterprise executives
Source: Genpact, 2026



That gap between governance and controls matters more as agents gain greater autonomy because risk can change after approval. A payment agent authorized to release routine transactions within a limit can meet a changed supplier bank account or an unexpected sequence of actions that no one assessed at deployment. Governance and security need to operate at the point of action, enforcing permissions, applying guardrails, and escalating when conditions require it.

Point-of-action control also needs someone answerable for each agent. In Genpact's experience, that means a named business owner accountable for the agent's outcomes and a technology owner accountable for its reliability, security, and integration.

Most enterprises have AI governance and security. Leaders are more than four times as likely to have governance embedded and enforced across the full lifecycle, not applied at the gate and left to periodic review. A deployment gate decides whether an agent may enter production. Lifecycle governance decides whether the enterprise can keep trusting it as data, context, users, and operating conditions change. When security and RAI are built into operations, governance can support more autonomy without sending every decision back through manual approval queues. For Leaders, it becomes an enabler of scale rather than a brake.



Case study

Rewiring a complex insurance inspection-review process for human-governed autonomy at scale

For a fast-growing specialty insurance brokerage, unidentified risks can mean lost margin and lost trust. A church kitchen, storm-prone trees over a roof, a pool that may or may not already be on the policy: each could hide in an inspection docket of up to 450 pages. And finding them could take a specialist more than two hours per file. The work demands judgment, not extraction. Rules-based automation could handle only about 20% of it.

The firm rebuilt the process rather than patching it. A multi-agent, multi-modal platform now reads text and photographs overnight and drafts an inspection-findings checklist from client-provided documents for operator review, with AI confidence indicators to support prioritization. Agents became the makers and operators became the checkers, reviewing and approving each morning. The system is governed by design: security controls and guardrails, a full audit trail, and a human-approved learning loop. Final decisions remain with people.

Early production results show handle time per inspection down 63%, from 136 to 50 minutes. Review accuracy rose from 80% to around 95%, and coverage grew from 4 to 29 risk types. Review capacity has been able to keep pace as the firm adds acquisitions.

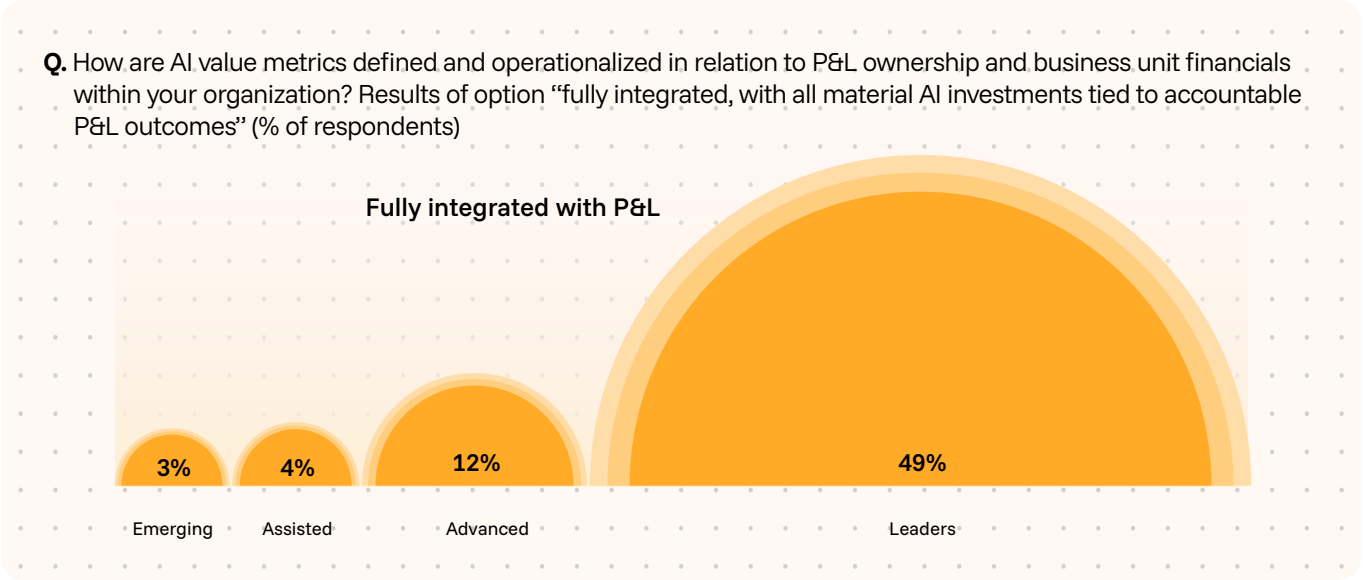
**The client provides the data, reviews all outputs, and makes all underwriting, pricing, and coverage decisions. Data used with consent and in accordance with applicable data protection requirements. Early production results are specific to this engagement. See the disclaimer section for more details.*

5. Value capture: most enterprises measure AI value. Leaders put it on the P&L

Figure 5 shows that only 15% of enterprises fully tie AI investment to specific P&L outcomes, and that most struggle to redeploy the capacity AI creates. Creating value and capturing it are not the same thing. An agent may cut processing time, add capacity, or improve a decision, but none of that becomes an economic outcome on its own. Without a named owner and a financial target, freed capacity has nowhere to go.

