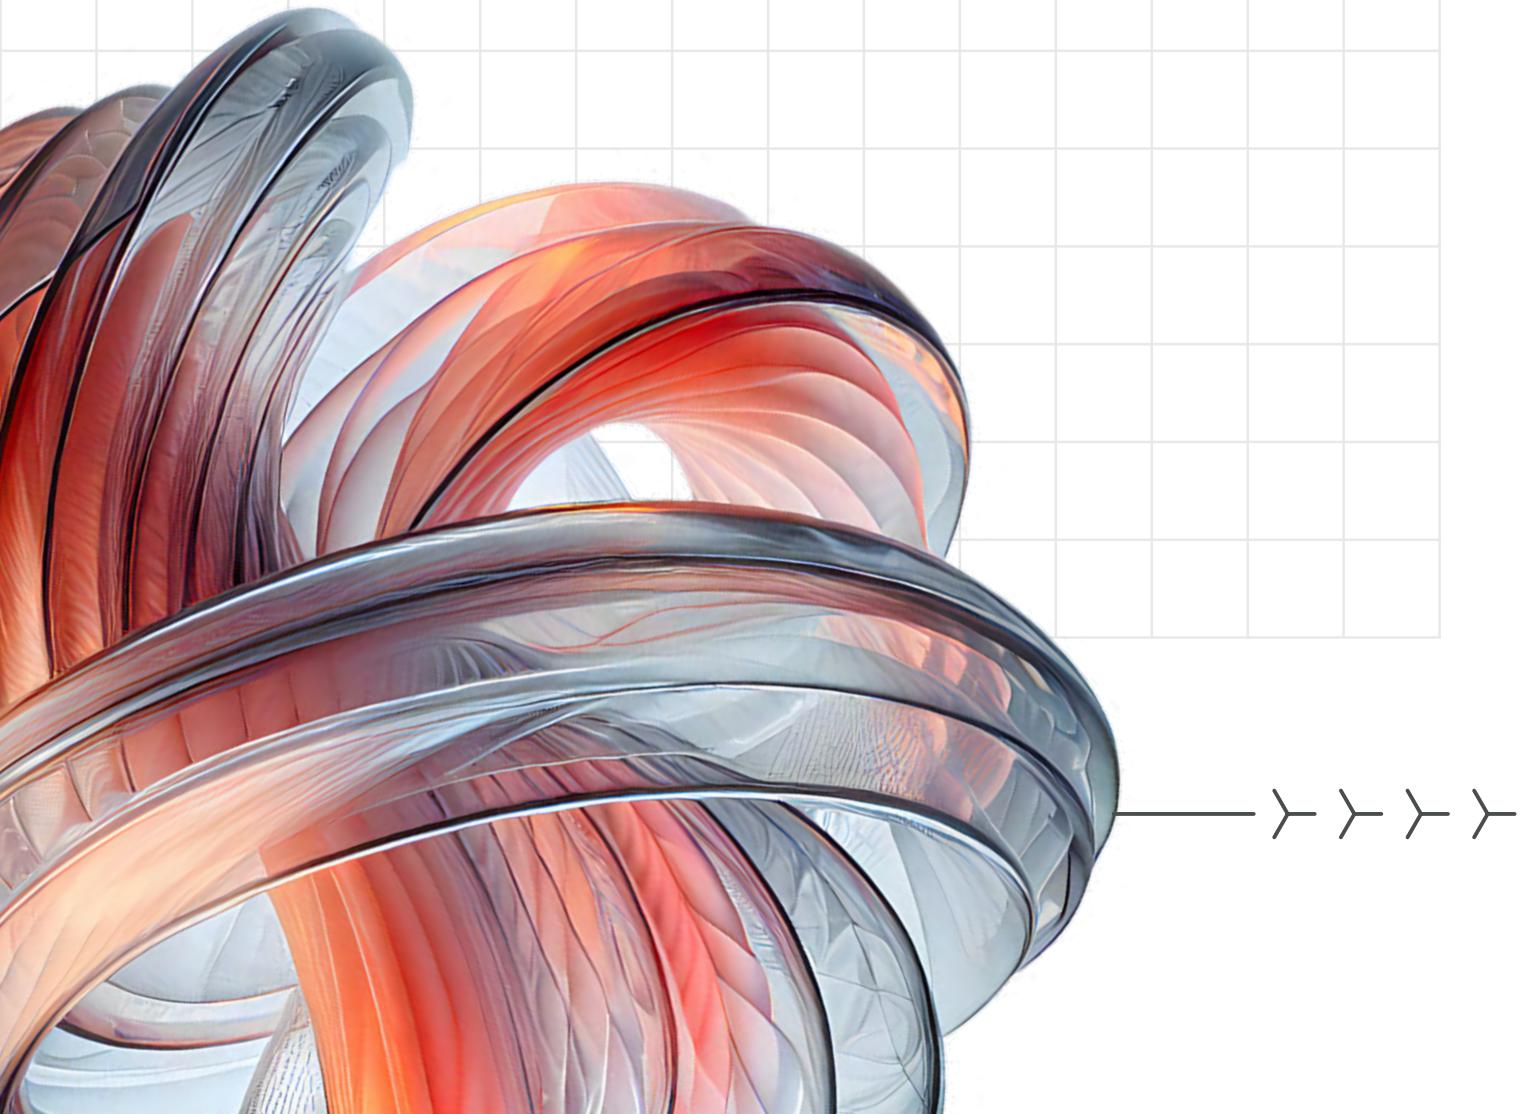




Autonomy by design

Scaling AI for enterprise value

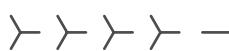


Foreword



Sanjeev Vohra

Chief Technology and
Innovation Officer,
Genpact



Artificial intelligence is at an inflection point. In the last couple of years, the conversation has evolved from how to adopt AI to how to embed it as a driver of enterprise-wide value – where systems decide, act, and learn alongside humans to unlock transformative outcomes.

At Genpact, we see this shift accelerating as companies integrate AI, including agentic AI, into their core operations' workflows, moving from siloed use cases to collective intelligence. They're making the shift to become autonomous enterprises, companies that have woven these technologies into their businesses, creating new growth opportunities and running agentic operations that make decisions in milliseconds and act with confidence.

Autonomy by design: Scaling AI for enterprise value explores that transition: its drivers, implications, and the questions leaders must now answer.

Our findings confirm that transformation goes beyond technology. It demands a parallel focus on process reinvention, data monetization, architecture, reskilling, and, more importantly, unlearning old ways of working. As our experience and this study show, there's no artificial intelligence without process intelligence. AI must operate as a connective tissue across the enterprise.

The enterprises at the forefront of this AI revolution demonstrate that autonomy isn't built through a single platform or algorithm – it's enabled by a blueprint for innovation.

I invite you to explore the insights that follow, and to imagine what your own autonomous enterprise could become.



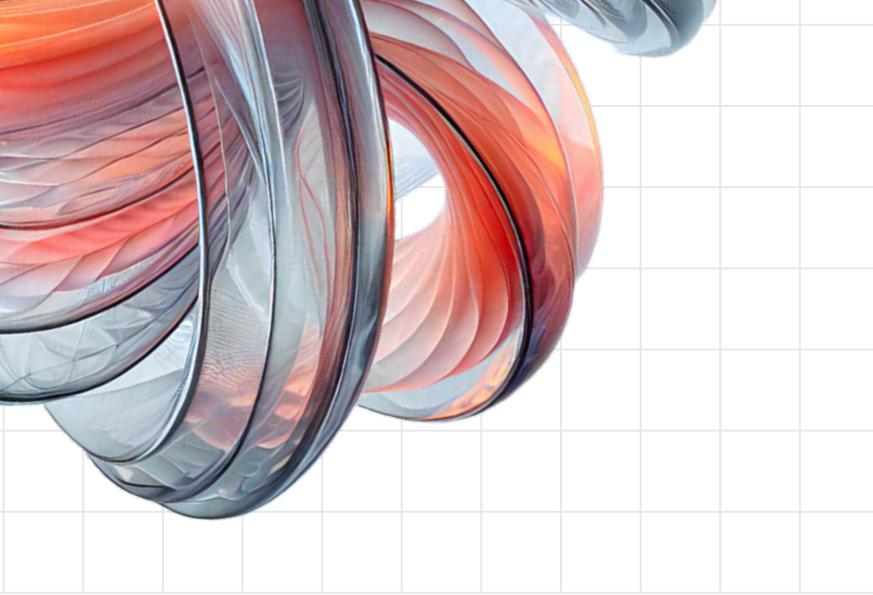


Table of contents

Executive summary	04
Introducing the levels of AI maturity	05
The AI of now	08
Value deficit	09
Technological complexity	11
Organizational inertia	12
The rise of the autonomous enterprise	15
A symphony of agents	17
The universal AI practitioner	19
Enterprise architecture redux	22
Governing at the speed of AI	25
The leaders' playbook: Translating vision into action	28
About this research	31





Executive summary

Across industries, organizations are under pressure to turn early AI gains into sustainable, enterprise-wide advantage. And with agentic AI, that future includes intelligent systems that act, decide, and learn alongside humans, redefining the limits of what people and companies can achieve.

For *Autonomy by design: Scaling AI for enterprise value* – Genpact's first annual report on enterprise AI – we surveyed 500 senior executives, including CEOs and C-suite leaders from enterprises with revenues of \$1 billion to more than \$50 billion across sectors and functions, supported by executive and industry expert interviews.

Barriers, blueprints, and playbooks

The study identifies four levels of AI maturity, enabling us to see how leaders are getting ahead and how the groups that follow can catch up. The barriers to progress with AI are real for all businesses, with some – like the skills gap and compliance – still challenging the most advanced companies. By seeking outcomes beyond productivity and overcoming technological complexity and organizational inertia, enterprises can build the foundations they need to fully embrace AI autonomy.

Taking that next step toward securing long-term advantage requires four interconnected enabling themes. Our research surfaces how leaders are harnessing AI's full value by:

1. Orchestrating a symphony of agents
2. Empowering the everyday AI practitioner
3. Reimagining enterprise architectures
4. Governing at the speed of AI

Collectively, these themes offer a strategic blueprint for translating AI ambition into measurable enterprise value. And they form the basis for our leaders' playbook (p28). Applicable across industries, the playbook offers practical guidance for any organization ready to transition into an autonomous enterprise.

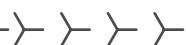


What is an autonomous enterprise?

“We believe organizations that are transitioning from ‘AI that generates’ to ‘AI that executes’ are on the path to becoming autonomous enterprises.

“To make the shift, humans and machines work together seamlessly and responsibly. The autonomous enterprise has AI **embedded across the business** as it connects workflows and orchestrates agents, making decisions in real time.”





Introducing the levels of AI maturity

AI's recent trajectory has been stratospheric. We see it through accurate medical diagnostics saving more lives, instant fraud detection keeping people's money safe, and supply chain agility deftly avoiding disruption to keep goods on shelves.

But the agentic AI opportunity opens new avenues that are not lost on our study's respondents. AI is stepping to the forefront, enabling systems to reason, decide, and execute autonomously.



Looking out three years, at least a quarter of senior executives expect self-managing business processes that run with minimal human oversight could become a reality. This may sound like an ambitious goal, especially considering that today, only 12% of respondents' organizations are clear leaders (figure 1). But these leaders are making AI autonomy an operational capability and hold important lessons that can help all businesses accelerate the move toward becoming autonomous enterprises.

What makes an AI leader

We assessed each enterprise across six dimensions: adoption effectiveness, decision autonomy, governance rigor, operational enablement, regulatory compliance, and technical capability (table 1). And it revealed four distinct levels based on AI maturity.

Leaders have both strategic foresight and executional rigor. While they may not yet reflect the full autonomous enterprise vision, they are enhancing processes and building architecture for scale. They have cohesive implementation strategies and are less burdened by the constraints of legacy infrastructure.

