Deloitte.

The *State of AI* in the Enterprise

Built differently: How AI is reshaping the way enterprises work, grow, and create value

October 2026

deloitte.com/us/state-of-ai

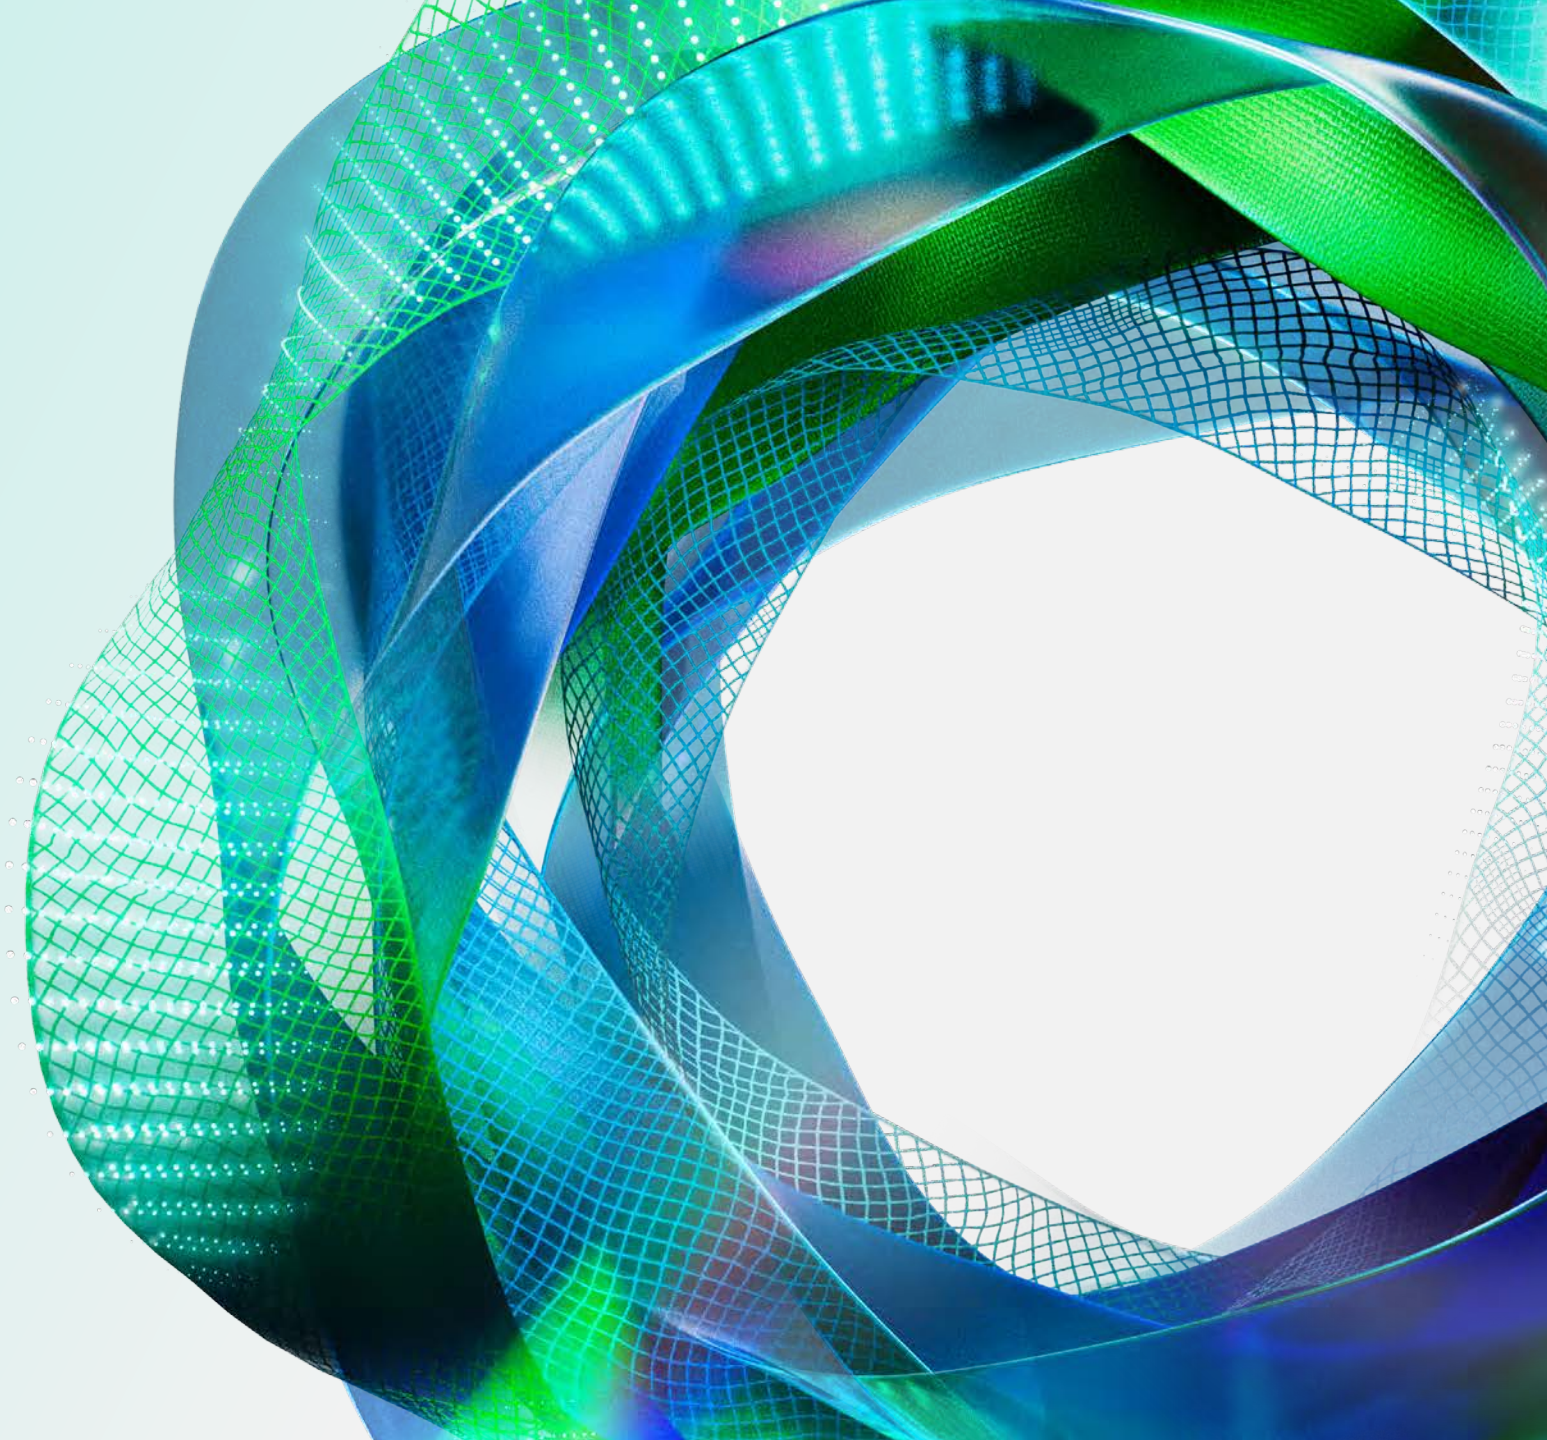


Table of contents

Introduction 03

Overview 06

About the report 09

Key findings 10

AI agility 10

AI economics 13

ROI paradox 16

Work redesign 18

Governance as operating system 20

Physical AI 22

Sovereign AI 25

How to build differently 27

Redesign the work, not just the workflow 28

Trace cost to the token, not just the total bill 28

Connect AI initiatives to the commercial decisions that produce revenue 29

Match governance to the autonomy AI is actually given 29

Build Physical AI expertise ahead of the deployment curve 29

Treat sovereignty as an ongoing strategic input, not a one-time fix 30

Acknowledgments 31

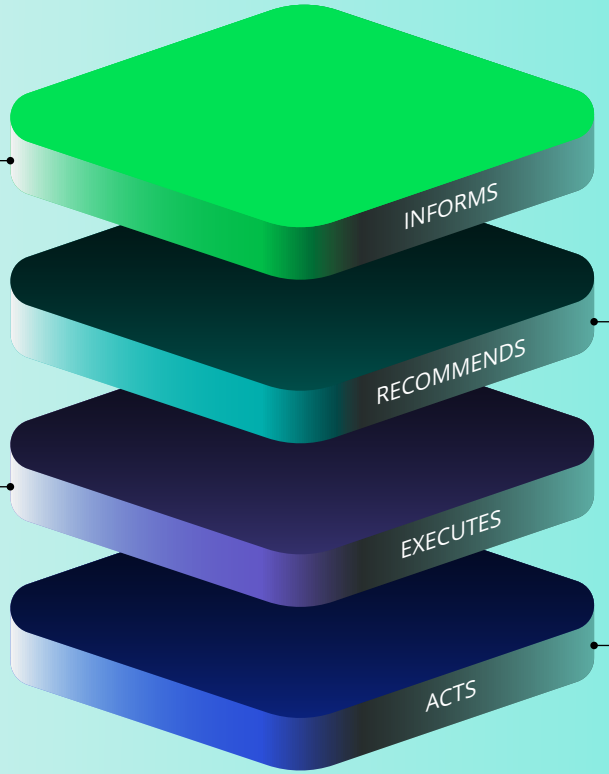
About the Deloitte AI Institute 34

Methodology 35

Introduction

Artificial intelligence (AI) has moved decisively into the enterprise, and Agentic AI is beginning to shift systems from answering questions and supporting decisions to taking bounded, multistep actions within workflows. The question is no longer whether companies will adopt AI, but whether they can turn its rapidly expanding use—and growing autonomy—into durable business value. After several years of experimentation and broader deployment, initial excitement is giving way to a more practical understanding of what it takes to integrate AI into the business. Technology continues to advance at extraordinary speed, but creating value at scale remains the harder challenge.

AUTONOMY LEVELS



01
Informs
- Explicit pre-approval
- Continuous/real time
- Escalation triggers
- Supervising decision-maker

02
Recommends
- Pre-action sign-off
- Milestone-based
- Escalation triggers
- Shared accountability

03
Executes within limits
- Pre-authorized guardrails
- Periodic/sampled
- Escalation triggers
- Functional lead/actor

04
Acts autonomously
- Zero prior approval
- Outcome and trend review
- Escalation triggers
- Individual operator

WORKFLOW

Approval required → Monitoring and oversight → Escalation triggers → Accountable owner

Introduction

To succeed in the future, companies will need to be built differently—with AI at the core of many or most processes and workflows. However, that challenge now has a sharper economic edge. As AI becomes more deeply embedded in decisions, workflows, products, and operations, leaders are moving beyond questions of access and adoption to questions of cost, token consumption, and return on investment (ROI). The enterprise must evolve not only how it redesigns work, measures value, and assigns accountability, but also how it manages the infrastructure, expertise, and economics required for broad deployment. This is no longer just a technology implementation conversation. It's a question of enterprise design: how work gets done, how decisions get made, and how AI becomes a productive and economically durable part of the operating model. That is less a matter of redrawing the org chart than of redesigning the work chart itself: how work gets done, where decisions are made, and how accountability is shared between people and AI systems.

The latest *State of AI in the Enterprise* report finds that progress has been meaningful, but uneven. It also points to why. The binding constraint on enterprise AI is no longer access to the technology or a willingness to experiment; it is the organization's capacity to evolve its operating model around AI—at the speed the technology now sets. Increasingly, that bottleneck also runs through AI economics: the ability to understand where tokens are going, what they are producing, and whether the value created justifies the cost.