Leaders close that gap. We find that 49% have AI value metrics fully integrated with specific P&L lines, against 7% of the rest of the study. The step from Advanced to Leaders is the steepest on the chart: 12% of Advanced enterprises have made the link, against 49% of Leaders (Figure 14).

Figure 14. Leaders stand apart in connecting AI value to the P&L



Sample: 2,005 global enterprise executives
Source: Genpact, 2026

Among Leaders, value is much more often tied to the business unit, financial target, or P&L owner expected to capture it. If agents create capacity but no business owner is accountable for where it is redeployed, monetized, or reflected financially, value stays visible at the task level without ever becoming enterprise value. Leaders link AI-generated gains to accountable outcomes up front, rather than measuring ROI after deployment. As *Genpact's Global Agentic AI Officer* Vijay Vijayasankar puts it:



The future of enterprise AI won't be defined by how many agents an organization deploys. It will be defined by how effectively humans and agents operate as one workforce. The leaders will be those that treat AI as an operating model transformation, not a technology program, embedding accountability for value capture into every process they redesign."

Vijay Vijayasankar
Global Agentic AI Officer, Genpact

Agentic operations in practice: bringing all five building blocks together

Taken together, the five building blocks are a blueprint for running agentic operations. Enterprises do not maximize autonomy everywhere. They decide where agents should act, where people should decide or audit, how each hands work off to the other, what systems and controls make that possible, and what business outcome the redesigned work is meant to produce. The unit of transformation is the work itself, redesigned so that people and agents each do what they do best, toward outcomes the business can measure. Autonomy in practice is what is achieved when it works.

Case study

Genpact Client Zero: aligning autonomy with value

When Genpact set out to transform its own operations, leadership set one rule: every AI initiative gets measured, and its impact must show up in the company's own results. Value capture was designed in from the start, not measured in hindsight.

The lens was the work, not the technology. Functional heads reimagined their functions for a human-agent workforce, process by process, deciding what agents would execute and what humans would elevate to. Where freed capacity would go was defined before agents went live, not negotiated afterward. According to Genpact's President and CEO, Balkrishan "BK" Kalra:



We walk the path ourselves first. We call it Client Zero. We do not ask our clients to be the experiment. We bring them what we have learned."

Balkrishan "BK" Kalra
President and CEO, Genpact

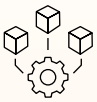
The dividend is landing in three ways:

Efficiency



Client Zero and other simplification efforts have generated tens of millions of dollars in savings that allow Genpact to make strategic investments back into the business. To date, 180 agents are in production across HR, finance, IT, and business operations. In finance, within one legal entity alone, around 40% of journals now flow through without manual intervention. As it scales more broadly, Client Zero is laying the groundwork to help Genpact operate 2X leaner in the future.

Business model



The operating disciplines and agents proven in Client Zero have become the blueprint for the agentic solutions Genpact now sells. Genpact AP Suite — the company's first productized agentic solution — anchors a portfolio that now spans solutions for accounts payable, record-to-report, source-to-pay, banking, and insurance.

Growth



Demand for these solutions is accelerating. As shared in Genpact's Q3 2026 earnings, agentic solutions bookings will exceed $1 billion in total contract value in 2026, more than five times 2025 levels.*

Client Zero puts the Leaders' playbook into practice: a CEO-level mandate, work redesigned ahead of deployment, governance built in, and value capture run as an operating discipline. This discipline continues as Client Zero expands.

See the section titled "Disclaimers" for information about forward-looking statements made by Genpact.

Make autonomy work in practice

AI has moved into the enterprise faster than the enterprise has moved to meet it. Among executives, 79% report agents live in production, and AI now leads or fully runs 40% of decisions across the stages we measured. Yet on average, only 49% of funded AI initiatives have delivered the value expected, and just 15% report having fully linked AI investment to P&L outcomes. The constraint is no longer AI capability alone. It is the ability to operate AI across the enterprise and capture its value.

Leaders show how this works in practice. They are the top tier of Genpact's AI Maturity Index, 17% of enterprises in our study, and their advantage is not simply that they deploy more AI. Backed by a CEO-level mandate, they connect the five building blocks of agentic operations as one system: redesigning work end to end, creating roles for people and agents, making data and AI capabilities reusable across workflows, governing agents while they run, and connecting AI-enabled gains to P&L outcomes.

What Leaders do differently

51% vs 17%

Redesign work across the enterprise

51% of Leaders pursue proactive, enterprise-wide operating-model transformation, compared with 17% of other enterprises.

51% vs 22%

Build roles for human-agent work

51% of Leaders create net-new roles specifically for human-agent operations, compared with 22% of other enterprises.

87% vs 54%

Give agents the access to act

87% of Leaders provide agents with timely read and write access across the systems their workflows require, compared with 54% of other enterprises.

47% vs 11%

Govern agents across the lifecycle

47% of Leaders have governance embedded across the full agent lifecycle, compared with 11% of other enterprises.

49% vs 7%

Connect AI value to the P&L

49% of Leaders fully integrate AI value metrics with specific P&L lines, compared with 7% of other enterprises.

Together, these differences point toward what autonomy in practice actually means. It is not the pursuit of maximum autonomy where agents run everything. It is the deliberate design of work around people and agents, with the right level of autonomy for the process, decision, and risk involved. Agentic operations is how enterprises make that possible using the five building blocks; autonomy in practice is what happens when it works at enterprise scale.