At least a quarter of senior executives expect self-managing business processes that run with minimal human oversight

Adoption effectiveness	Decision autonomy	Agentic governance rigor	Operational enablement	Emerging AI regulation readiness	Advanced technical capability
How broadly and successfully AI applications have been deployed across the organization	The degree to which AI systems participate in decision-making	The ability of current governance structures to manage autonomous and agentic AI responsibly	The investment in dedicated teams or systems to coordinate AI technologies across the business	The maturity of compliance practices as new AI regulations emerge	Engagement with next-generation AI technologies such as agentic orchestration and real-time strategic simulations

Table 1: How we benchmarked AI maturity



Advanced firms show broad adoption with strong technical foundations and responsible AI training, while assisted organizations are still building the basics. Emerging companies remain fragmented with limited adoption, focused on foundational elements like data platforms.

It's through this maturity lens that we assess the findings from our study and ask, can a self-managing organization be a reality by the end of the decade? We identify the routes that lead to scale and value, while addressing the structural and strategic potholes to avoid in the pursuit of an autonomous enterprise.

First, we examine the impediments that are holding companies back.

Benchmarking AI maturity identifies four clear levels

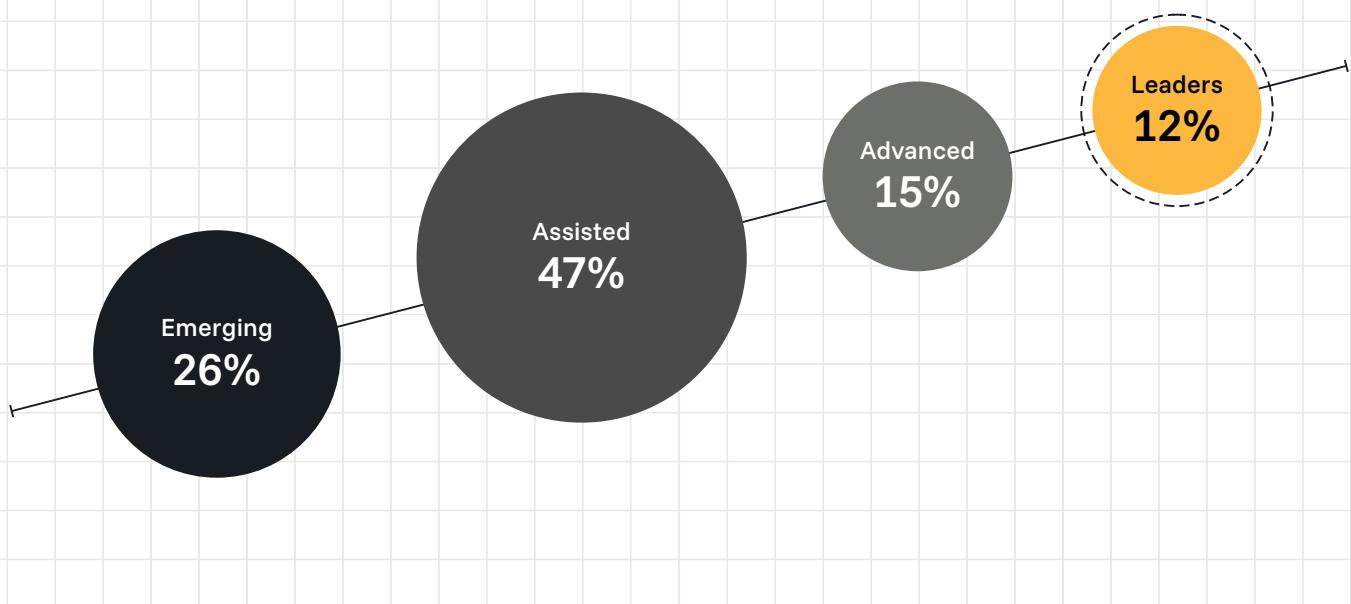
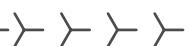
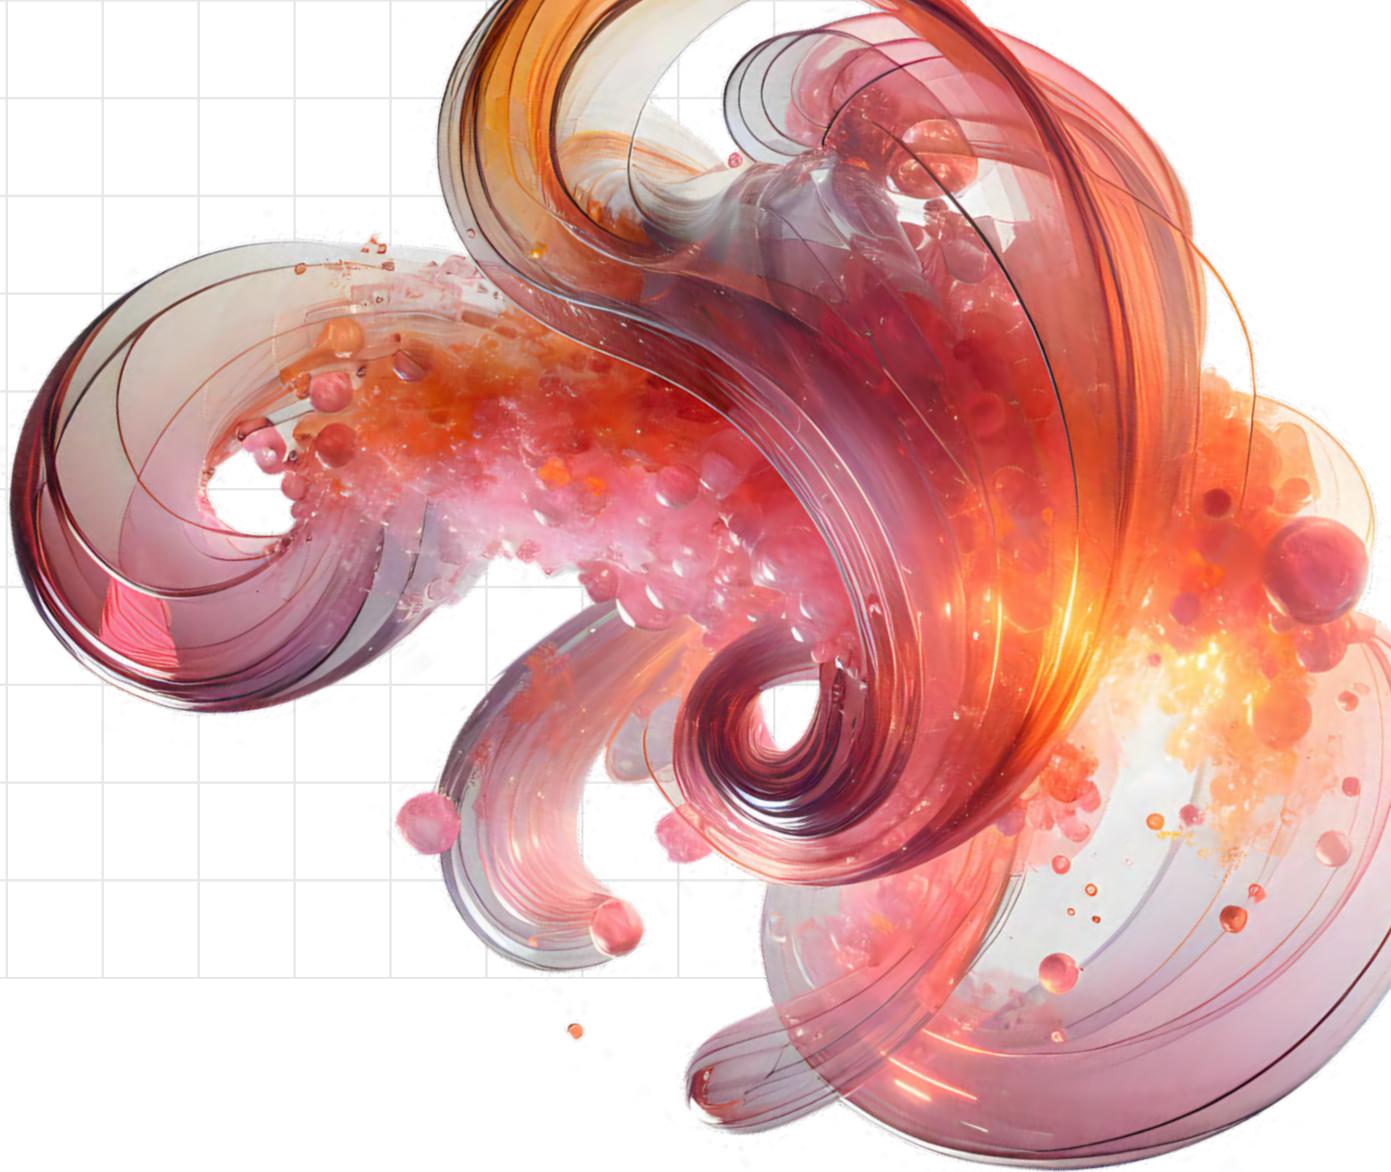


Figure 1: The levels of AI maturity

For allchart questions and N sizes, refer to p36





The AI of now

AI has evolved from generating insights to executing actions, paving the way for the autonomous enterprise. However, all organizations encounter the same hierarchy of obstacles – a value deficit, technological complexity, and organizational inertia – which underscores that progress toward AI autonomy is neither linear nor guaranteed.



Value deficit

The use of AI is now pervasive across enterprises, yet only 35% of responses indicate that select AI applications are very effective at delivering measurable business value. Our experience shows that translating AI investments into confirmed financial outcomes remains a significant challenge, underscoring the magnitude of progress still needed to realize tangible impact.

As a chief information security officer at a leading US insurance firm says, “Early commoditized tasks are ripe for the taking, but the ROI is lower because it’s easier to pick up slack there. The big bang for the buck is more complex, intricate processes involving multiple teams where AI has an opportunity to shine.”

Organizations are comfortable deploying AI in the intermediate decision-making stages – such as assisting humans by generating options and reallocating resources – where pattern recognition and optimization deliver efficiency gains. But when it comes to high-stakes, judgment-driven decisions, namely problem framing and final decision-making, executives remain cautious about putting AI fully in the driver’s seat (figure 2). Strategic decision-making continues to be people-led, reflecting a deep-rooted trust in human intuition and accountability.

This pattern underscores a deeper challenge: most organizations have yet to fundamentally reimagine how to operate in an AI-enabled world. The result? Progress is visible, but true financial impact remains elusive.



When it comes to high-stakes, judgment-driven decisions, executives remain cautious about putting AI fully in the driver’s seat