This iteration of the report brings that tension into focus. Organizations are pushing for more AI use, even as many confront a widening gap between what AI costs to run and what it demonstrably returns. The new imperative is not to slow adoption; it is to pair ambition with operating rigor so that more AI use produces more value rather than simply more spend.



The binding constraint on enterprise AI is no longer access to the technology or a willingness to experiment; it is the organization's capacity to evolve its operating model around AI—at the speed the technology now sets.

Introduction

At the same time, the AI agenda is widening. Agentic AI is moving AI systems from supporting decisions to taking bounded actions within workflows, while Physical AI is bringing intelligence into machinery and operations. Geopolitical pressures are raising new questions about geographic dependence, resilience, and control, moving Sovereign AI higher on the enterprise agenda. Together, these developments are pushing AI beyond individual tools and use cases and into the core of how companies are designed and run.

This report examines how enterprises are building differently around AI across seven key finding areas—AI agility, economics, value realization, work redesign, governance, Physical AI, and Sovereign AI—and discusses what it will take for companies to move from adopting AI to operating differently because of it. Across those findings is a clear corporate shift: from using more AI at any cost to using more AI wisely.

Enterprise AI will not be defined by technological capability alone. It will be shaped by the choices leaders make about how their companies invest, govern, design work, and adapt. The priority is not simply to do more with AI, but to build an enterprise capable of turning AI's unique and expanding potential into enduring value. The organizations that move ahead will be those that combine technical progress with the management systems needed to make AI productive, accountable, and economically durable.



Geopolitical pressures are raising new questions about geographic dependence, resilience, and control, moving Sovereign AI higher on the enterprise agenda.

Overview

AI is becoming part of how companies run their business, not just a technology they deploy. The progress is real: AI is shortening decision cycles, producing more measurable value, transforming jobs, and moving into physical operations.

But progress remains uneven: Companies are accelerating workflows faster than they are reshaping the business around AI. The technology is more ready than the organization is; that gap is the bottleneck.

Key findings

Faster, not different: 81% say AI speeds decisions, but only 25% are deeply transforming the business.

Speed is the baseline. Redesign is the vehicle and the constraint.

Of the companies surveyed, 81% say AI has improved the speed of decision-making and execution, making faster decisions one of AI's most widespread benefits. Yet only 25% are using AI for deep transformation—meaning to create new products or services, reengineer core processes, or fundamentally reinvent their business model. Speed is becoming common. Deeper transformation is not.

71% see the total bill. ***Only 44% know the unit cost.***

Visibility often stops at the top line.

Companies can increasingly see what they are spending on AI, but many still cannot see what is producing the bill. Seventy-one percent track total AI inference and application programming interface (API) spend, while only 51% track token volume and 44% track per-unit inference cost. The total bill is visible. What sits underneath it is much less so.

AI revenue realization rises to 28%, as the gap between investment and return narrows.