Where to go from **here**

The right next step depends on where an enterprise stands today. The five building blocks will not all advance at the same pace and there is no single implementation path. But the maturity patterns in our research point to a practical progression. The steps below are Genpact's point of view, informed by what Leaders do.

Emerging to Assisted: prove it

Name one executive owner and pick one process with heavy manual work and clear value. Embed agents with human oversight defined, baseline the process, and name a financial owner before go-live. Fix the data and integration points that block that first use case.

Assisted to Advanced: scale it

Define what agents own, what people own, and how exceptions escalate, then extend to more processes. Create the first human-agent roles, such as exception specialists, agent trainers, and output auditors. Give agents read and write access to the systems their workflows need, and move governance from deployment gates to controls that work while agents run. Set value metrics and a plan for where freed capacity goes.

Advanced to Leaders: run it as one system

Take the mandate to the CEO and the board. Align business and technology leaders to one enterprise operating agenda, with clear ownership of architecture, agent performance, risk, and value. Redesign across functional boundaries instead of function by function. Make data access, integration, identity, and monitoring reusable so each new workflow builds on what already exists. Enforce security and responsible AI across the full agent lifecycle, and tie AI value to P&L lines and the budget cycle.

At every stage, progress should be measured by what the enterprise can reliably operate, not by how much AI it has deployed. The current stage should be running in production, with human oversight defined, accountable ownership, and a measurable business outcome. Progress will rarely be even across all five building blocks. Enterprises should identify the building block that most constrains progress and strengthen it as they advance. Genpact's AI Maturity Index provides a way to understand that starting point.

Johan Koekemoer, *Director of Shared Services, Africa and Australia at AngloGold Ashanti*, describes the importance of working across functional boundaries:



We have an opportunity to change the business from a shared services perspective. We really would like to do that, but you can't do that in isolation. It's critical that you do that with the rest of the organization."

Johan Koekemoer
Director of Shared Services, Africa and Australia, AngloGold Ashanti

Operations become the **advantage**

The next phase of enterprise AI will not be defined by how much autonomy organizations can achieve, but by where autonomy creates the greatest value. Different work, decisions, and risks will require different combinations of people and agents. The goal is the right autonomy for the work and the outcome.

That changes the unit of transformation from the AI use case to the work itself. Enterprises that redesign work across functional boundaries, deliberately combine human judgment with agent execution, and connect the resulting capacity to business outcomes can turn AI capability into a new way of operating.

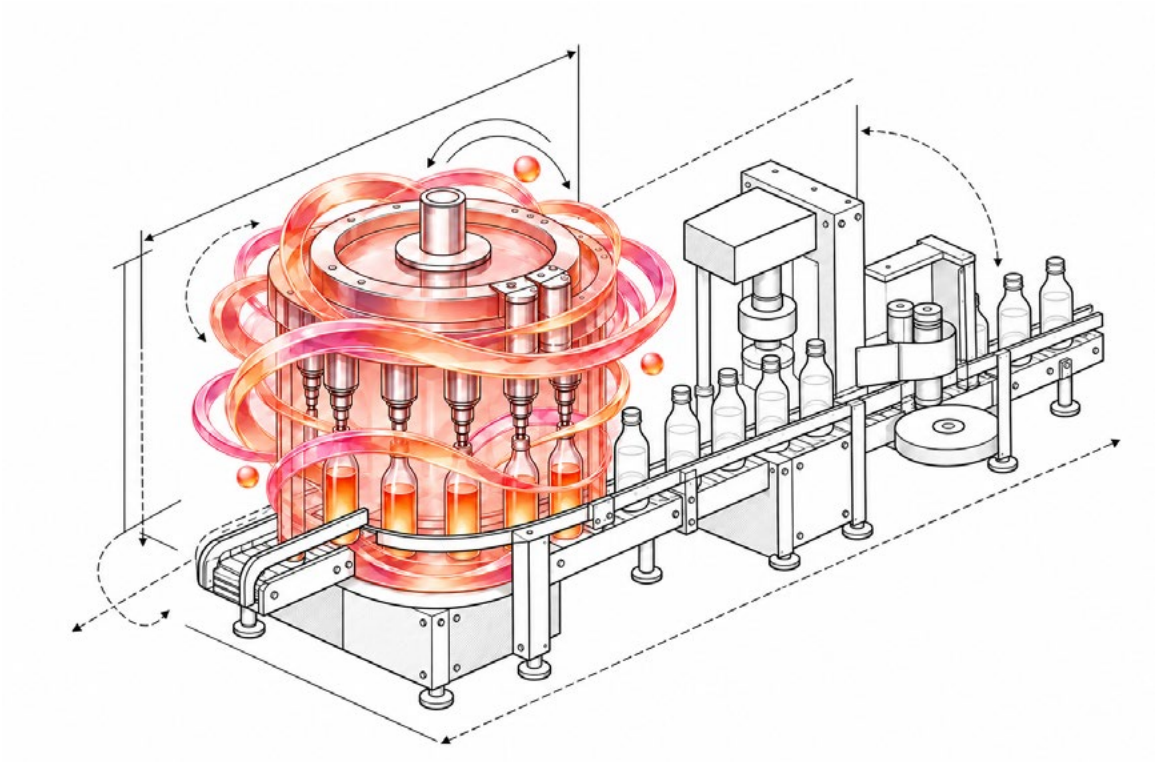
This elevates the human dimension of transformation. As Nelson P. Repenning, *Faculty Director of the MIT Leadership Center, and the School of Management Distinguished Professor,* puts it:

Deploying agentic AI creates potential. Agentic operations helps turn that potential into business value: autonomy designed around the work and accountable to business outcomes. That is autonomy in practice.