Strategic decision-making still sits with humans

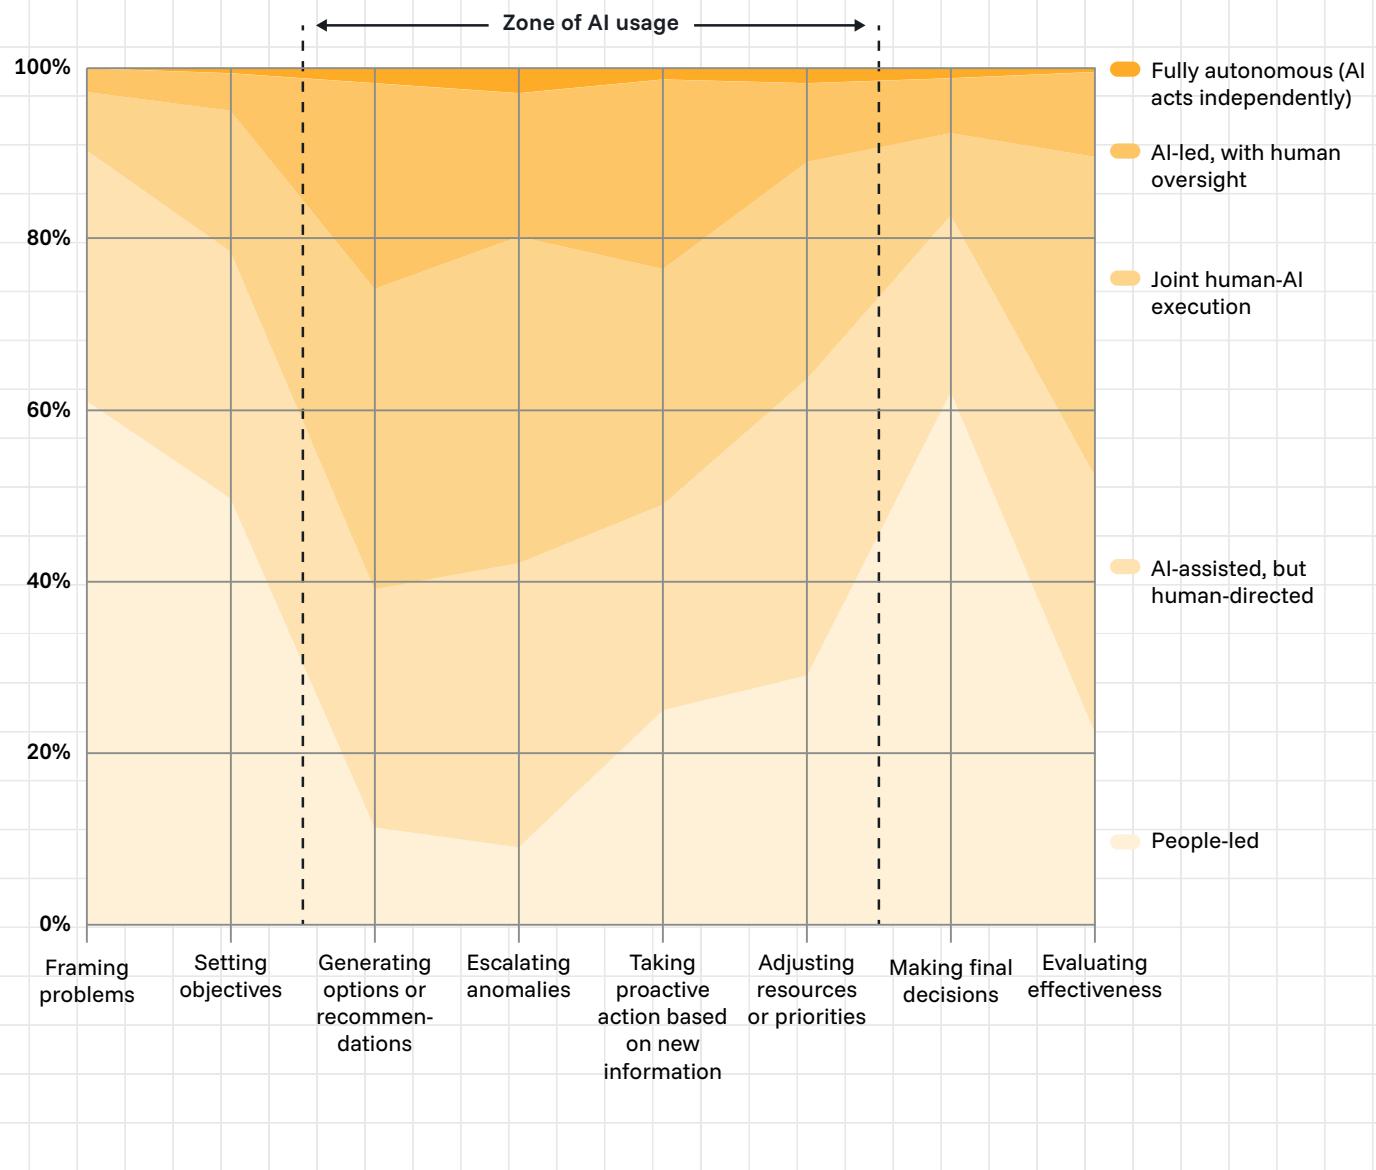


Figure 2: Human and AI responsibility across the decision-making process



Technological complexity

The explosion of technology – especially with the rise of agentic AI – has intensified architectural complexity, shifting focus toward integrating agents seamlessly into core workflows.

Companies are inundated with offerings from AI startups promising solutions to major business challenges. But reviewing their security, compliance, and how they fit with business systems takes resources and time and increases complexity.

And barriers do not disappear once companies select solutions. Among non-leaders (88% of respondents), the top technology-related challenge is difficulty integrating AI into existing workflows, cited by 32% of respondents. Technology limitations follow closely at 31%, while only 11% point to data quality as a major pain point – and leaders report similar patterns (figure 3).

Integration is the bigger barrier

Although poor data quality is often identified as a major barrier in industry discussions, our study shows integration is a significantly bigger challenge. Enterprises with process maturity can quickly integrate AI into their workflows. “AI doesn’t fail because the models are weak; it fails because integration is,” says Anu Dixit, global chief customer officer at Resolution Life, a global life insurance group. “Fragmented data and immature enterprise wiring prevent AI from scaling.”

Asha Poulose, chief digital officer at GE Healthcare, a leading global healthcare solutions provider, adds that getting both data and processes right is also vital for the future: “The value you gain from agentic AI will be determined by the process maturity and quality of the data that goes into it.”

Legacy integration and fragmented data ecosystems block companies from realizing the full benefits of new AI products and platforms. When integration fails, AI cannot scale.

Ultimately, much of a company’s technological complexity stems from legacy processes that were never designed for AI-driven operating models. Scaling AI demands that businesses rethink their technology stacks with process intelligence.

The top technology-related challenge is difficulty integrating AI into workflows



Integrating AI into processes a bigger challenge than data

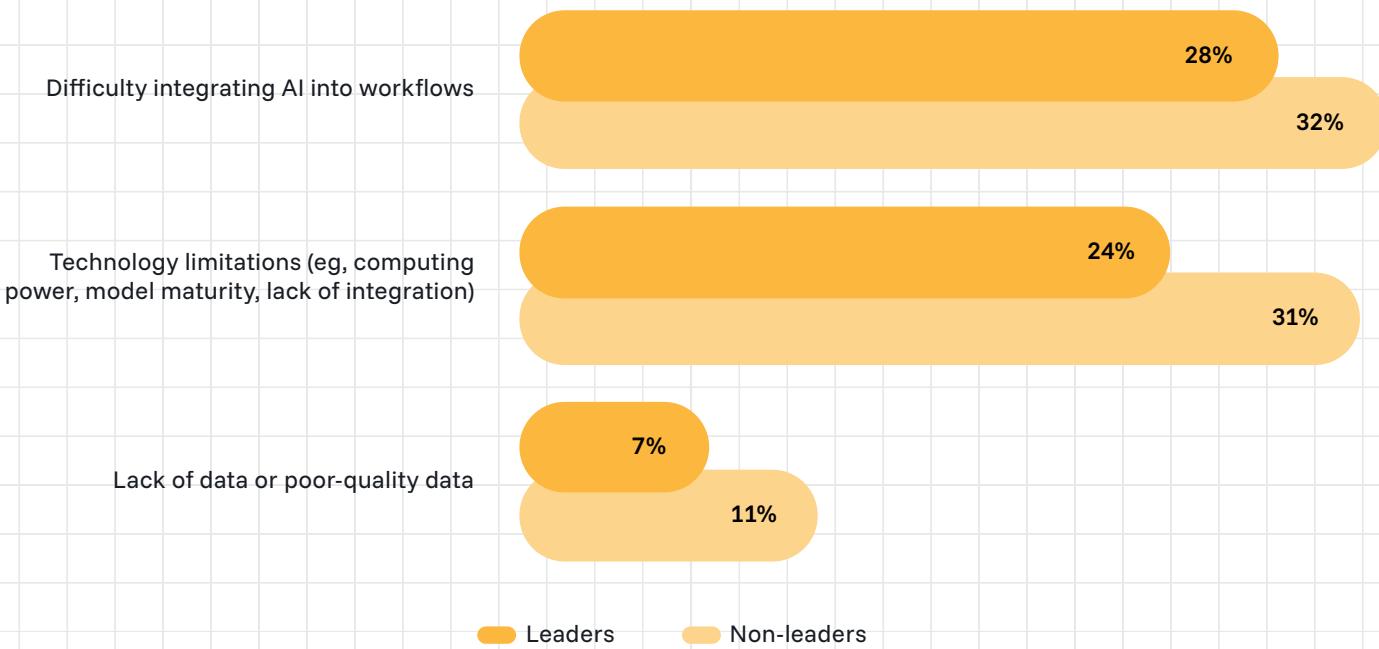


Figure 3: Significant AI implementation pain points – technology complexity

Organizational inertia

Misaligned structures and slow decision cycles are consistent barriers to scaling AI. These problems mainly stem from a workforce that lacks clarity and the capability to collaborate effectively with AI.

Workforce capability gaps continue to be the most frequently cited organizational constraint to AI adoption, as reported by 6 in 10 executives – yet only 45% say their organizations offer AI training for all employees. In tandem, 40% identify fragmented ownership and accountability as key challenges. While leaders have done more to overcome these barriers, they've not eliminated them yet (figure 4).



Build skills, set ownership, and manage change to overcome inertia

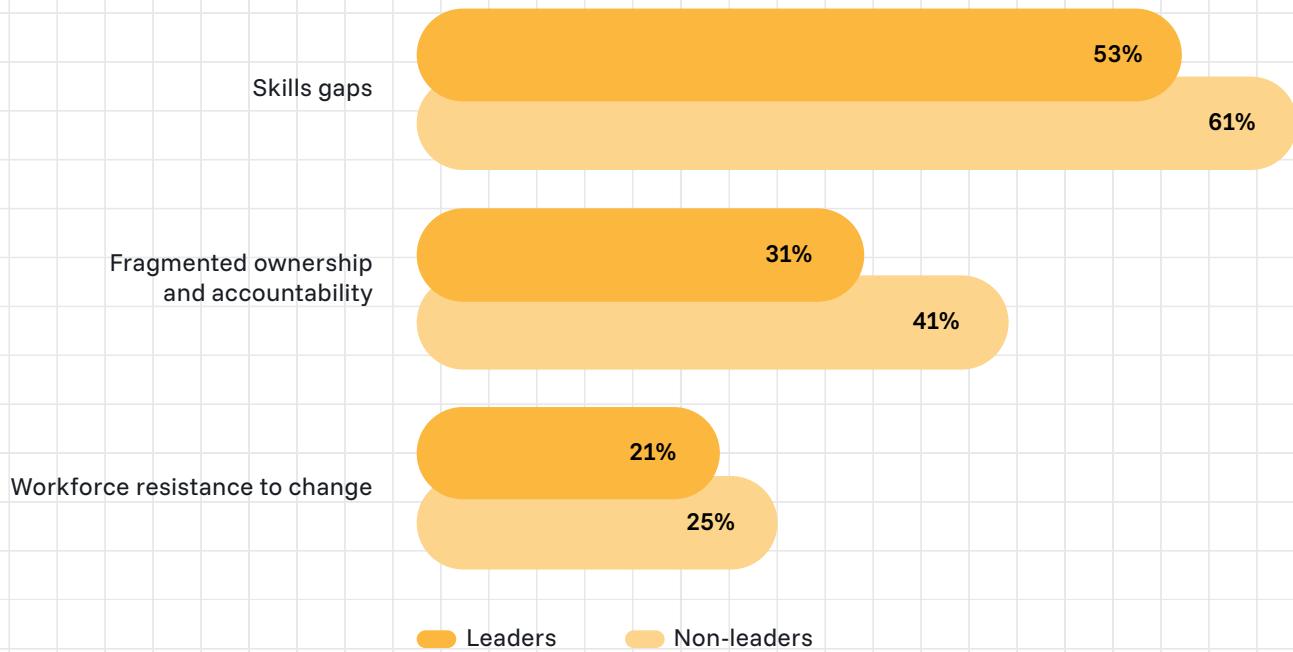


Figure 4: Significant AI implementation pain points – organizational inertia

“The true value of implementing AI solutions lies not solely in the technology itself, but in how it enables us to fundamentally redefine our operating model.”

Francesco Tinto, Global CIO, Advantage Solutions



Where leaders excel is in how they manage change. Only 36% of non-leaders have a structured change management framework with a dedicated team in place (versus 55% of leaders). While resistance to change appears low among barriers, without institutionalizing change management as a core capability and reconfiguring operating models to AI at scale, companies will remain at a strategic crossroads – stuck between ambition and execution.

As Francesco Tinto, global CIO at Advantage Solutions, a leading omnichannel retail solutions agency in North America, says, “The true value of implementing AI solutions lies not solely in the technology itself, but in how it enables us to fundamentally redefine our operating model.”

Advantage Solutions is adopting AI at two levels: “First, we focus on driving enterprise value by developing scalable solutions that optimize company-wide processes and workflows. Second, at the employee level, we empower our team members by adopting AI tools such as Copilot to enhance personal productivity.”