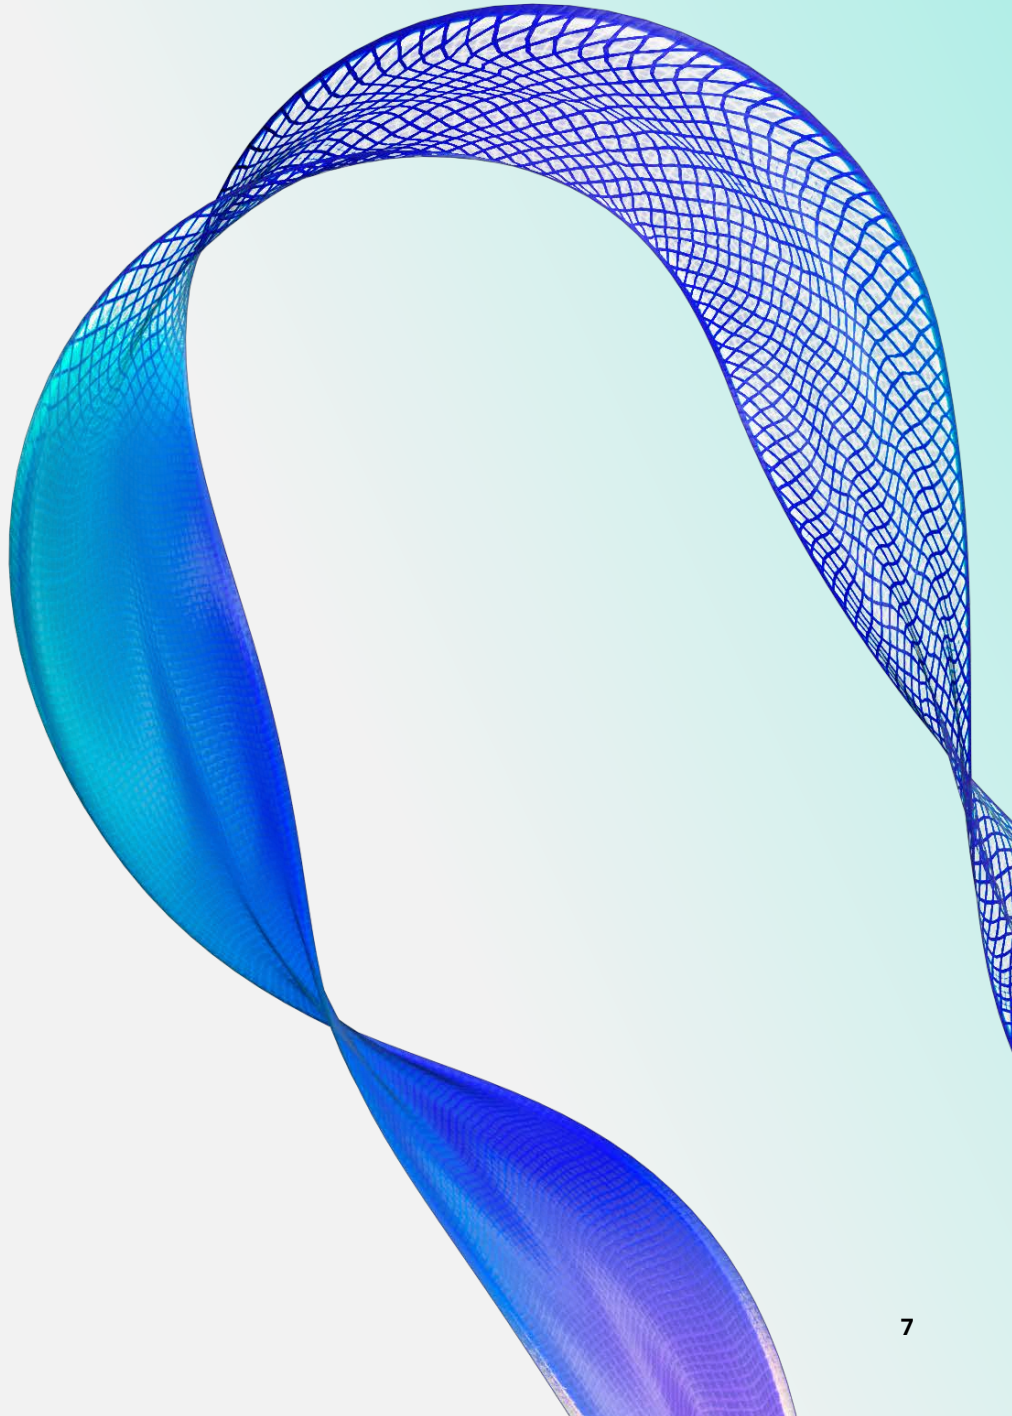
The value story is moving beyond productivity.

Twenty-eight percent of companies now report revenue gains from AI, up from 20% in the January report. Progress is also visible across efficiency, decision-making, cost reduction, customer relationships, and innovation. Revenue remains the return that companies expect most in the future, but it isn't the only one that counts. AI's business case is no longer just productivity and cost savings; it now includes a wider set of operational, customer, and innovation outcomes.

60% have redesigned jobs around AI, but many are still just scratching the surface.

Work is changing, but full transformation remains rare.

Of the companies surveyed, 60% have already redesigned jobs around AI capabilities to at least a moderate degree. This is a good first step. However, only 15% describe those changes as extensive or complete. Similarly, only 15% report a significant or fundamental shift in decision-making authority as a result of AI.



From guardrails to accountability: *Governance must be built in, not bolted on.*

Accountability is becoming more important as autonomy expands.
Forty percent of companies allow AI to execute decisions or actions without human approval to at least a moderate degree. Most can identify who is accountable when an AI system produces a negative outcome, but nearly one in four, 23%, still say accountability is unclear. As AI moves from recommending actions to taking action, that remaining ambiguity becomes more consequential.

Physical AI is scaling. Expertise will determine who captures the value.

Adoption is rising faster than expertise.
Today, 25% of companies report at least moderate use of Physical AI. Within two years, that figure is expected to reach 48%. However, expertise remains scarce: Only 19% rate their Physical AI expertise as high or very high. Adoption is expected to grow faster than the skills needed to deploy these systems well.

Sovereign AI stays on the strategic agenda for 74% of organizations, as they diversify where and how their AI is built and operated.

Control is becoming part of the AI strategy. Seventy-four percent of companies say data residency and in-region compute are at least moderately important to their AI strategy, while 62% are at least moderately concerned about reliance on foreign-owned AI technologies and compute. Infrastructure, geography, and control are becoming part of the AI strategy alongside models, applications, and use cases.

About the report

This report is based on a survey of 3,500 senior leaders across 25 countries and six industries: consumer; energy, resources and industrials; financial services; life sciences and health care; technology, media and telecommunications; and government and public services.

Respondents included board and C-suite members and leaders at the president, vice president, and director levels. The sample was roughly evenly divided between IT and line-of-business leaders. All participating companies have one or more working AI implementations in daily use.

Conducted between May and June 2026, the survey was supplemented by 15 interviews with executives and AI and data science leaders from large companies across industries and countries.

For full details, see the Methodology section. This report is part of an ongoing series from the Deloitte AI Institute designed to help business, technology, and public-sector leaders track the pace of AI change and adoption.