The companies that survive and thrive in the next decade will be those that understand that implementing artificial intelligence is a human problem."

Nelson P. Repenning
Faculty Director, MIT Leadership Center, and the School of Management Distinguished Professor



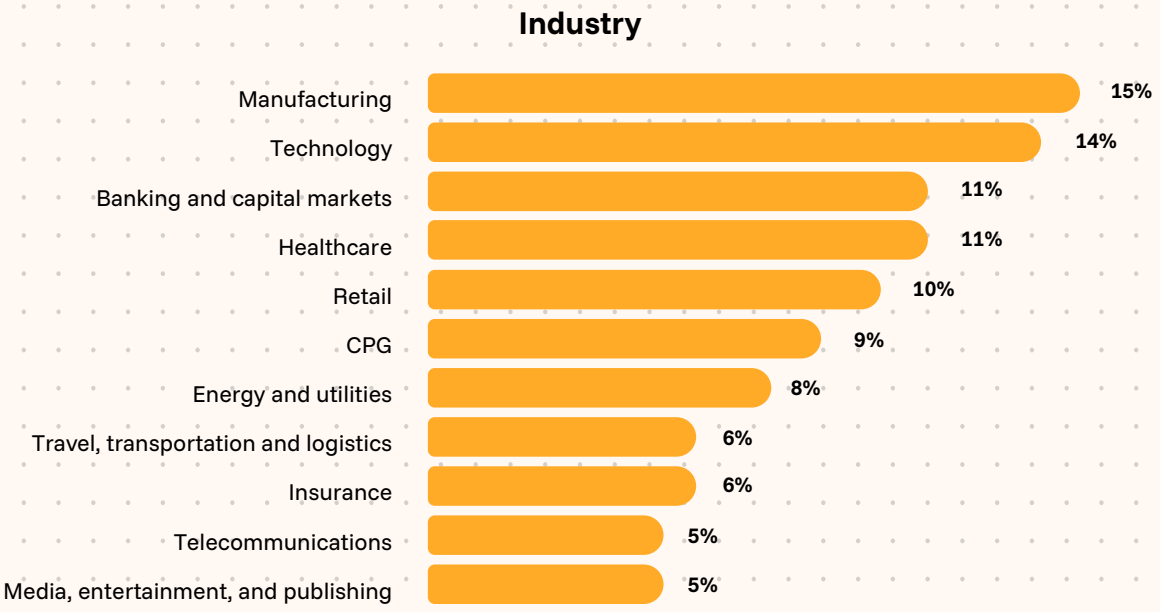
Research methodology

About the study

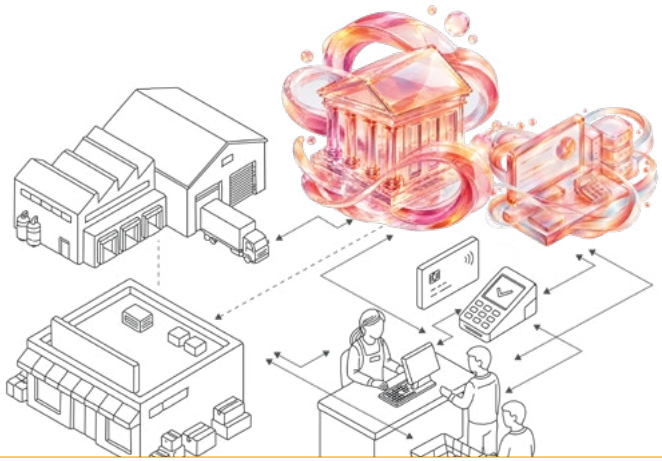
This report is based on primary research conducted by Genpact in 2026. The study surveyed 2,005 enterprise executives globally, spanning 17 industries and 14 primary business functions. Respondents were senior decision-makers across business and technology at companies with annual revenues from $1 billion to more than $50 billion. The quantitative research was supplemented by in-depth interviews with enterprise leaders and technology executives. Survey findings reflect self-reported responses from 2,005 enterprise executives and are presented for market context. Percentages are rounded and may not sum to 100%.