People at the core

Organizational inertia arises when processes are not designed around people. In the rush to embrace AI as the future, enterprises risk overlooking the human foundations required to sustain that journey. “AI poses an interesting paradox,” says Nelson Repenning, Faculty Director of the MIT Leadership Center, and the School of Management Distinguished Professor of System Dynamics and Organization Studies at the MIT Sloan School of Management.

“On the one hand, its main benefit will come from automating the work that humans currently do. On the other, however, you can’t automate the process of implementing automation. Humans must still make the tough, messy decisions about where and how to apply AI.”

Without leadership alignment, skill development, and employee cooperation, progress will continue to stall, regardless of technological advancements.



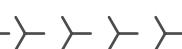
Only
36%
of non-leaders

vs

55%
of leaders

follow a
structured change
management
framework as they
adapt to AI





The rise of the autonomous enterprise

Even though organizations continue to face obstacles when scaling AI, leaders are making real progress. They have strategic approaches that all businesses can learn from in pursuit of an autonomous enterprise.



Our research reveals four enabling themes that differentiate those who are advancing toward enterprise autonomy from those still setting their foundations. While not definitive truths, these patterns offer a compelling framework for understanding how early leaders are operationalizing AI at scale.

Taken together, these themes represent a working model of autonomy in practice among the 12% of respondents currently at the frontier.



1

A symphony of agents

Orchestrating multi-agent systems integrated into business processes



2

The universal AI practitioner

Democratizing AI fluency and adoption across roles and functions



3

Enterprise architecture redux

Building data-centric foundations capable of supporting AI as it scales



4

Governing at the speed of AI

Establishing the oversight and mechanisms required to scale AI quickly and responsibly





A symphony of agents

There's no artificial intelligence without process intelligence – it's central to every successful AI initiative. Many enterprises are deploying dozens, even hundreds, of specialized AI agents. While each is designed to perform a specific task, they risk working at cross-purposes without coordination across business functions and processes.

What's missing is a unifying layer: a "conductor" to orchestrate their efforts. This orchestration is essential not only for addressing isolated challenges through point solutions but also for embedding agentic AI holistically across end-to-end processes.

"How are agents coming together for impact?" asks Jinsook Han, global agentic AI officer at Genpact. "How can you orchestrate agents, people, data, AI, and technology, while also measuring performance? This is why orchestration platforms are so important, giving companies the ability to manage them in one place."

Agentic orchestration is nascent

Only 3% of organizations – and 10% of leaders – are actively implementing agentic orchestration. This limited adoption signals that orchestration is still an emerging discipline. Among those leaders, we see respondents from high-tech, banking, and insurance represented most heavily (figure 5).

The scarcity of orchestration is a litmus test for both internal capability and external strategic positioning. Successful orchestration requires integrating AI into workflows, systems, and decision loops with precision and accountability. Those who have achieved AI agent harmony are building intelligent systems that scale.

Only
3%
of
organizations
are actively
implementing
agentic
orchestration





Case study

Scout – a cohesive, connected family of agents

When a new hire joins Genpact, Scout is on hand to smooth their onboarding so they can quickly navigate processes, find who they need to speak with, and get up to speed. When a sales executive is preparing for a client presentation, Scout's there to find the information that will make them shine.

Scout is Genpact's family of AI agents that work together to support all employees across IT, HR, finance, sourcing, and other functions. They're conversational and collaborative.

Using natural language understanding, Scout's master agent interprets a query's intent, context, and type. It breaks down complex questions into actionable components and can retrieve relevant information from previous interactions

or knowledge bases. Once connected to the best internal agents, functional agents resolve queries.

How do they work together? The system's hierarchical multi-agent ecosystem enables Scout to interact with the master and functional agents. Looking ahead, we're introducing a fully orchestrated agentic ecosystem – a digital fabric.

"This is where AI agents collaborate seamlessly across functions, platforms, and data, so our employees can interact naturally with agents across channels," says Vidya Rao, chief information and transformation officer at Genpact.

"We're building an enterprise nervous system where Scout has multi-agent orchestration behind it delivering intelligence, productivity, and compliance at scale."

Even leaders are yet to fully embed agentic orchestration across their businesses

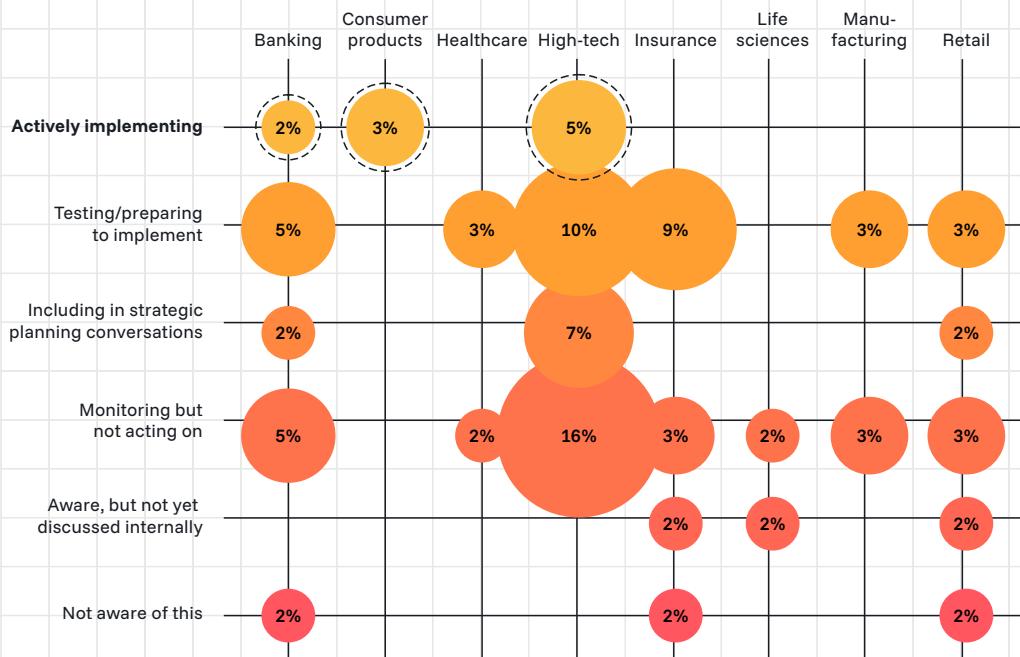


Figure 5: Leaders' adoption of agentic orchestration by industry





The universal AI practitioner

The shift to autonomy requires every employee to operate as an AI practitioner. Empowering employees starts with rethinking how people and AI systems collaborate within the business.

As more companies invest in AI training for all employees, leaders anticipate a more customized future. Within 3 years, 59% of leaders expect AI-enabled personalized training and support programs – dynamically tailored to skills, roles, and performance – to be a reality. However, some organizations are already there: “Our training programs are now customized for each person’s role thanks to AI insights,” says a senior executive from a life sciences firm.

Engaging employees

In addition to training, an organization’s ability to scale AI hinges on how effectively it aligns with, rather than disrupts, its workforce, especially as it introduces new roles, ways of working, and technology. As Tinto says, “The line between IT and non-IT roles is fading. Fewer finance staff will handle routine tasks like journal entries, while strategic thinking and analytical skills will become more important. As finance leaders create their own AI agents, the required skill set will continue to change.”

Change management continues to be critical to successful AI adoption. While resistance among the workforce may not be widespread (24% across all respondents), its prevalence is enough to slow progress if left unchecked. Only careful, coordinated change management will help employees get comfortable with the evolution, a practice seen most clearly among 55% of leaders (figure 6).

“Successfully implementing new technology – even one that automates human work – is ultimately about getting humans to behave differently,” says Repenning. “Put simply, implementing artificial intelligence requires a heavy dose of real intelligence.”

By reframing how humans and AI collaborate, companies are laying the foundation for enterprise-wide AI integration, supported by structured training and a culture of AI-powered augmentation. Together, these elements reduce hesitation, build confidence, and position AI adoption as an aspirational growth opportunity that translates to measurable impact.

59%

of leaders expect AI-enabled personalized training and support programs to be a reality within 3 years



Leaders are building a workforce of AI practitioners

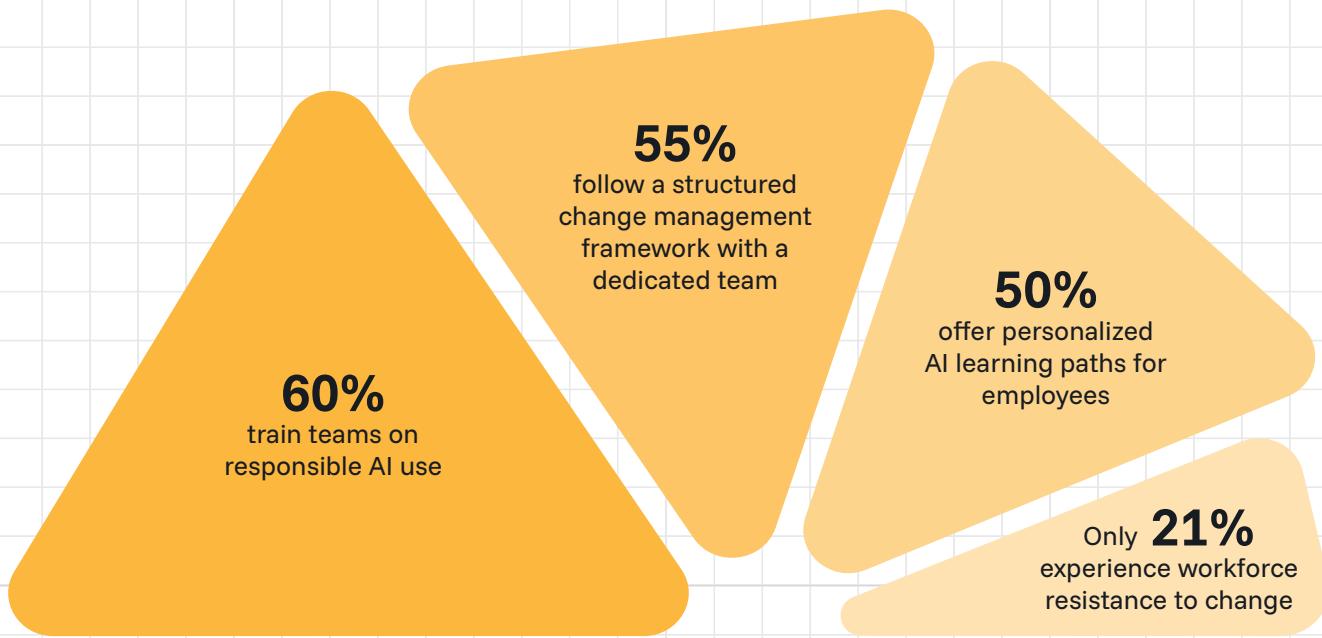


Figure 6: How leaders are preparing their workforces for AI adoption





Case study

Building an AI-fluent workforce with AI at scale

From entry-level joiners to the global leadership team, Genpact is building an AI-fluent enterprise. “Our goal is for every employee to become more fluent in AI – not just by learning about it, but by using it in the flow of their everyday work,” says Piyush Mehta, chief HR officer at Genpact.

Everyone has access to a role-based AI learning platform that defines and delivers the skills leaders, builders, and practitioners need. This includes vibe-coding parties and Genome Shorts – social-style, 90-second videos from Genpact leaders on how they use AI every day to inspire users and engage different learning behaviors.

But the real difference lies in how AI augments the program to refresh learning at speed.

Digital twin

To give learners greater access to subject-matter experts' expertise, Genpact built a digital twin – the AI Guru – using a corpus of 20,000 topics built on our specialists' knowledge. AI Guru has now answered almost 1 million learner questions with 96% positive feedback. Conversational AI bots provide users with contextual answers, role-based simulations, nano-training recommendations, and personalized assessments, all based on our experts' knowledge.

AI-based simulations

Take a client account leader preparing for a conversation with a bank's chief data officer. With the Data and AI School's voice-based simulation, an AI coach makes recommendations not only on the presentation's content but also on the account leader's active listening and storytelling, and how well they articulate the solutions.