Learn more. →

Key finding

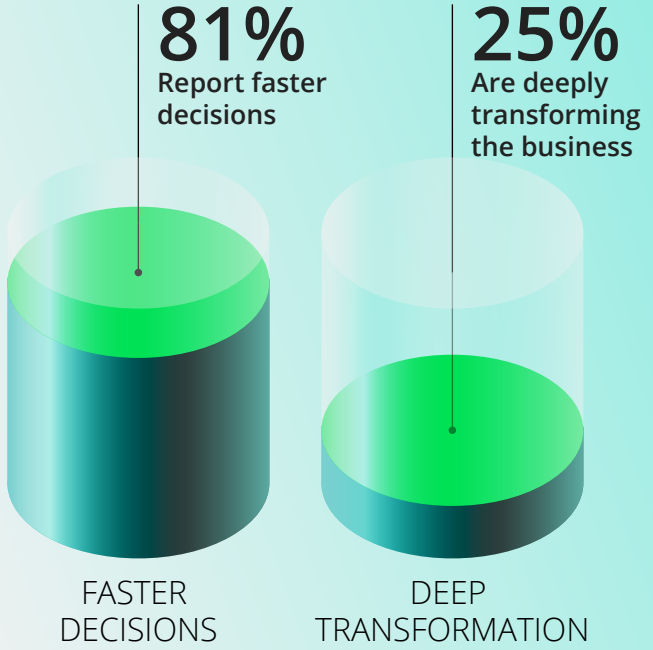
AI agility

Faster, not different: 81% say AI speeds decisions, but only 25% are deeply transforming the business.

AI enables companies to move faster; however, it has yet to change the way they operate. Within that figure, 43% report a slight improvement, 31% a moderate one, and 7% a significant one. Increased speed is no longer a differentiator; it is now the baseline. If an organization isn't seeing faster decisions from AI today, it's behind the curve. The constraint is no longer whether AI can accelerate work. It's whether the organization is prepared to redesign the work it accelerates.

The gains are consistent across industries. Across most industries, faster decision-making falls within a narrow band of 78% to 82%, with financial services at the low end of that range and technology and energy tied at the top of it. The spread is too tight to call a real pattern—speed has become a common benefit of AI, not one limited to a particular type of company.

SPEED vs. TRANSFORMATION*



Increased speed is no longer a differentiator; it is now the baseline.

Question: *How has AI impacted your organization's speed to decision-making and execution?*

Question: *Which best describes your organization's current approach to process transformation during AI implementation?*

Base: all respondents (N=3,500)

* Figures may not sum to 100% due to rounding.

Faster does not mean transformed

Greater speed does not always mean deeper change. Only 25% of companies are currently using AI to deeply transform their businesses: creating new products and services, reinventing core processes, or fundamentally changing their business models. That is down from 34% in the January report, even as speed gains have become the norm. The other three-quarters are keeping the fundamentals of the business largely intact, with some either automating existing tasks, making minor changes, or redesigning selected processes without changing the broader model.

Generative AI has made it much easier to develop and test sophisticated tools without building everything from scratch. An AI architect described a sharply lower barrier to experimentation: Business teams can now test a use case quickly, see what works, and decide whether to continue. That is valuable, but it also makes it easier to stack up small improvements without stepping back to rethink the full process or business model.

For many companies, the immediate focus is still on practical gains. An IT director described a balanced objective: pursue business transformation while continuing to move faster. The leader identified productivity gains, faster reporting and analysis, better forecasting, and less administrative work as priorities. These benefits matter, but they mostly improve the way the business already operates.

AI can shorten a process or improve a decision without changing the underlying workflow, ownership, or customer experience.

That is the risk of mistaking speed for transformation. An inefficient process may run faster after AI is added while keeping the same handoffs, approvals, and constraints. A team may get to an answer more quickly without gaining the authority to act on it. A company may build a long list of successful pilots without changing how the business works as a whole.

The bigger opportunity is to use AI as an enabler to rethink the work itself. That means asking whether a process should be simplified, combined with another, or removed altogether. It also means reconsidering where decisions are made, which activities create value, and whether AI makes new products, services, or ways of serving customers possible.

The advantage is in how AI is applied

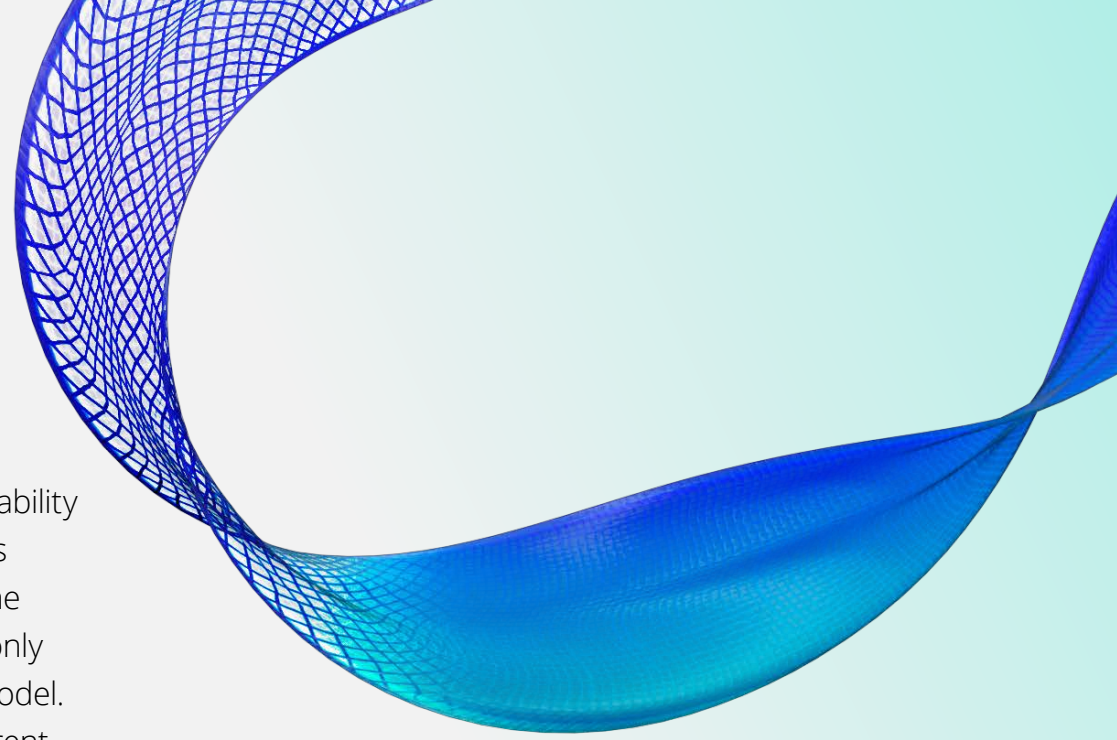
The advantage will not come from access to AI alone. A technology executive argued that applying AI is itself a skill and can remain a differentiator even when the underlying technology is widely available. Companies will still differ in how well they use the tools, redesign processes, and turn new capabilities into business results.

That difference becomes clearer when AI is used for more strategic decisions. The same executive pointed to cash forecasting as an example. The value is not simply that fewer people may be needed to produce a forecast. A more accurate forecast can improve how hundreds of millions of dollars are managed. In that case, the business result matters far more than the productivity savings.

Agility requires redesign, not just optimization

Agility is more than moving quickly. It is the ability to change direction, redesign how work gets done, and create new sources of value as the technology evolves. Companies that focus only on speed may improve today's operating model. Companies that go further can build a different one that is dramatically better.

The leadership challenge is knowing when optimization is enough and when it is time to redesign. Not every process needs to be reinvented, and not every use of AI needs to transform the business. But companies should be clear about the difference. Otherwise, a series of local improvements can look like transformation without changing the fundamentals of how the company competes.