Respondent demographics

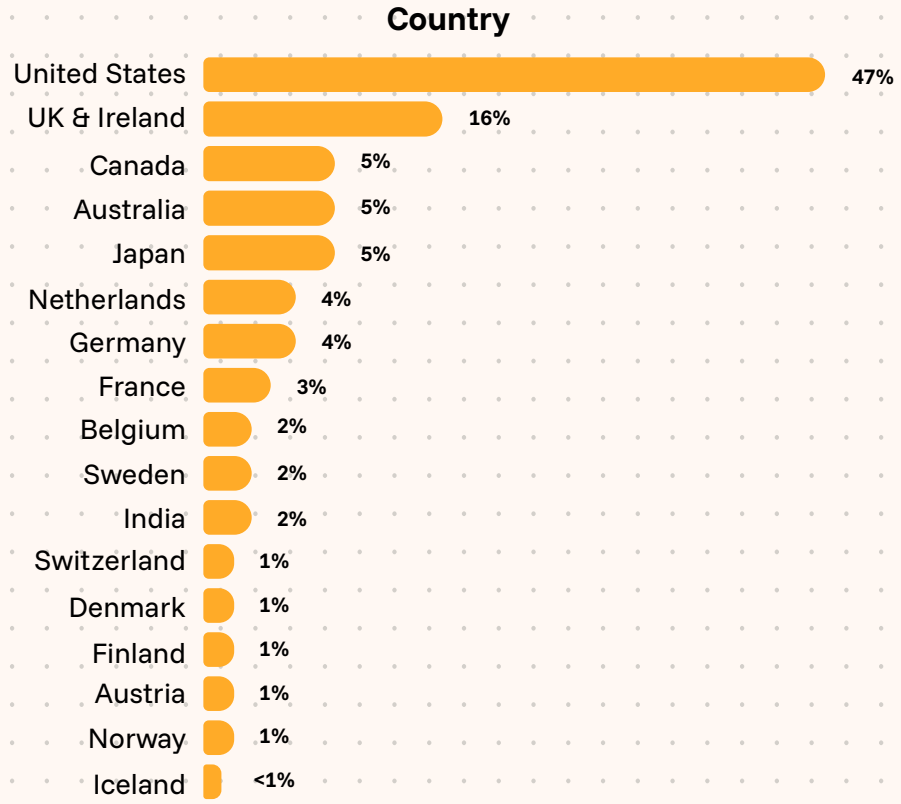
Q. Which of the following choices best describes your company's primary industry? (% of respondents)



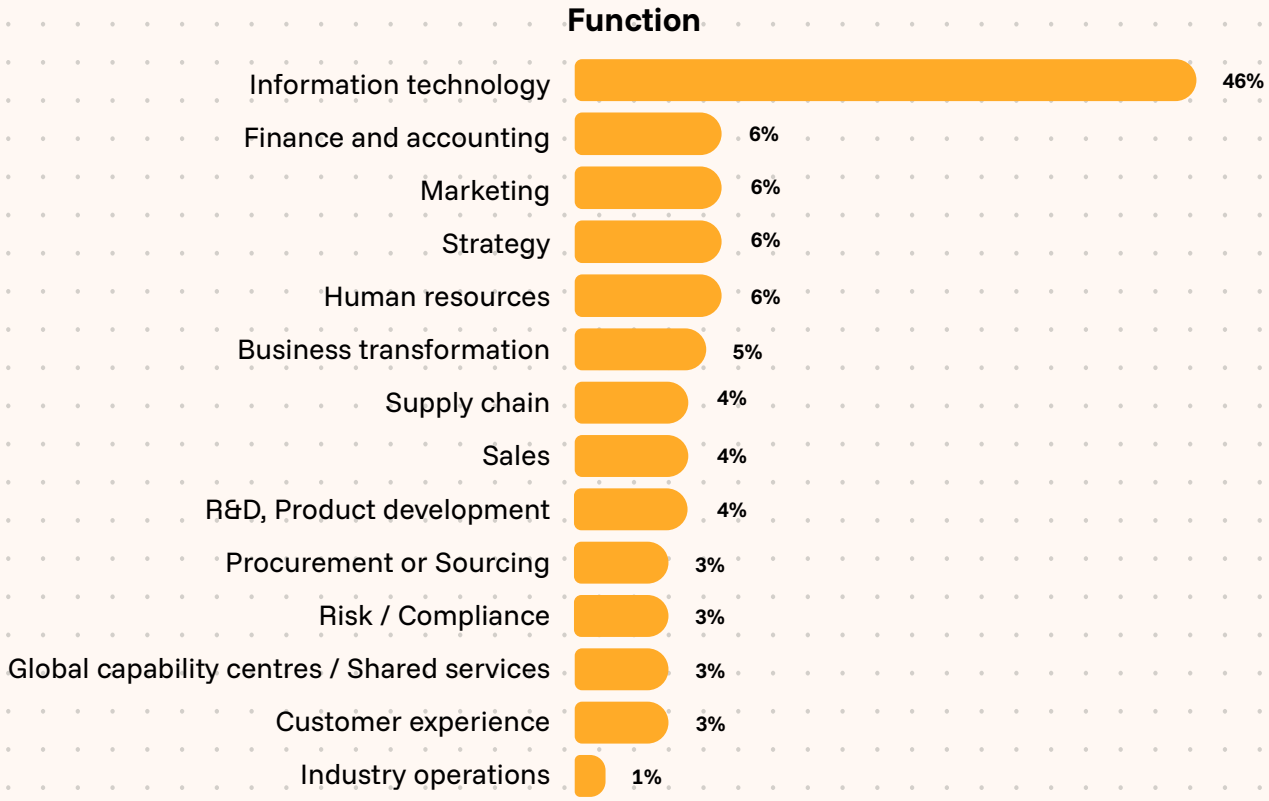
Sample: 2,005 global enterprise executives
Source: Genpact, 2026