The return on learning

The result is capability that shows up in the business. “Learning is translating into growth,” says Mehta. “As our people become AI practitioners, confident working with advanced technologies, we've filled nearly 60% of our open advanced tech roles with internal talent.” Through the Data and AI School, Genpact has increased its advanced tech sales conversation quality by 56%, which links to a 9% increase in deal inflows.

In 2024 alone, employees logged 11 million learning hours, averaging 82 hours per person. And more learning means people stay longer. Genpact has five times greater retention among people who learn continuously and engage with our training programs. Combining AI training with AI capabilities is creating value for clients, the company, and people's careers.





Enterprise architecture redux

Despite years of investment, achieving the transformative results AI promises is not straightforward. A core challenge lies in a company's enterprise architecture. Many organizations still rely on application-centric designs when what they need is a data-centric foundation capable of supporting scale.

Business leaders know that they need a new approach, and technology professionals and enterprise architects feel it most. When it comes to scaling AI, 61% say that the complexity of their technology architecture is a major or moderate challenge (figure 7).¹ Yet many are overcoming this to progress toward a data-centric approach.

This progress is evident in how leaders from our study are deploying enterprise platforms to scale AI systems. For example, 86% of leaders have centralized mechanisms for sharing and reusing AI tools, prompts, and use cases. This shows proof of investment and helps set the foundations for successful AI adoption, but companies need to go further.

Only
25%
of leaders have
fully adopted a
real-time data
infrastructure

Technical complexity is holding companies back

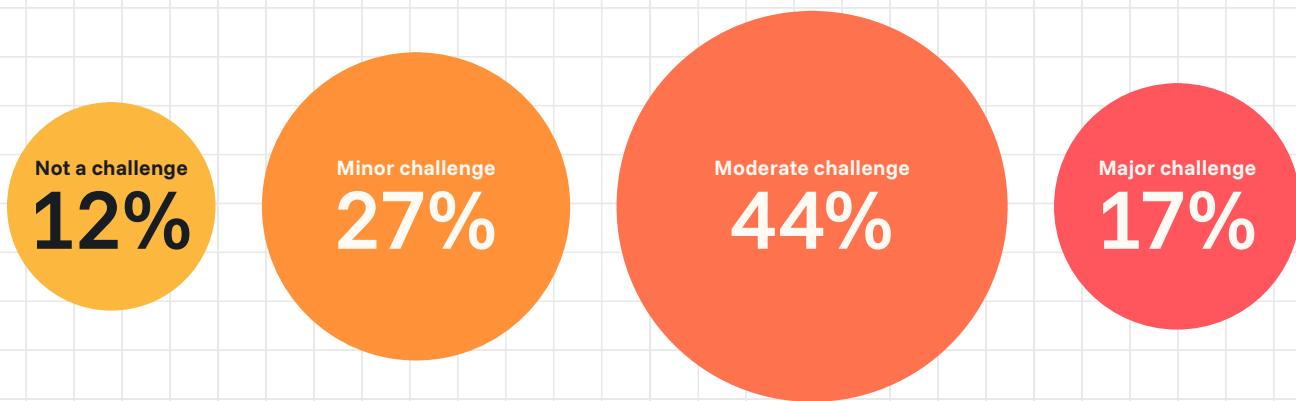


Figure 7: The extent to which the complexity of technology architectures is a challenge to scaling AI¹

¹Poll conducted at Microsoft Ignite, November 2025. Data captured at the Genpact booth from 52 technology executives from companies with revenues of over \$1 billion



True scale requires real-time responsiveness, a capability that is still limited. For instance, only 25% of leaders have fully adopted a real-time data infrastructure, enabling them to react and make decisions dynamically (figure 8). The next frontier remains largely unexplored, with just 3% of companies having active agentic orchestration, a sign of deep architectural maturity and confidence in their AI capabilities.

This limited orchestration underscores a critical truth: while investments are moving in the right direction, the architectural backbone required for enterprise-wide transformation is still emerging, and without it, orchestration cannot scale beyond isolated pilots. This imbalance highlights why the pivot to data-centric architecture is now essential to how organizations design, govern, and operate their systems.

As Ben Haklai, national technology officer at Microsoft Israel, says, “In this coming agentic wave, your ability to derive significant value from AI technology will depend on the baseline of your architecture on the data side and your organization’s ability to govern and manage these new agentic entities.”



Case study

Rewiring the enterprise for data at scale

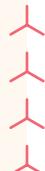
A global energy company was struggling to generate AI-driven, real-time insights at scale thanks to its aging technology, siloed data platforms, and fragmented analytics. Its outdated tech stacks were also increasingly expensive to maintain. The impact? Delayed business decisions and limited value from AI.

To break free from its legacy shackles, the firm built an AI-ready data foundation. It introduced a unified enterprise data intelligence platform to collapse silos and established a consistent governance and metadata framework across the enterprise. This was one of the industry’s largest migrations. Pulling fragmented data from hundreds of tech stacks, the firm moved more than 100,000 legacy objects from over 20,000 dashboards.

By drawing on a core execution engine, domain-based pods, and AI-enabled accelerators, the company reduced its migration timeline by approximately 25%.

The firm has also cut operational complexity thanks to consolidated data, an analytics backbone, and democratized data and AI. And with improved governance, enhanced self-service, and enterprise-wide data interoperability, it expects to reduce technical debt by up to 30% and accelerate time to insight by up to 40%.

With these data and technical foundations in place, the company is ready to move beyond isolated AI use cases and unlock enterprise-scale AI transformation.



Leaders invest in technologies that power the enterprise architecture

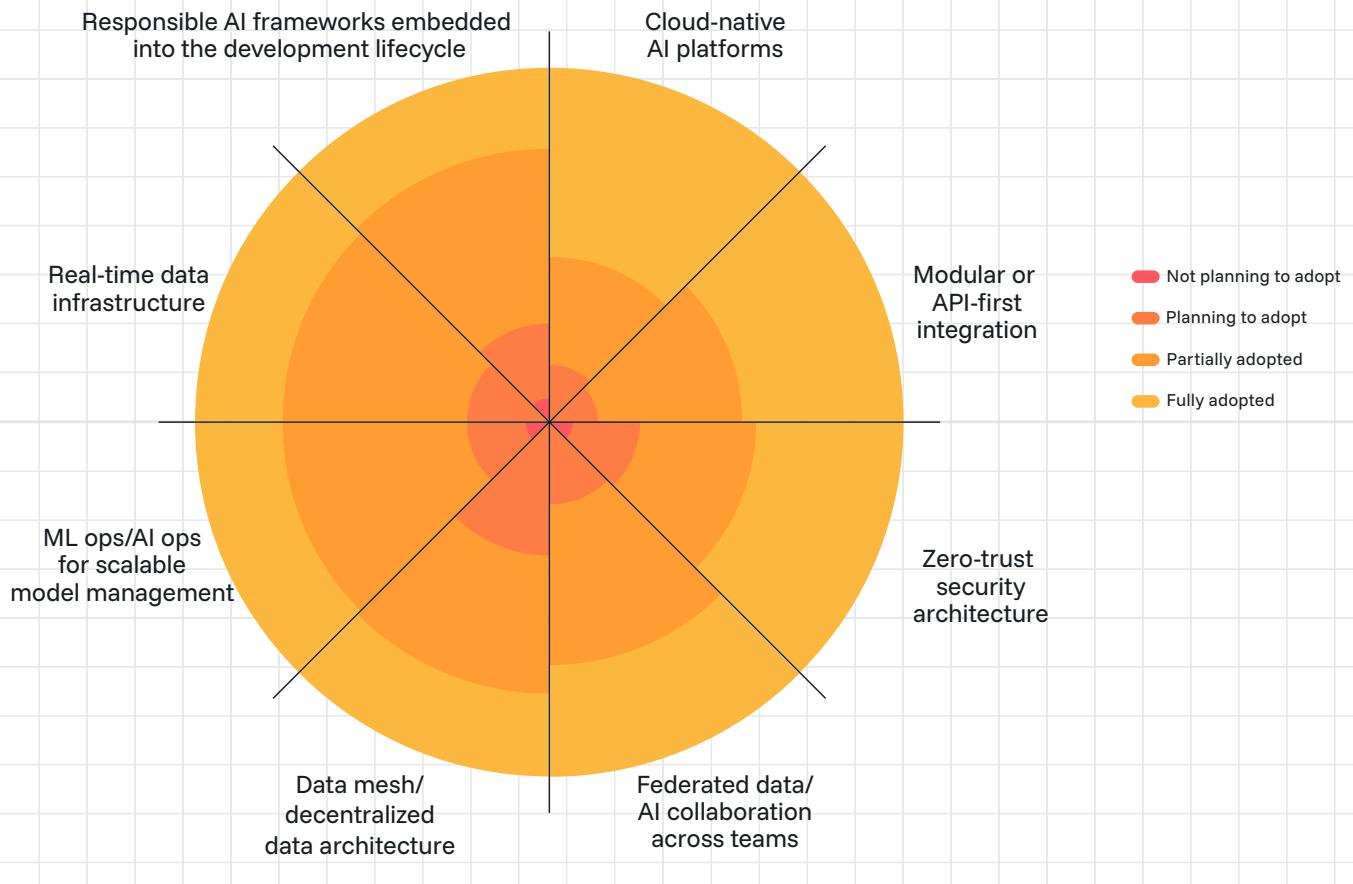


Figure 8: Leaders' adoption of the technical foundations to support advanced AI deployment





Governing at the speed of AI

Governance sets the pace – and limits – of AI innovation. When thoughtfully designed and consistently applied, it builds trust, secures executive alignment, and enables AI to deliver sustained enterprise value.

An overwhelming 99% of executives indicate they don't have adequate governance models and structures in place for autonomous or agentic AI systems and associated risks.

"Agentic AI requires new approaches to governance," says David Shrier, professor of practice, AI and innovation, Imperial College London.

"For example, when decisions are made by human/AI hybrid systems, 'idea provenance' becomes necessary. Companies must use AI itself to bring extra speed to monitoring and cybersecurity, to enable governance, and to keep pace with AI and innovation. But central to these approaches is the need for a robust governance model and framework."

99%

of executives indicate they don't have adequate governance models and structures in place for autonomous or agentic AI systems and associated risks.

Governance leadership

Acknowledging the need for robust models, leaders are building internal AI frameworks ahead of regulatory demands and are more likely to empower frontline teams with clear guardrails than route decisions through AI coordination councils. In addition, 67% have some level of AI review boards for compliance and risk evaluation, compared to 45% of non-leaders. The reported prevalence of AI review boards suggests growing awareness, though in many cases these structures may be nascent, more indicative of intent than of fully operational governance.

In parallel, our research highlights clear patterns in the way governance is structured and deployed. Centralized governance offers consistency, control, and streamlined decision-making, but it can be slow to adapt and less responsive to local needs. In contrast, federated governance enables agility and contextual responsiveness across domains yet requires strong coordination between central oversight and decentralized units to prevent fragmentation.