AI is already helping most companies move faster. The next test is whether they can use that momentum to change how they operate, make decisions, and create value.

Key finding

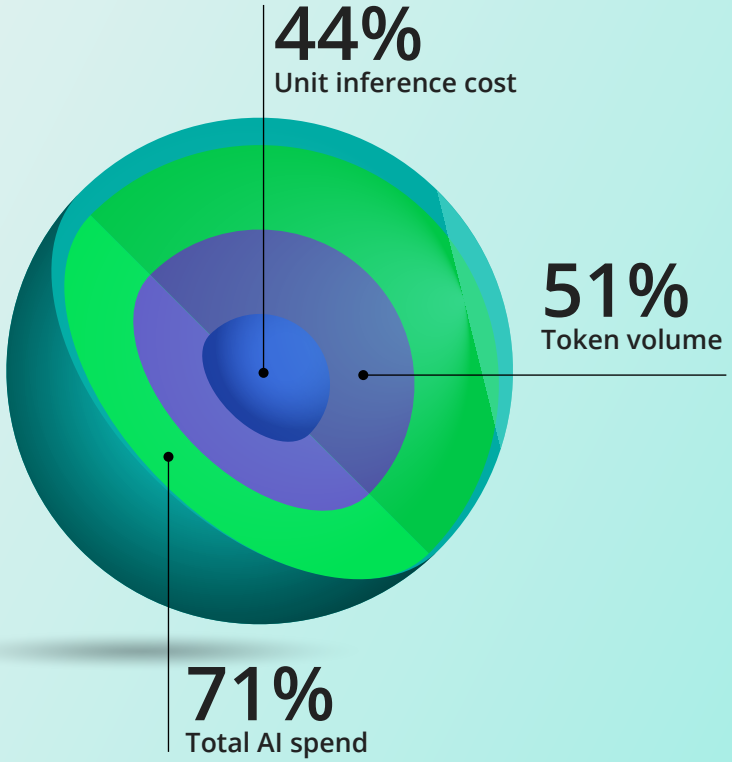
AI economics

71% see the total bill. ***Only 44% know the unit cost.***

A few terms are worth defining up front. A token is the basic unit AI models use to process text, roughly a few characters or part of a word, and every prompt and response is billed in tokens. The total bill is what a company spends in aggregate on AI usage across the organization. The unit cost is the price of a single token or API call, the number that reveals what a specific task, workflow, or product actually costs to run. Most companies can see the total bill. Far fewer can see the unit cost underneath it.

AI has become one of the fastest-growing expenses in corporate technology budgets, and the nature of that cost has changed. Where past technology waves were priced in licenses, seats, or virtual machines, AI cost accrues token-by-token. That shift matters because token consumption does not scale in a straight line: It moves with model complexity, usage patterns, and infrastructure choices, making AI spend

AI COST VISIBILITY*



Visibility thins as you go deeper.

Question: *To your knowledge, does your organization currently track the following?*

Base: all respondents (N=3,500)

* Figures may not sum to 100% due to rounding. Percentages represent separate measures and may reflect different companies.

inherently more volatile and harder to predict than the technology costs finance teams are used to managing. AI economics is the discipline of understanding what drives that cost, at the level of the token, the inference call, and the workflow, not just tracking the total. Not all AI usage creates value. Some organizations reward employees for adoption alone, without connecting that activity to measurable business impact. However, even for well-intentioned use cases, the cost of the AI can stealthily exceed the cost of the labor it replaces. That makes token economics a management issue, not merely a procurement or infrastructure issue.

The total bill is real, rising, and closely watched

If the first constraint is an organization's readiness to redesign work, the second is its ability to recognize what AI actually costs. AI investment has reached a scale that makes its cost a strategic issue for leadership, not just a technology budget line to revisit once a year. Eighty-six percent of companies expect their investment in AI technologies to increase in the coming fiscal year, with 34% reporting that it will increase by 20% or more. Amid this scale of investment, the total bill has become the most closely watched measure of AI economics: 71% of companies track total spend on AI inference and API consumption, the highest level of tracking across the metrics in the survey. That holds across industries, with Technology, Media & Telecom leading at 75%.

86%

of companies expect their investment in AI technologies to increase in the coming fiscal year.

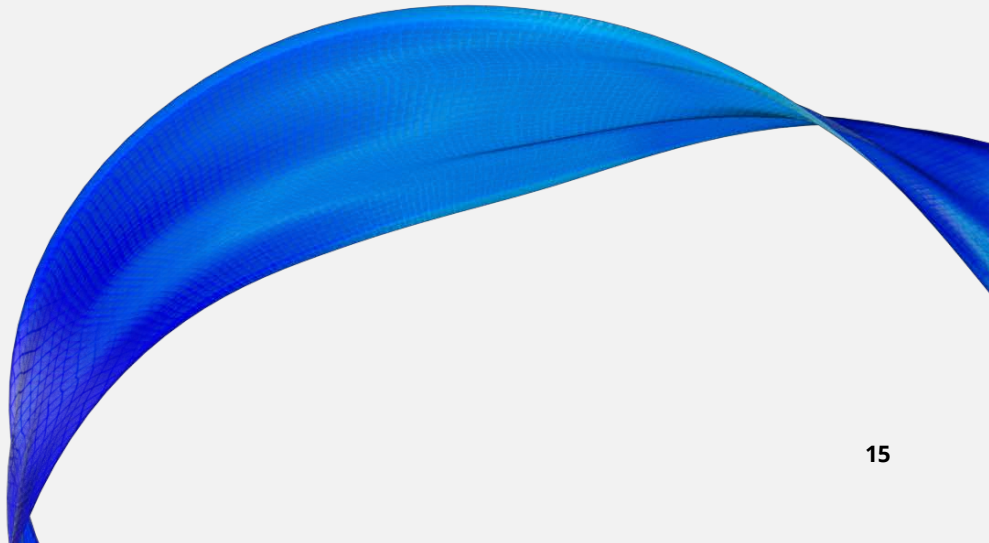
The cost visibility hierarchy

Unfortunately, only 51% of companies track their total volume of AI-generated tokens, the usage layer that actually drives spending. Fewer still, 44%, track the per-unit cost of AI inference, the cost per token or per API call that explains what activity is producing that spend in the first place. The pattern across these three metrics is a hierarchy. Visibility is highest at the aggregate level, thins out at the usage level, and is thinnest at the unit level—the layer closest to the actual driver of cost. Companies have made the top of the bill visible. What sits beneath it remains relatively opaque, and unevenly so: Technology tracks companies' per-unit cost at about 1.6 times the rate of Consumer, the widest gap between any two industries in the survey.

Governing AI as an economic system

Governing AI cost effectively means actively managing unit economics, not just watching the total bill. Traditional cost frameworks were built for spend that moved in predictable increments: more seats, more licenses, more virtual machines, each with a fixed, known price. Token-based consumption breaks that logic. The same workload can consume a different number of tokens depending on how a prompt is engineered, how long a reasoning chain runs, and which model handles the task, while the price per token itself shifts as model capabilities and infrastructure evolve. That combination makes AI spend behave less like fixed technology overhead and more like a variable economic system, one that must be carefully managed, not simply paid.