Q. Which country are you based in?

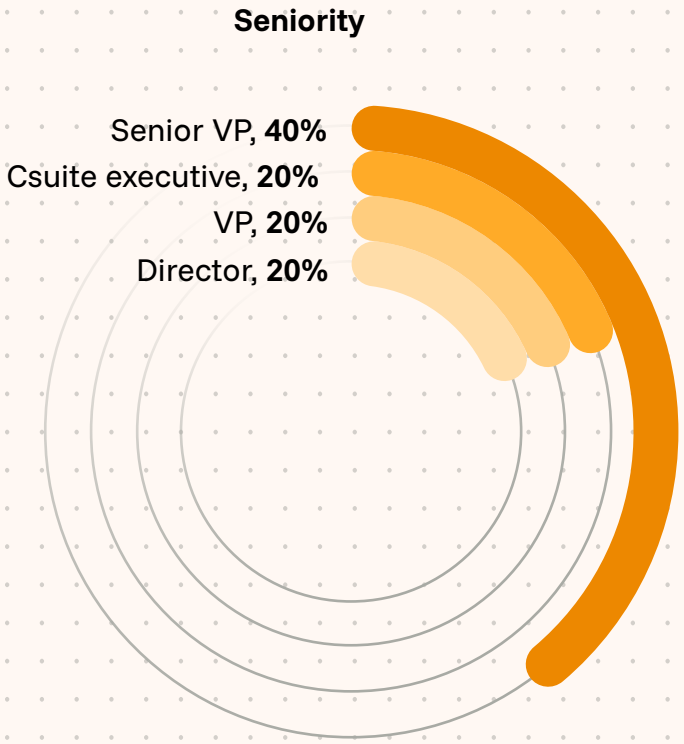


Q. Which of the following best describes your primary functional area within your organization?

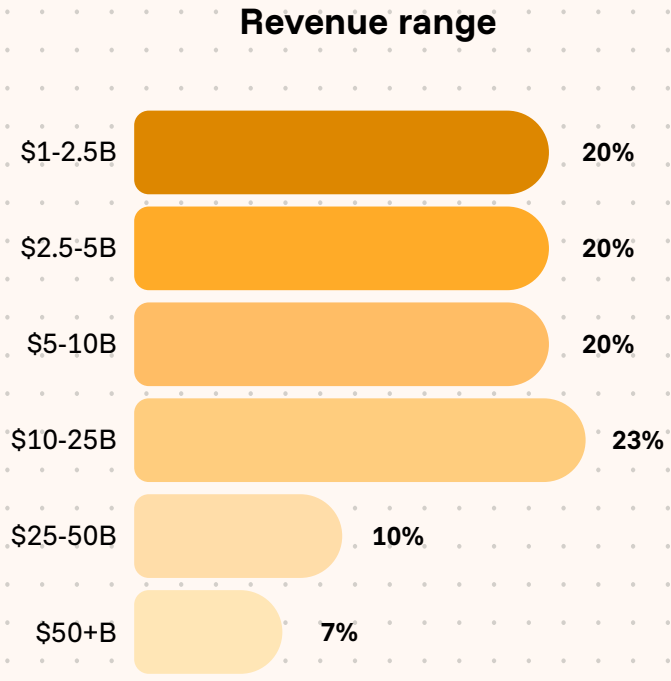


Sample: 2,005 global enterprise executives
Source: Genpact, 2026

Q. Which of the following best describes your seniority at your current role?



Q. Which of the following categories best describes your company's total annual revenue?



Sample: 2,005 global enterprise executives
Source: Genpact, 2026

How the **AI Maturity Index works**

Genpact's AI Maturity Index assesses how coherently enterprises are advancing AI across seven dimensions. For 2026, we retained the six dimensions used in the original index to preserve year-on-year comparability, while updating them to reflect the movement of AI from experimentation and deployment into enterprise operations.

Operating model transformation now captures how work, roles, decision rights, and processes are changing around AI. Advanced technical capability places greater emphasis on orchestration, agent-appropriate security, data foundations, and full-stack architecture. We also added value realization to assess whether AI-enabled gains are connected to measurable enterprise and P&L outcomes.

The seven dimensions of Genpact's 2026 AI Maturity Index

2026 maturity dimension	What it captures
Adoption effectiveness	How broadly and consistently AI applications are deployed across the organization.
Decision autonomy	The degree to which AI participates in, leads, or executes enterprise decisions.
Governance rigor	The ability to govern autonomous and agentic AI responsibly across the lifecycle.
Compliance readiness	The maturity of controls and practices required to meet emerging AI regulations.
Operating model transformation	The extent to which the enterprise has a deliberate strategy, operating changes, and investment to coordinate AI technologies and work across the business.
Advanced technical capability	Strength in agentic orchestration, agent-appropriate security, data foundations, full-stack architecture, and other capabilities needed to operate agents at scale.
Value realization	The discipline for measuring AI-generated value beyond productivity and linking it to enterprise and P&L outcomes.

Source: Genpact, 2026

AI maturity reflects how coherently an enterprise brings these dimensions together. An organization may be strong in adoption but weak in redesigning work, technically sophisticated but conservative in extending decision authority, or well governed but unable to translate AI-enabled gains into enterprise value. The index assesses the combined maturity of these capabilities rather than any one dimension in isolation.

Leaders treat autonomy as an enterprise-wide discipline. Agents run a meaningful share of processes, decision authority extends further, and operating model transformation, technical capability, lifecycle governance, and P&L-linked value disciplines work together. At the other end of the index, Emerging enterprises have limited or fragmented AI adoption, with foundational operating model, technology, governance, and value disciplines still at an early stage.