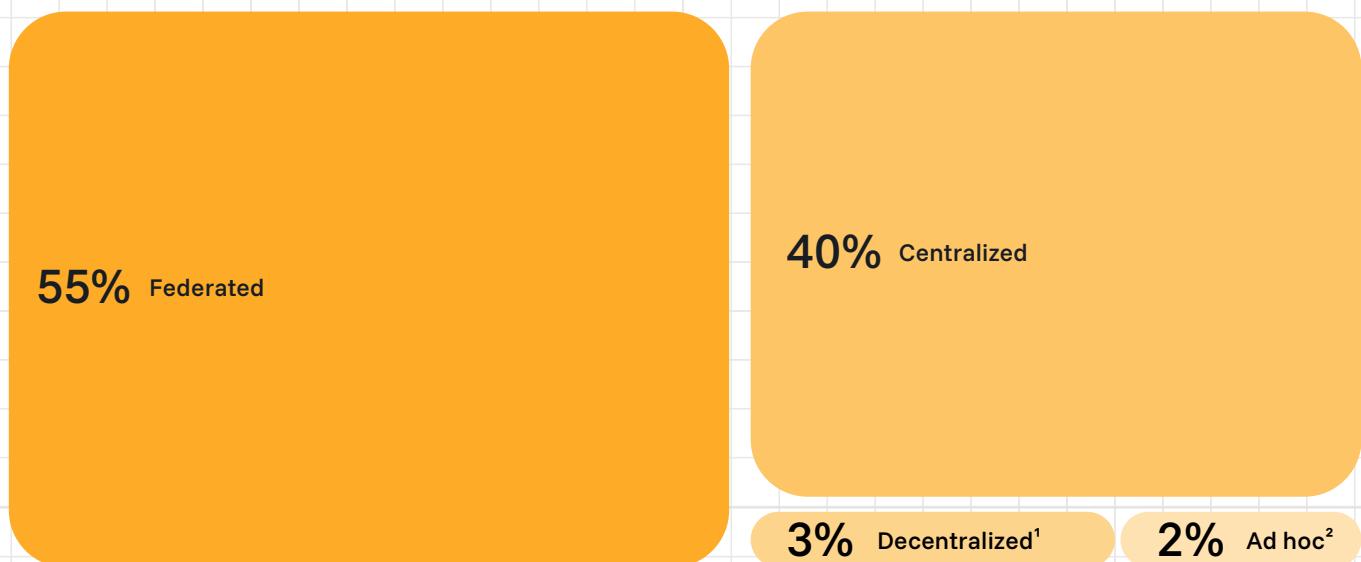


Federated or centralized governance?

Over half of leaders (55%) now operate under federated models, empowering business units to innovate within shared standards (figure 9). However, governance remains deeply contextual. For executives in highly regulated sectors such as life sciences (47%) and insurance (45%), centralized oversight prevails to enable compliance and risk management. As an operations director of a large healthcare firm notes, “We want to be pioneers [in AI governance], but we are also coming across some barriers with regard to finding the gold standard AI protocol at the moment, because there isn’t one per se.”

The implication is clear: effective AI governance must be adaptive – balancing autonomy with control while aligning with both industry demands and organizational design.

Federated governance structures deliver agility and context



¹Decentralized: Business units choose and manage their own AI use.

²Ad hoc: Governance varies widely across teams or functions.

Figure 9: How leaders are structuring governance in their organizations





Case study

Governance that fuels innovation

The CEO of a rapid-growth manufacturing engineering firm had a strong vision for AI. But the company needed to match its AI goals with AI governance to scale successfully while complying with ethical and regulatory standards to manage risk. And with intense competitive pressure, it had to safeguard intellectual property, too.

The company redesigned its AI governance operating model, assigning clear roles across departments and creating AI project workflows to maintain control without slowing innovation.

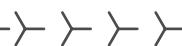
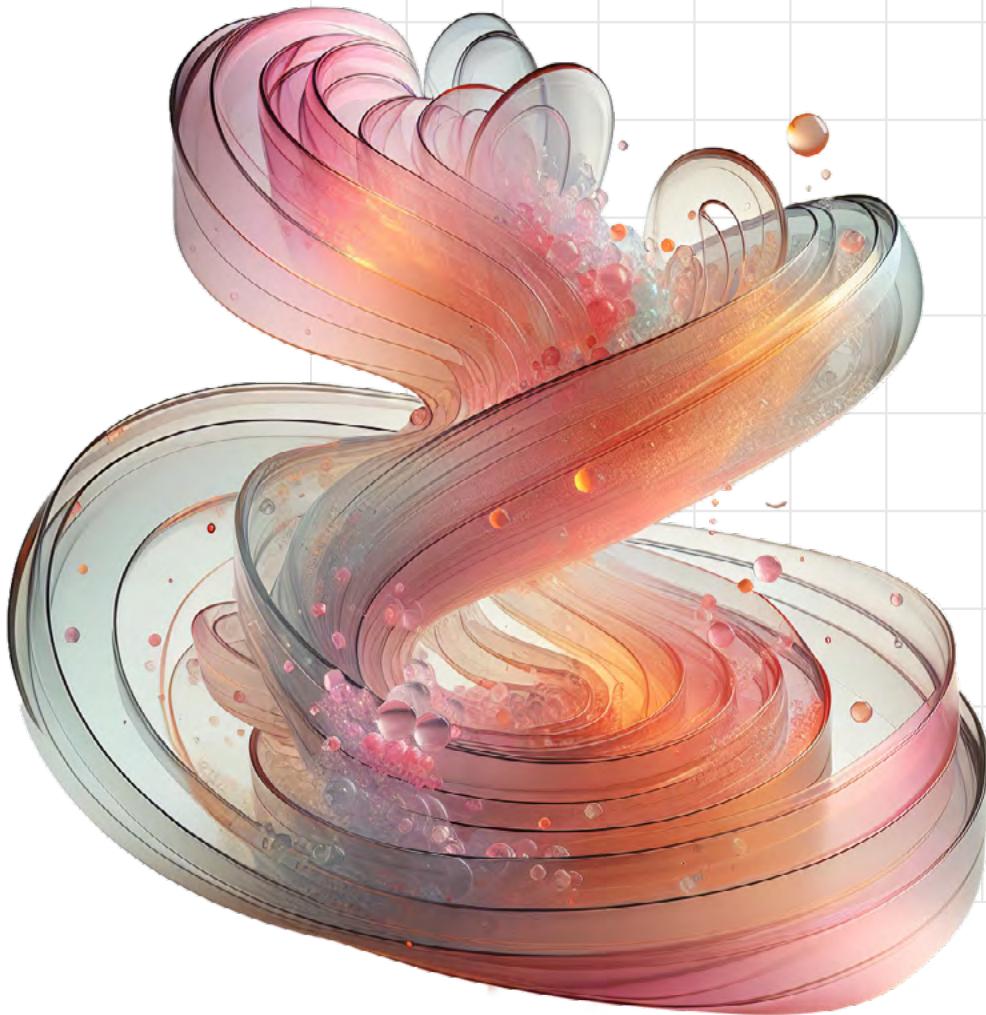
The company's culture of department-centric management was a challenge, so Genpact created a hub-and-spoke model to connect governance with the company's cultural identity. The firm set

up a centralized committee and embedded people in the business to help teams get comfortable with and see value from AI governance.

AI governance guardrails have now given the business the freedom to experiment. Governance is now seen as a resource to propel the business further, not hold it back. And teams have been innovating at speed. With a robust governance operating model in place, the company has launched nine AI projects that have unlocked up to \$350 million in value, with many more in the pipeline.

Just as critical: sensitive data stayed secure. No leaks. No incidents. Only value.





The leaders' playbook: Translating vision into action

Before advancing your organization's AI journey, get the basics right. Start with a baseline assessment of your current AI maturity across leadership, processes, technology, and workforce, and set a pragmatic ambition that matches this readiness.



Our leaders' playbook distills how leaders do it with practical actions for organizations at any stage of their AI journeys that are keen to move beyond AI adoption and build an autonomous enterprise. This playbook is grounded in Genpact's first-hand expertise and experience deploying agentic solutions across industries and functions.

1 A symphony of agents

Beginning your AI journey

Embed AI agents into core processes

Start with quick wins – focus on operations with extensive manual work, unstructured data, and rule-based workflow. These are low-hanging fruit for agentic AI.

Key considerations

- Which processes (or subprocesses) are ready and simpler to embed AI in, with a clear value opportunity?
- Where do you need 'humans in the loop' engaged with AI agents on critical decisions and complex actions?

As you mature

Redefine your processes to lead with AI

As AI becomes pervasive within the organization, use agents as interconnected systems for seamless workflows across business functions. Stand up a unified orchestration layer with process intelligence, making every agent traceable, ROI-accountable, and managed like a workforce asset.

Key considerations

- What are the new value parameters for AI-led processes?
- What's your agentic governance model – from 'humans in the loop' to 'humans on the loop'?

2 The universal AI practitioner

Beginning your AI journey

Elevate AI fluency across your organization

Every employee should be able to interpret, challenge, and escalate AI outputs. They must also shift their mindsets to make AI fluency, critical thinking, and oversight core competencies, which will help calm job-loss fears.

Identify which roles AI will augment, which it will absorb, and which new roles you must create, such as AI orchestrators, agent managers, and AI risk stewards.

Key considerations

- If AI is already reshaping work, which roles will you redesign first, and how will you enforce enterprise-wide fluency?

As you mature

Redesign your workforce and operating model to be AI-first

Redraw your operating model, workforce mix, and talent strategy to bring in AI agents. Job architecture, performance metrics, and accountability must shift to how humans and agents create value together, not how people perform alone. Position employees as coauthors of AI change.

Key considerations

- What cultural shifts are required to turn employees into coauthors of AI change rather than passive adopters?
- How will you redesign your organization's performance architecture now that AI agents are part of the workforce?



3 Enterprise architecture redux

Beginning your AI journey

Build foundation before innovation

Agents cannot be layered on fragmented systems. They need a coherent data and integration spine: systems of record + integration + agentic layer.

A hybrid architecture with guardrails and federated delivery creates central standards for security, cost, control, and interoperability, while allowing local teams to innovate within boundaries.

Key considerations

- What are the points of friction in your data flow and app interactions, and how can you mitigate them?
- What's your build vs. buy strategy? Don't build what you can buy.

As you mature

Build an architectural backbone for AI-led workflows, high-velocity data, and open integration

Use an agentic development lifecycle, instead of a traditional software development lifecycle, to guide all new design decisions. Next, establish an AI orchestration platform as your core intellectual property that enables seamless agentic interconnection across vendor platforms and custom applications, enabling plug-and-play functionality.

Key considerations

- How do you retain control of your security and risk architecture as you democratize the use of new agentic platforms and open integration?

4 Governing at the speed of AI

Beginning your AI journey

Establish a centralized governance backbone before federating

Create a light but authoritative center of excellence that sets standards and guardrails and prevents shadow AI before federating delivery.

In addition, define what 'safe to scale' means by codifying minimum conditions such as data integrity, audit trails, risk thresholds, and human override points. While regulators will impose external rules, internal discipline must set the enterprise standard.

Key considerations

- Which guardrails must be nonnegotiable to balance speed, cost, and risk in AI adoption?

As you mature

Embed automated monitoring, traceability, and policy enforcement into your AI architecture

Make sure AI oversight spans the full chain – harmonized data inputs, governed model decisions, and controlled agentic orchestration – with leaders held accountable for outcomes driven by both humans and agents.