Left unaddressed, this gap exposes companies to margin compression that only shows up after the fact because it is impossible to optimize workloads when their cost drivers remain hidden. Building the discipline to trace cost from workflows down to tokens is what turns AI spend into measurable and meaningful ROI at the activity level.



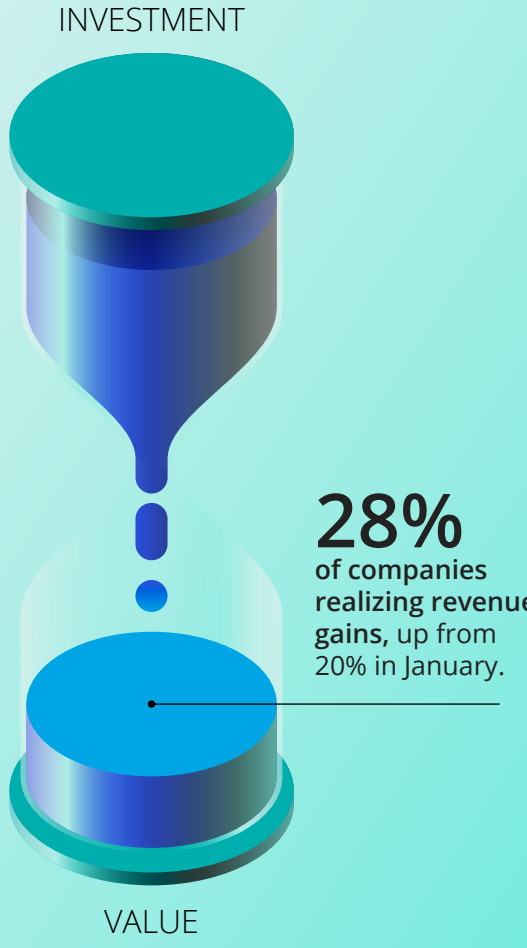
Key finding

ROI paradox

AI revenue realization rises to 28% as the gap between investment and return narrows.

For years, the AI return on investment paradox has left leaders navigating a disconnect between capital outlay and financial returns that take much longer to materialize. That's because AI produces several distinct kinds of value that arrive on different timelines. Some benefits, such as faster task completion or improved decision-making, surface early while others, such as revenue growth, take longer to realize. The effect of the paradox is real: Capital continues to go out while the return that matters most, top-line growth, has lagged. However, the gap is starting to narrow, with surveyed companies reporting current revenue realization rising from 20% to 28% since January.

TIME-SHIFTED VALUE*



Question: *Which benefits are you currently realizing?*

Base: all respondents (N=3,500)

* Figures may not sum to 100% due to rounding.

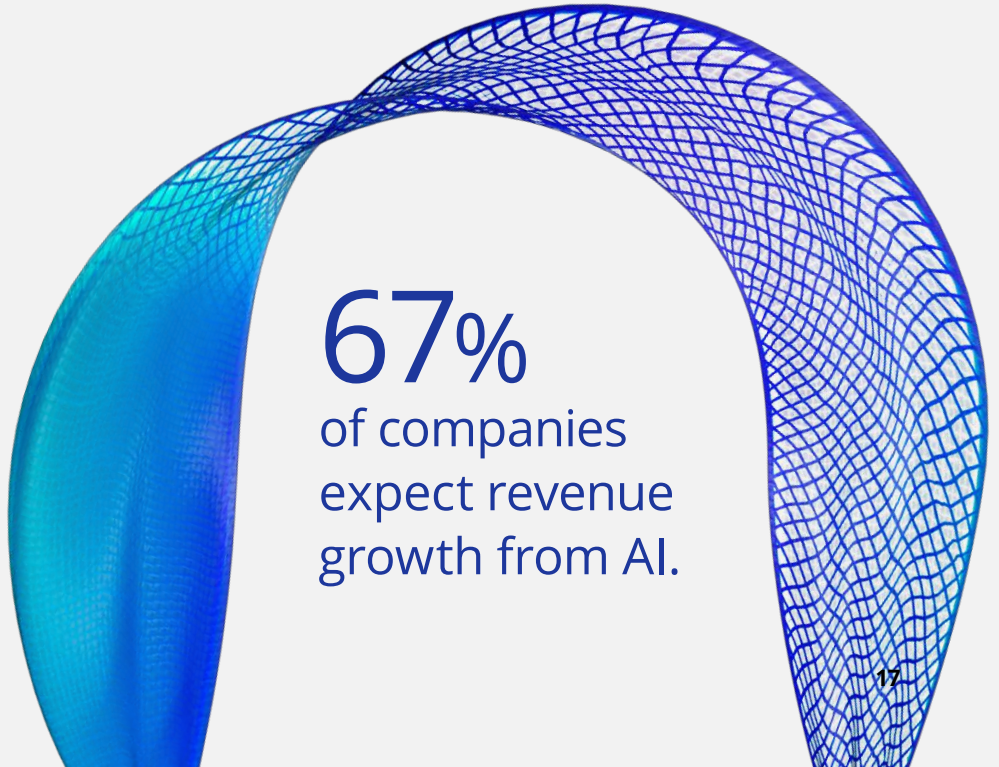
This report measures business outcomes across revenue, efficiency, cost, decision-making, customer relationships, and innovation, not enterprisewide Earnings Before Interest and Taxes (EBIT) alone. This broader view matters because AI value often appears first in specific business activities before it is visible in an aggregated financial measure.

AI-driven gains extended across every type of business benefit. Improvements in efficiency and productivity reached 74%, up from 66% in the January report. Enhanced decision-making and data-driven insight reached 60%, up from 53%. Cost reductions reached 45%, up from 40%, while improvements in client relationships reached 44%, up from 38%. Every benefit that companies pursued through AI made strides since the January report.

Revenue remains the longer-term return

The clearest indicator of what remains ahead is companies' future expectations. Revenue growth is the most anticipated AI benefit, with 67% expecting it to materialize. That places revenue ahead of every other future outcome measured, including cost reduction and improvements in products and services, customer relationships, decision-making, and risk management.

Revenue doesn't follow a fixed sequence from productivity or cost reduction. Rather, it tends to accompany broader value realization across the enterprise. Companies currently reporting revenue growth from AI also report a wider-than-average range of other AI benefits. That is thematically consistent with the pattern across this report: ROI and revenue growth occur when an organization significantly changes around the technology, not simply when the technology is deployed. From here, the task is to sustain the gains already visible while connecting AI initiatives to the commercial choices, customer outcomes, and accountability needed to translate investment into durable returns.



Key finding

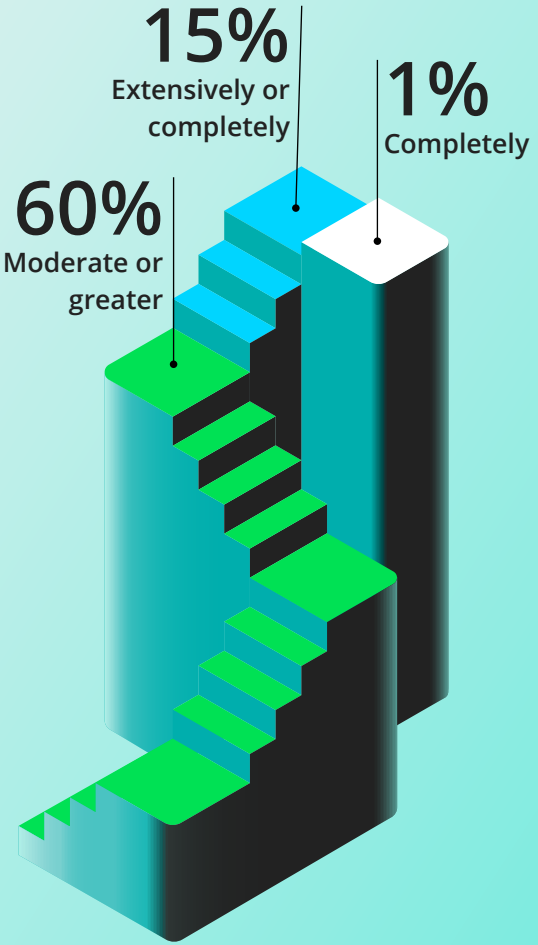
Work redesign

60% have redesigned jobs around AI, but many are still just scratching the surface.