About the **authors**

Lead author



Elena Christopher

Global Head of Genpact Intelligence

Elena Christopher is Global Head of Genpact Intelligence, Genpact's research and thought leadership organization. She leads original research on AI, agentic operations, and enterprise transformation, translating data into provocative points of view for C-suite leaders. Elena brings more than three decades of technology and professional services experience, including senior research leadership roles at HFS Research and Gartner.

Contributors



Ajay Vasal
Global Lead for AI and Agentic Services

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



Jayant Swamy
Chief Enterprise Architect

Jayant Swamy is Chief Enterprise Architect at Genpact and the founding director of Genpact Labs, where he leads the innovation charter. His focus is to help our clients make effective use of AI and related technologies in their business transformation. He is regarded as a data and AI evangelist in the industry, building deep experience in enterprise architecture, helping clients shape their AI and data strategies, build new capabilities, and design complex transformation programs. He has worked across several industries and brings a point of view on the future of AI and how companies should adopt it.



Ojas Save
Technology Strategy Lead

Ojas leads strategic initiatives in the Office of Genpact's Chief Technology and Innovation Officer, where he turns technology strategy into decisions, capabilities, and market-facing impact. He works across Genpact's advanced technology portfolio — connecting strategy, offerings, and go-to-market with industry analysts and enterprise clients. With two decades of experience in the data and AI space, he has led marketing, corporate strategy, business planning, and client relationships for both global MNCs and startups, playing a key role in helping organizations through their growth phases. Ojas has an MBA from the Indian School of Business, Hyderabad.



Dr. Pramukh Nanjundaswamy Vasist
Chief Data Scientist

Dr. Pramukh drives AI strategy engagements with clients at Genpact, while also leading research and thought leadership initiatives. A Data and AI practitioner with over 18 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. His thought leadership spans AI-enabled enterprises, disinformation, responsible AI, the future of data-driven organizations, and the societal impact of digital technologies. He also teaches at leading business schools, including the Indian Institutes of Management.



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

Zainab leads AI strategy and programs in the Office of Genpact's Chief Technology and Innovation Officer, serving as a strategic partner to the CTIO and the company's senior leadership. Her focus is turning strategy into action: helping leaders align on priorities, make informed decisions, and deliver on initiatives across the business. She helps lead Genpact's primary research on enterprise AI adoption, including *Autonomy by Design*, the prequel to *Autonomy in Practice*. She also designs and builds AI solutions that help the business operate more effectively. Working across sales, finance, delivery, marketing, and advisory, she brings a practical, cross-functional view of what it takes for organizations to adopt AI and make it work.



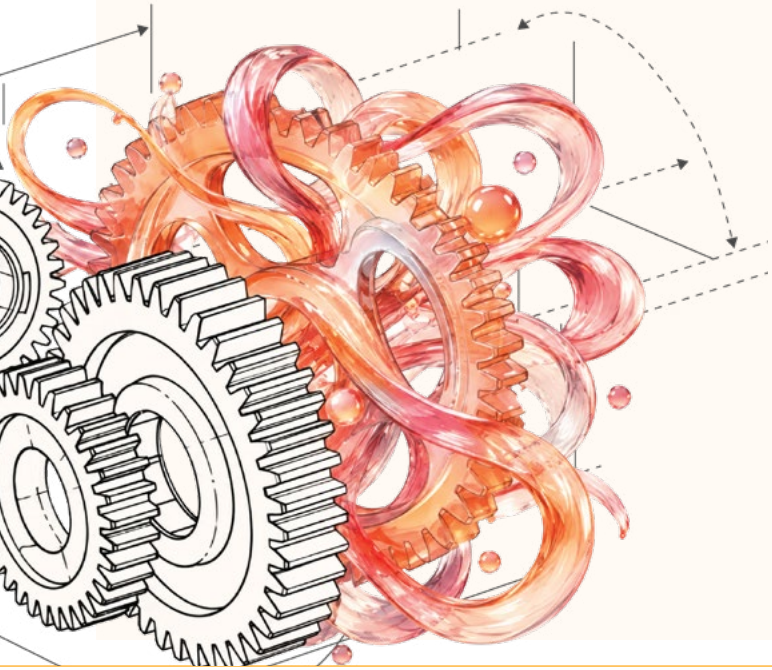
Aanchal Tyagi
Enterprise Narrative and Storytelling Lead

Aanchal leads enterprise narrative and strategic storytelling at Genpact, working across marketing, business, and leadership teams to translate complex technology and business themes into clear, compelling narratives through corporate messaging, executive thought leadership, and research-led content. With more than two decades of experience spanning television production, corporate communications, and executive storytelling, she brings a multidisciplinary perspective to building narratives that connect business strategy with audience priorities. Her work focuses on strengthening Genpact's voice and positioning across global audiences.



Michael Schneider
VP Communications

Michael leads strategic communications for Genpact's Office of the CEO. In this role, he drives the company's positioning in AI-powered business and agentic solutions, while elevating CEO thought leadership and strengthening stakeholder presence worldwide. He brings more than two decades of communications leadership to the role, with senior positions at Analog Devices, Fiserv, and ADP. He also serves on the board of directors at Bay Atlantic University and has taught at Rutgers and Fairleigh Dickinson University.