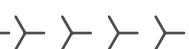
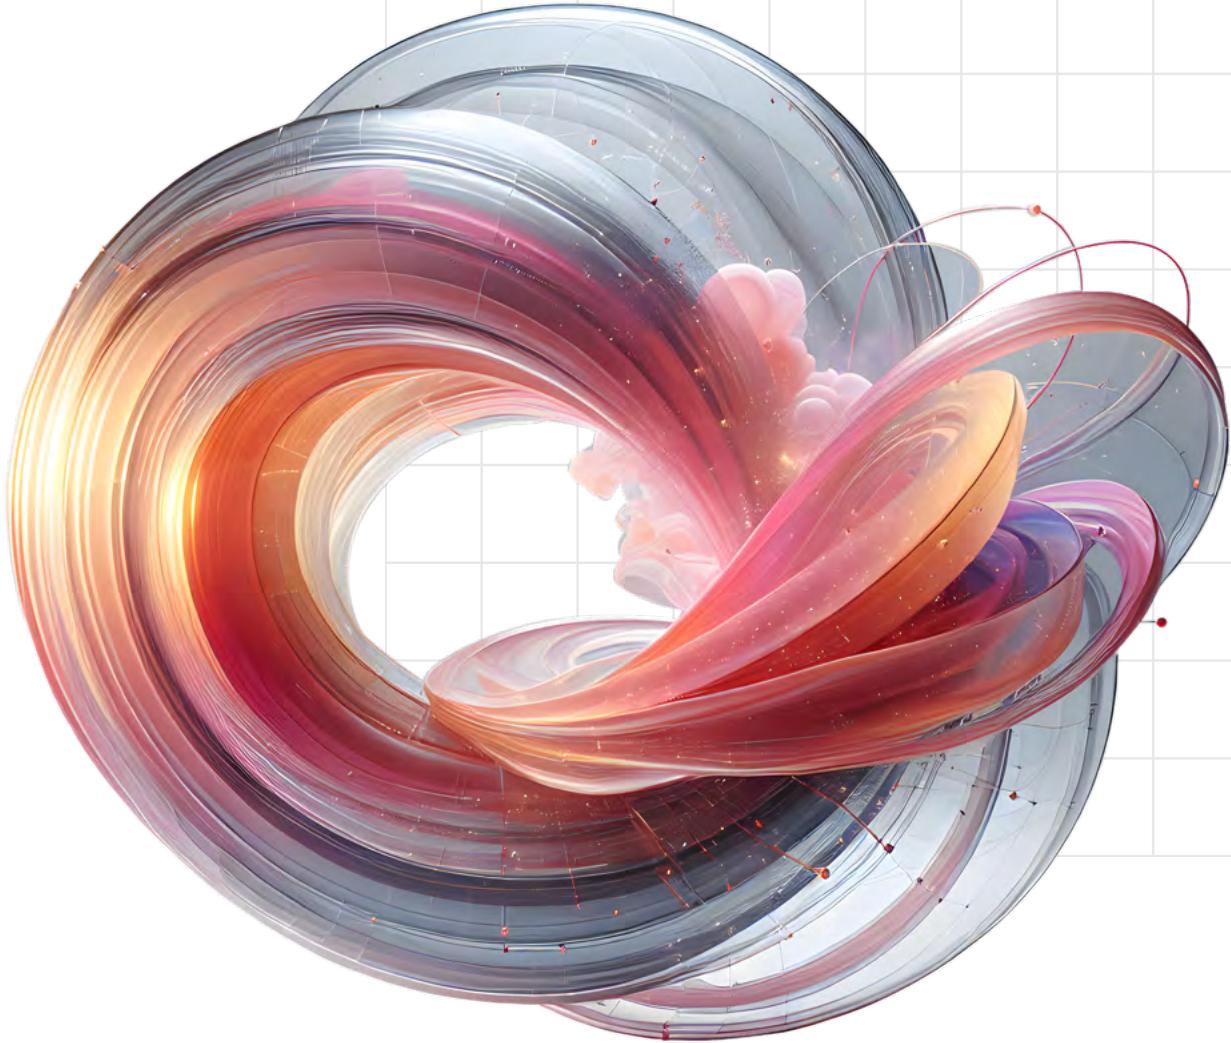
Key considerations

- How will governance evolve from compliance to a strategic advantage in scaling AI responsibly?

Deliver sustainable value with AI

The four autonomous enterprise enablers – a symphony of agents, the universal AI practitioner, enterprise architecture redux, and governing at the speed of AI – are interconnected capabilities that must evolve together. Mastering these enablers will help your business transition into an autonomous enterprise and turn AI into a lasting competitive advantage.





About this research

Understand who we surveyed and interviewed, how we benchmarked AI maturity, and the questions we asked.



Methodology

This report draws insights from a global survey of 500 senior executives working across 13 business functions in organizations generating over \$1 billion in annual revenue, spanning 8 industries and 13 countries.

We also conducted qualitative interviews with 10 senior executives and AI experts across industries to provide expert perspectives on the context of enterprise AI today. We ran the survey and interviews in Q3 2025.

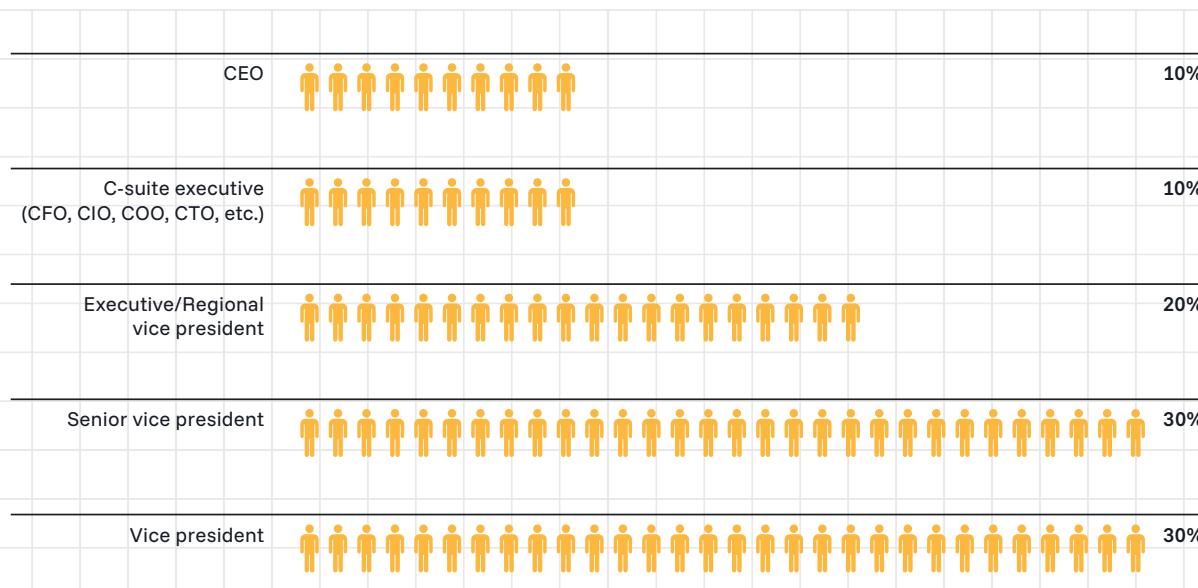


Figure 10: Respondents' role levels

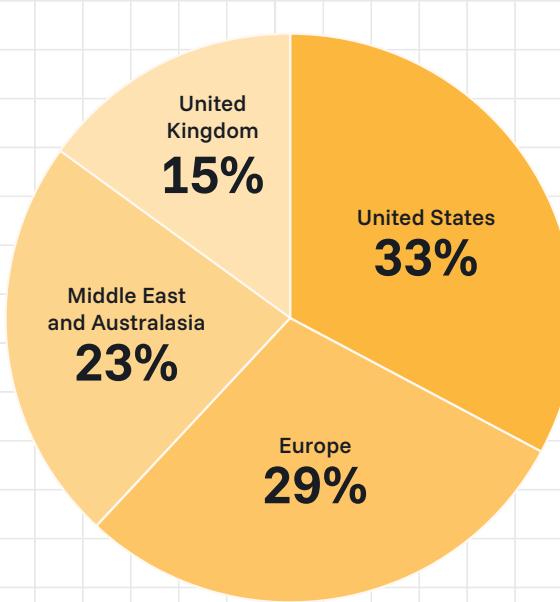
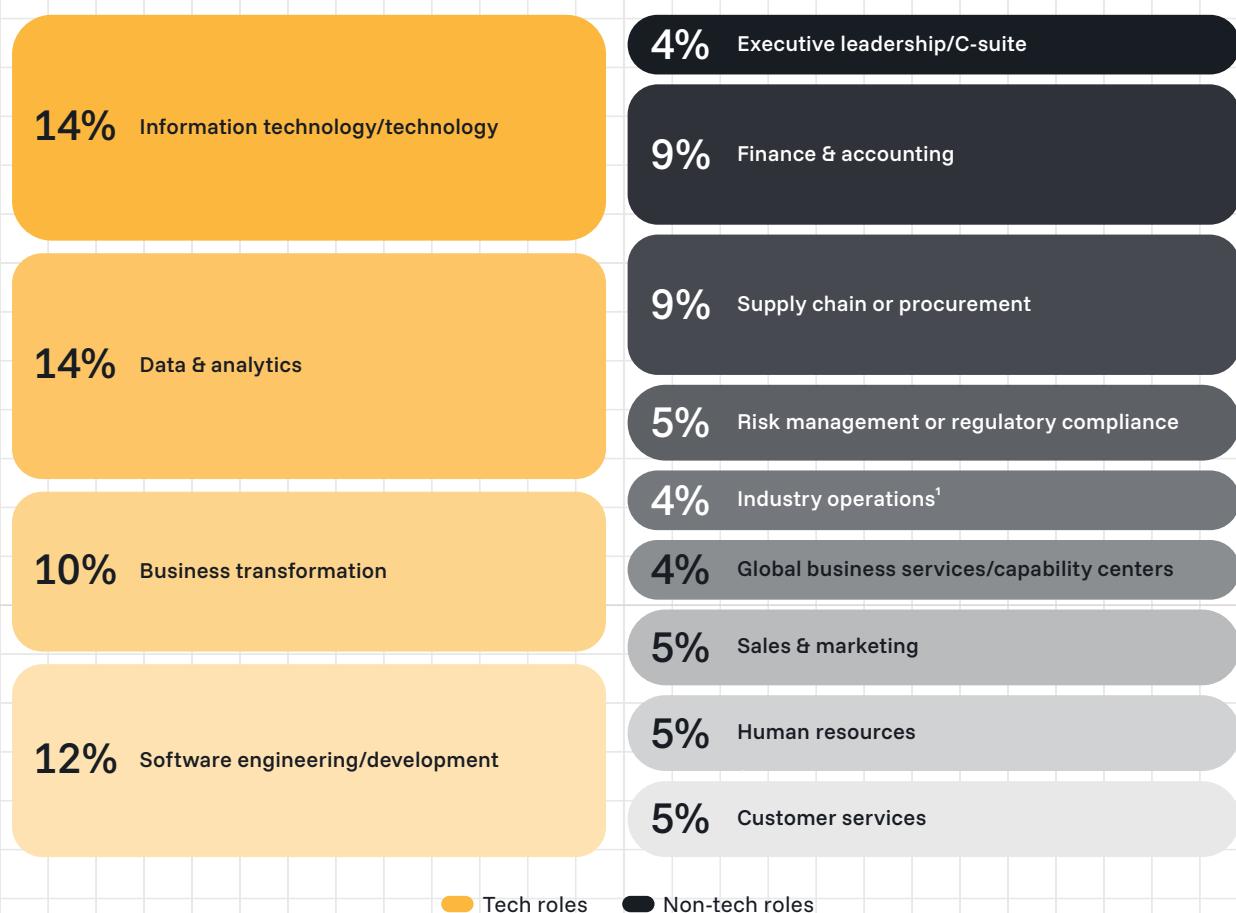


Figure 11: Location of company headquarters





Figure 12: Organization size by revenue



¹For example, claims, underwriting, financial crime.