For AI to become an integral part of the business, companies need to fundamentally rethink jobs, responsibilities, and who makes decisions. In the survey, 60% of companies have already redesigned jobs around AI capabilities to at least a moderate degree. This is a good first step. However, only 15% describe those changes as extensive or complete, a share essentially unchanged since January.

The main challenges are scope and depth. To date, many companies are primarily using AI for niche activities or general productivity. They aren't fundamentally redesigning their core operating models and workflows around AI's increasingly powerful capabilities.

JOB REDESIGN*



Question: *To what degree has your organization redesigned jobs around AI capabilities?*

Base: all respondents (N=3,500)

* Figures are cumulative, so they do not sum to 100%.

Work redesign

Similarly, only 15% of companies report a significant or fundamental shift in decision-making authority as a result of AI. Giving people AI tools can affect the quality of information they have access to, how quickly they can analyze it, and which actions they can take. But when approval processes and decision-making authority remain unchanged, companies are simply stuffing new capabilities into old operating models. Employees are being asked to work differently without being given the power to act differently. Teams armed with AI may support faster decisions, while leaders continue to make decisions at the same old pace. Legacy approval steps stay in the same places, adding delay rather than judgment.

The next phase of job and work redesign is not principally about deciding which tasks AI can perform. It is about deciding which people and systems have the authority to act on AI-enabled insight, where in the process human judgment remains essential, and when decisions should be escalated.

Authority is not just about lines of reporting. It shapes whether a redesigned role feels empowering or strained. When companies introduce AI tools but keep all the old approval steps, they can create more work rather than more capacity. Conversely, when they combine automation with clearer decision-making authority, people can spend less time gathering information and more time using their expertise, applying judgment, and handling exceptions that require human involvement.

Work and job redesign is entering a harder but more valuable stage. Many companies have started changing the work itself. Now they have an opportunity to redesign the business elements around it, including roles, authority, accountability, career paths, and the ways people and machines work together to multiply each other's value rather than simply divide the labor. Ultimately, this will separate companies that simply added AI to old jobs from those that built something genuinely new.

For example, a well-known furniture retailer that deployed an AI tool for room design found that the system did not eliminate the need for human designers. Instead, an executive said the company repurposed and upskilled designers, increasing the value of their work. The executive characterized the change as optimizing the workforce rather than reducing it.



Key finding

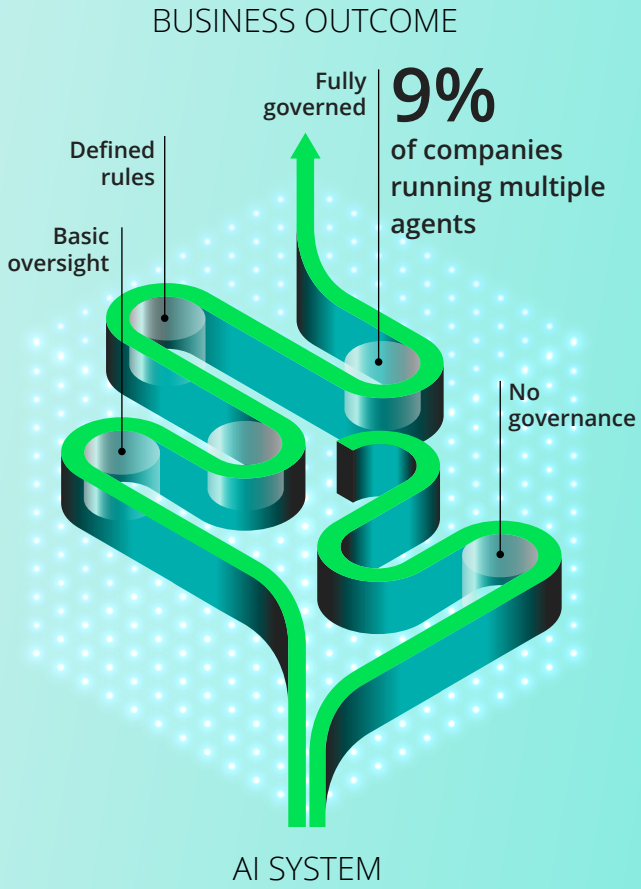
Governance as operating system

From guardrails to accountability: *Governance must be built in, not bolted on.*

As AI becomes more integral to everyday decisions and actions, the critical question is no longer simply what the technology can do. It is whether authority, accountability, and oversight are keeping pace with the level of autonomy the enterprise is actually granting to AI. Agentic AI makes that question more immediate, as systems move beyond generating recommendations to actually carrying out tasks. This report's findings show that AI is rapidly gaining more freedom to act, but governance is not always keeping pace.

Forty percent of companies already allow AI to make decisions without human approval to at least a moderate degree. But the systems needed to oversee those decisions are lagging behind. Among companies using multiple AI agents in the same workflow, half either have no formal governance or rely only on basic oversight. Just 9% say those workflows are fully governed.

GOVERNANCE OF MULTI-AGENT WORKFLOWS*



Question: *When your organization runs multiple AI agents working together on the same process or workflow, which best describes how that is governed?*
Base: companies running multiple AI agents on the same workflow (N=2,372)

Question: *To what extent is AI executing decisions or actions without human approval?*
Base: (N=3,500)

* Figures may not sum to 100% due to rounding.

The issue is not just whether governance exists; it is whether governance matches the authority companies are giving AI. A system that recommends an action does not require the same level of oversight as a system that actually carries out the action. With AI taking on more responsibility, companies need to decide which actions require human approval, which can be reviewed after the fact, and which should trigger an immediate escalation.

Companies also need to be clearer about ownership. When an AI system produces a negative outcome, 36% of companies hold the team that owns it responsible, while 24% assign responsibility to a central governance function and 10% place it with the employee who prompted the system. Nearly one-quarter, 23%, say accountability is unclear. Different ownership models may make sense for different uses, but uncertainty does not. When responsibility is unclear, problems can move between teams without anyone having the authority or obligation to resolve them.

This is where governance needs to go beyond policies and periodic review. It should be built into how AI-enabled work gets done, with clear decision rights, escalation paths, monitoring, and accountability. The goal is not to apply the same controls to every use case. It is to match the level of oversight to the level of risk and autonomy.

Done well, governance should not slow adoption; it should make responsible adoption easier and prevent decision-making from becoming a bottleneck. Clear rules help teams understand where they can move independently, when they need approval, and who owns the result, reducing the need to escalate every decision. In that sense, governance becomes part of the company's operating model, not a review added after the work is complete.

Done well, governance should not slow adoption; it should make responsible adoption easier and prevent decision-making from becoming a bottleneck.