Disclaimers

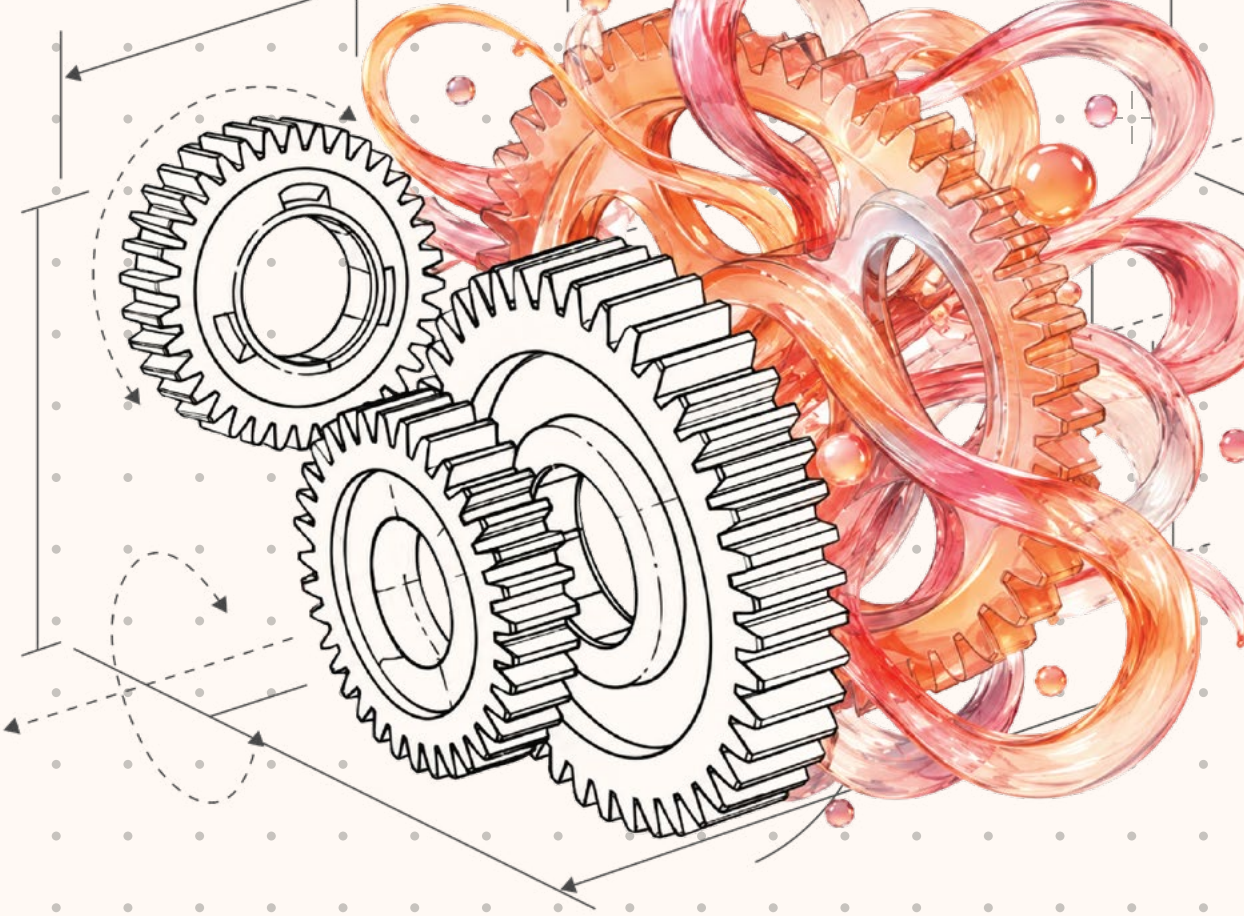
Client examples and results. Case studies describe specific client engagements. Results depend on client systems, process maturity, data quality, implementation scope, and configuration, and actual results may vary. Governance and control approaches described in client examples are illustrative.

Forward-looking statements. Statements regarding Genpact's ambitions, business plans, expectations, and anticipated revenue or bookings are forward-looking statements as defined in the safe harbor provisions of the U.S. Private Securities Litigation Reform Act of 1995 and involve a number of risks, uncertainties, and other factors that may cause our actual future results to be materially different from those expressed in our forward-looking statements. These uncertainties and other factors are discussed in more detail in our filings with the U.S. Securities and Exchange Commission, including our most recent annual report on Form 10-K, available at *www.sec.gov*. Although we believe that these forward-looking statements are based on reasonable assumptions, you are cautioned not to put undue reliance on these forward-looking statements. Genpact undertakes no obligation to update these statements. Historical figures are based on Genpact internal reporting.

Third-party names. Third-party names, trademarks, and quotations are the property of their respective owners, and their inclusion does not imply endorsement of Genpact or its offerings. *Find out more about Genpact's partner ecosystem here: *https://www.genpact.com/partners*

Insurance examples. Solutions described in insurance examples are intended for use in commercial and specialty insurance for business only and may not be used, in whole or in part, as a factor in determining a consumer's eligibility for insurance or for any other purpose set forth under the Fair Credit Reporting Act, 15 U.S.C. 1681 et seq.

AI disclaimer. Certain features use AI/ML models to support workflows. Outputs are subject to configuration, business rules, and human review.



genpact

Genpact (NYSE: G) is the Agentic Operations company, where frontier AI meets context-rich process intelligence. For decades, we have run mission-critical operations for the world's leading enterprises across finance, risk, supply chain, banking, insurance, and procurement. Today, with Agentic Operations, agents run the work, and our experts own the outcomes, so clients grow faster, operate leaner, and convert more to cash.

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