Figure 13: Respondents' roles



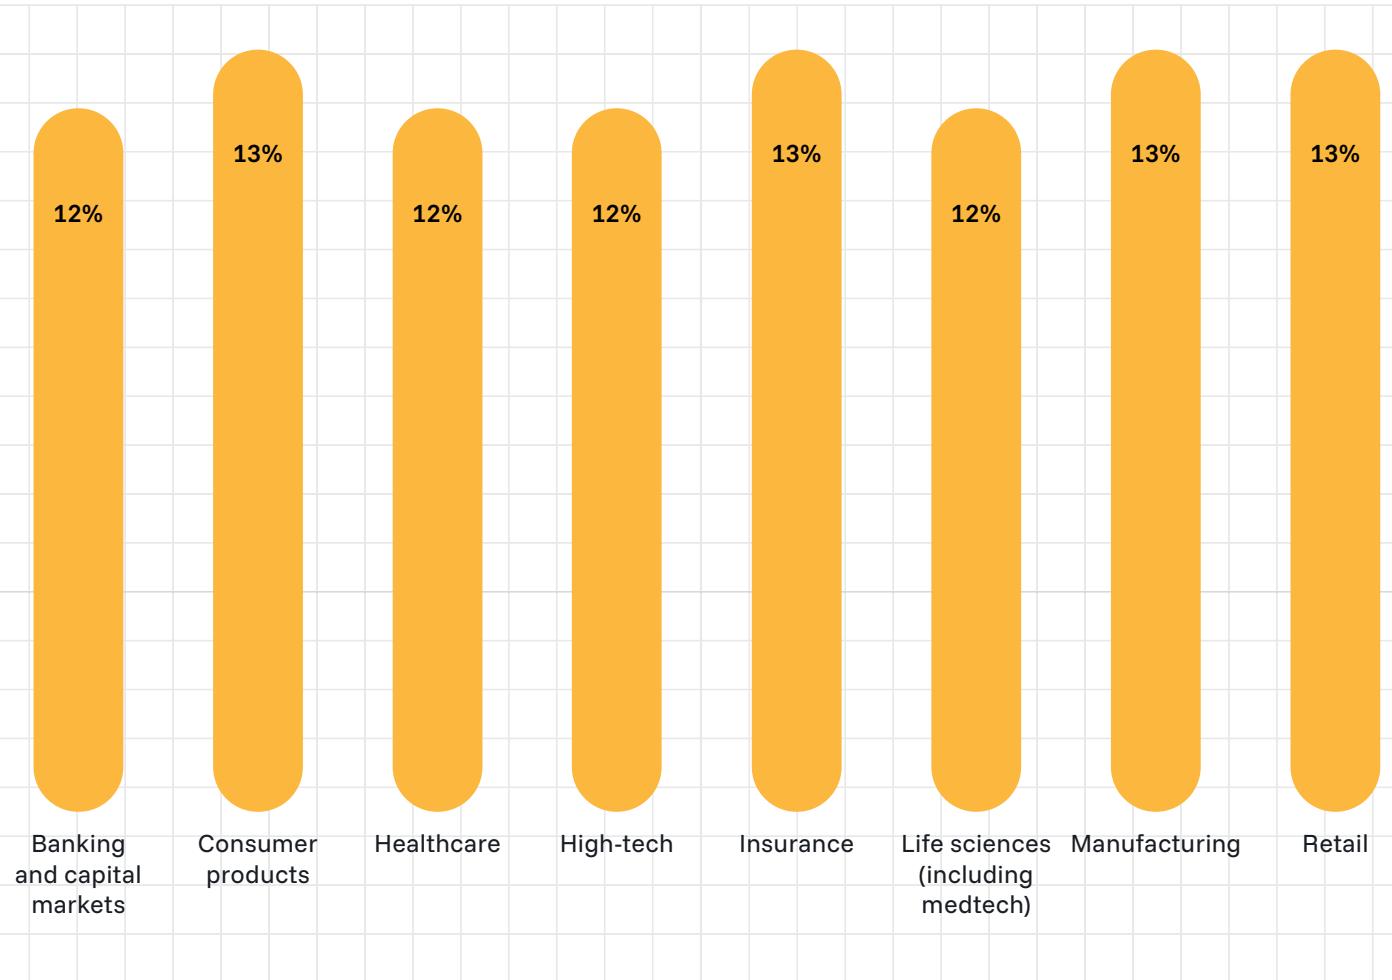


Figure 14: Respondents by industry



How we benchmarked autonomous AI maturity

This report uses a proprietary framework to assess how effectively enterprises are evolving toward self-directing, scalable, and well-governed AI systems – the hallmarks of autonomous AI maturity – and to create an index of autonomous AI maturity for each respondent.

We assigned each question a defined scoring scale and weight, calibrated to reflect its influence on enterprise readiness.

Dimension	Question	Scoring
Application effectiveness How broadly and effectively AI is applied	Which types of AI applications are currently in use at your organization, and how effective has each been at delivering measurable business value?	Responses scored by effectiveness across seven AI use cases, with higher weight for autonomous agents
Governance readiness Whether organizational structures can responsibly support agentic and autonomous AI	To what extent will your organization's AI governance structure need to be revisited to accommodate autonomous or agentic AI systems?	Weighted scale from minimal to extensive structural change, rewarding established governance maturity
Decision autonomy The degree of human and AI collaboration in enterprise decision-making	In your department, how are the following decisions and processes typically carried out and how much autonomy do AI systems have today?	Points awarded by the level of AI-led decision-making across core business functions
Operational enablement Whether dedicated teams or systems exist to manage autonomous AI operations	How is your organization coordinating AI systems across different departments?	Points awarded if dedicated agent operations teams are in place
Compliance mechanisms The maturity of responsible AI practices and accountability framework	What are the main ways your organization is preparing to comply with emerging AI regulations over the next 12 months? (Select up to three)	Points awarded for the number of responsible AI practices adopted
Advanced capabilities Deployment of higher-order capabilities such as orchestration, coordination, and multi-agent ecosystems	How is your organization engaging with the following capabilities? (Agentic orchestration, synthetic data, synthetic workforces, autonomous business units, real-time strategic simulation, self-healing systems)	Weighted by maturity of implementation across six next-generation AI capabilities, with orchestration double-weighted

Table 2: Benchmarking scores and weights for determining AI maturity



Segmentation and distribution: We ranked all 500 executives by their total score and used percentile thresholds to segment them into distinct maturity bands:

- **Leaders (≥ 90 th percentile / 11.6%):** Organizations furthest along the path to autonomous enterprise maturity
- **Advanced (75th–89th percentile / 14.8%):** Strong foundations in place, scaling autonomy across functions
- **Assisted (26th–74th percentile / 47.2%):** Developing capabilities and operational alignment
- **Emerging (≤ 25 th percentile / 26.4%):** Early-stage experimentation or limited deployment

Limitations: This benchmarking does not represent an absolute performance rating but a relative measure of progress toward enterprise autonomy. It is based on self-reported results from respondents, some of whom may work at different locations of the same organization.

Figure	Questions asked
1	See: How we benchmarked autonomous AI maturity on p34. N=500
2	In your department, how are the following decisions and processes typically carried out – and how much autonomy do AI systems have today? N=500
3	What are the most significant pain points relating to AI implementation at your organization? N=500
4	What are the most significant pain points relating to AI implementation at your organization? N=500
5	How is your organization engaging with the following capabilities? Agentic orchestration. N=58 (all leaders)
6	1 What are the main ways your organization is preparing to comply with emerging AI regulations over the next 12 months? N=58 (all leaders) 2 What are the main ways your organization is managing change as it adapts to AI? N=58 (all leaders) 3 Which of the following does your organization offer to support AI innovation? N=58 (all leaders) 4 How is your organization coordinating AI systems across different departments? N=58 (all leaders)
7	To what extent is the complexity of your technology architecture a challenge for your organization's ability to scale AI? N=52 (from Microsoft Ignite poll)
8	To what extent has your organization adopted the following technical foundations to support advanced AI deployment? N=32 (leaders in technical roles)
9	How is AI access and governance structured in your organization? N=58 (all leaders)

Table 3: Survey questions and respondent numbers used in charts



Contributors



Jayant Swamy
Chief Enterprise Architect,
Genpact



Ajay Vasal
Data and AI Advisory Lead,
Genpact



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



Dr. Pramukh
Nanjundaswamy Vasist
Thought Leadership and
Innovation, Genpact



Ojas Save
Technology Strategy Lead,
Genpact



David Shrier
Professor of Practice, AI and
Innovation, Imperial College
London, and Managing Director,
Visionary Future



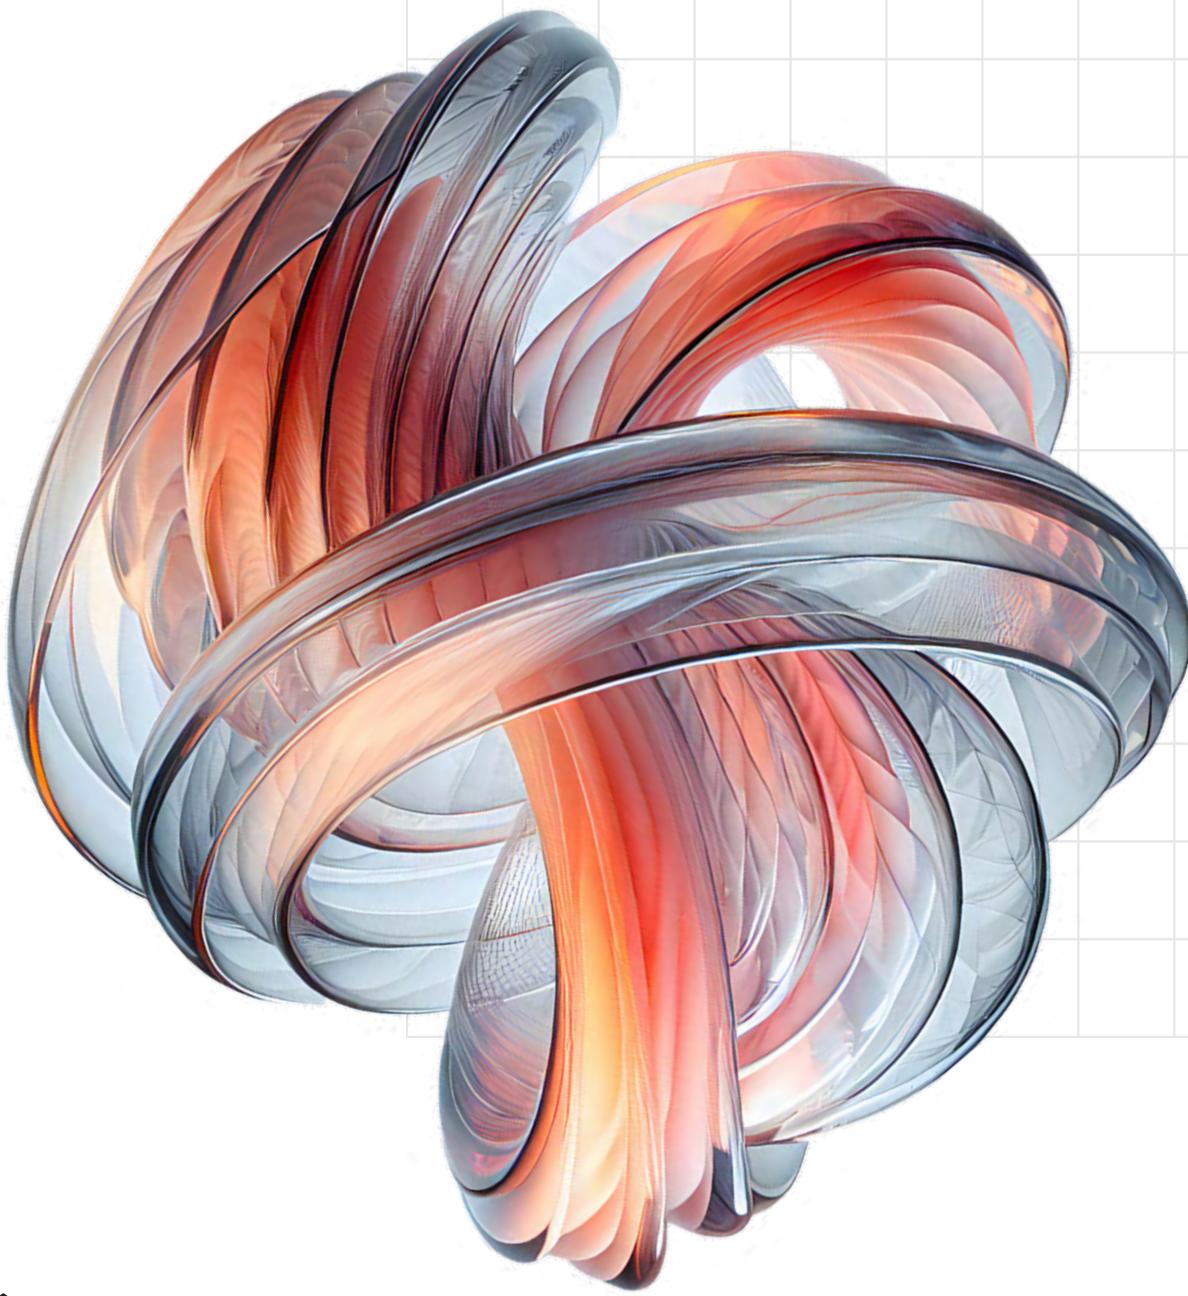
Sandra Higgison
Editor, Genpact Intelligence,
Genpact



Genpact Intelligence: our curated collection of insights inspired by our work with advanced technologies and enterprises around the world. Use our research, case studies, and perspectives to get ahead and stay ahead.

genpact.com/intelligence





Genpact (NYSE: G) is an agentic and advanced technology solutions company. We leverage process intelligence and artificial intelligence to deliver measurable outcomes. With a strong partner ecosystem and decades of client trust, we provide innovative solutions that transform how businesses run. Powered by a team with an active learning mindset and client centricity at its core, we deliver lasting value for the world's leading enterprises.

Get to know us at genpact.com and on [LinkedIn](#), [YouTube](#), [X](#), and [Facebook](#).