Key finding

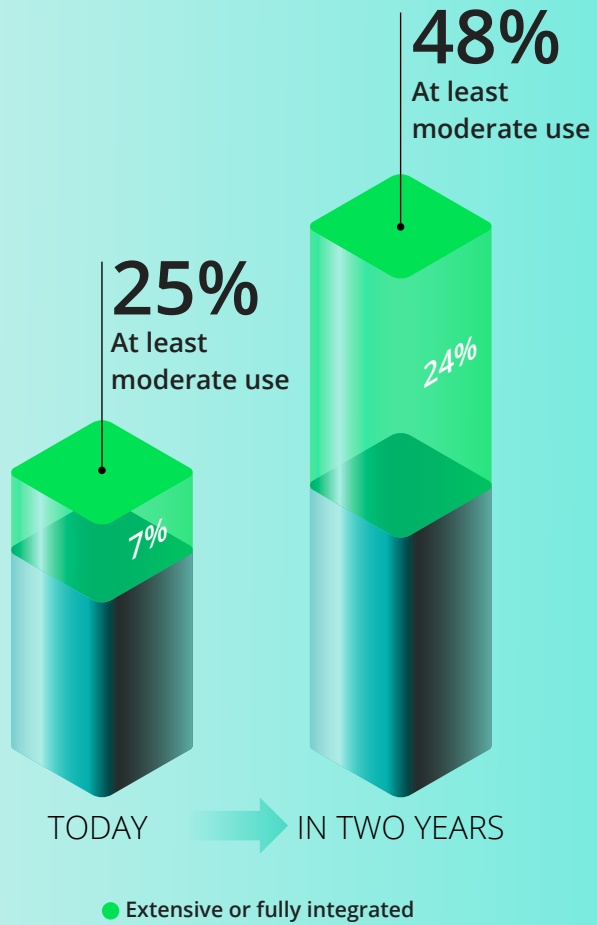
Physical AI

Physical AI is scaling. Expertise will determine who captures the value.

Physical AI systems combine AI and machine learning with sensors, controls, and robotics to perceive the real world, adapt to it in real time, and act through machines—a step beyond the preprogrammed automation that came before. Although still early in their operational maturity, these systems are expected to take on a much bigger role in companies' operations over the next two years. Today, 25% of companies report at least moderate Physical AI use, including 7% reporting extensive or fully integrated use. Within two years, however, respondents expect those shares to rise to 48% and 24%, respectively.

Asia Pacific is further along than other regions: 31% of companies there report at least moderate use of Physical AI today, compared with 25% in the Americas and 23% in Europe, the Middle East, and Africa (EMEA). The same pattern holds over the next two years, when 56% of companies in Asia Pacific expect at least moderate use, compared with 50% in the Americas and 44% in EMEA.

PHYSICAL AI USE*



Question: *To what extent is your organization utilizing Physical AI (e.g., robotics, automated machinery) in its operations? Asked for today and in two years.*

Base: all respondents (N=3,500)

* Figures may not sum to 100% due to rounding.

Physical AI

Even with that growth, adoption will not be uniform. In two years, 52% of companies expect to have no more than minimal Physical AI use, pointing to a two-speed market in which some companies are preparing to deepen deployment while many others remain at the exploratory edge.

The broader transformation horizon is longer still: Only 41% expect Physical AI to substantially transform their organization within three years, including 8% who say it already has. Another 34% place substantial Physical AI transformation more than three years away, while 25% do not expect it at all. This slower, more uneven path reflects a barrier profile distinct from other forms of AI. Cost is the most-cited barrier, at 40%, followed by unclear ROI (32%), difficulty identifying applicable use cases (31%), talent and skills gaps (31%), and infrastructure limitations (30%). That's a broader, more evenly distributed set of obstacles than Generative AI faces, where governance and data readiness dominate, or Agentic AI, where talent and governance lead.

Addressing these barriers extends the enterprise AI bottleneck into the physical world by requiring expertise that goes beyond AI and model development to include integration of sensors, robotics, control systems and real-world equipment. This unique expertise is what will set Physical AI leaders apart, and it is currently in very short supply. Today, just 19% of companies rate their own Physical AI expertise as high or very high, but that gap already shows up in the data: Such companies are far ahead on adoption, with 64% reporting at least moderate use, and 28% reporting extensive use or full integration. The head of AI at a European chemical company described a robot used to monitor plant operations as operationally capable of locating itself and interpreting its surroundings, but not intended to fully replace employees.



19%

of companies rate their own Physical AI expertise as high or very high.

Companies also trust Physical AI more in some settings than others. Among current users, 52% are comfortable letting it run without someone checking every action in a controlled setting like a warehouse or factory floor, the kind of environment Deloitte's Physical AI research consistently points to as the technology's proving ground before it moves into less controlled settings. That drops to 45% in logistics, 28% in customer-facing settings, and just 16% in health care. One infrastructure company's head of AI explained why: Full autonomy is often possible, but the company chose not to use it. Instead, they automate the repetitive parts of a job and keep a person in charge of approving the work, not doing it.

Physical AI is going to scale. The real test for leaders is whether their organization has the expertise to choose the right applications, build a credible business case, and run it confidently enough to get value out of it.



Key finding

Sovereign AI

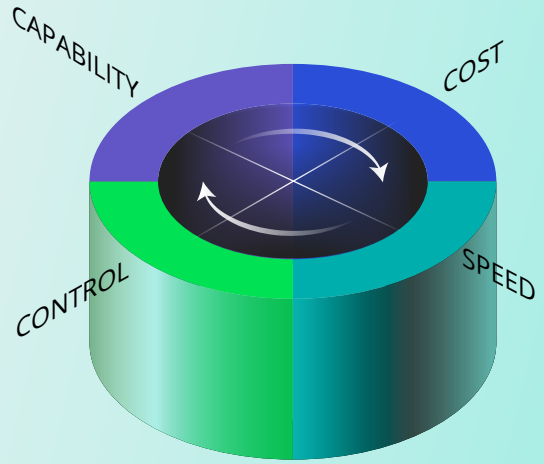
Sovereign AI stays on the strategic agenda for 74% of organizations, as they diversify where and how their AI is built and operated.

Sovereign AI continues to be a strategic concern on two fronts. Most companies say data residency and in-region compute considerations are at least moderately important to their AI strategy, and 62% are at least moderately concerned about reliance on foreign-owned AI. Leaders are navigating a practical tension between access to the leading AI technology, regardless of where it originates, and the need to retain control over the systems, infrastructure, and suppliers their AI strategies depend on. Here the fundamental bottleneck is control: not whether leading technology is available, but whether the organization can retain authority over the technology it depends on—and what it will need to give up to retain that authority.

The tension shows up directly in how executives describe their own decision-making, not just in the survey's headline numbers. One business leader said that both data sovereignty and intelligence sovereignty are active considerations, and that businesses increasingly weigh data sovereignty and the origin of AI models in their decisions.

Origin isn't an afterthought for these leaders; it is an active input into vendor decisions, weighed alongside cost and capability rather than subordinated to them.

SOVEREIGN AI IMPORTANCE AND CONCERN*





74%

Report at least moderate importance of Sovereign AI/in-region computing.



62%

Report at least moderate concern about reliance on foreign-owned AI technology and computing.

Question: *How important are Sovereign AI/in-region compute requirements to your organization's AI strategy?*

Question: *To what extent are you concerned about your organization's reliance on foreign-owned AI technologies and compute provision?*

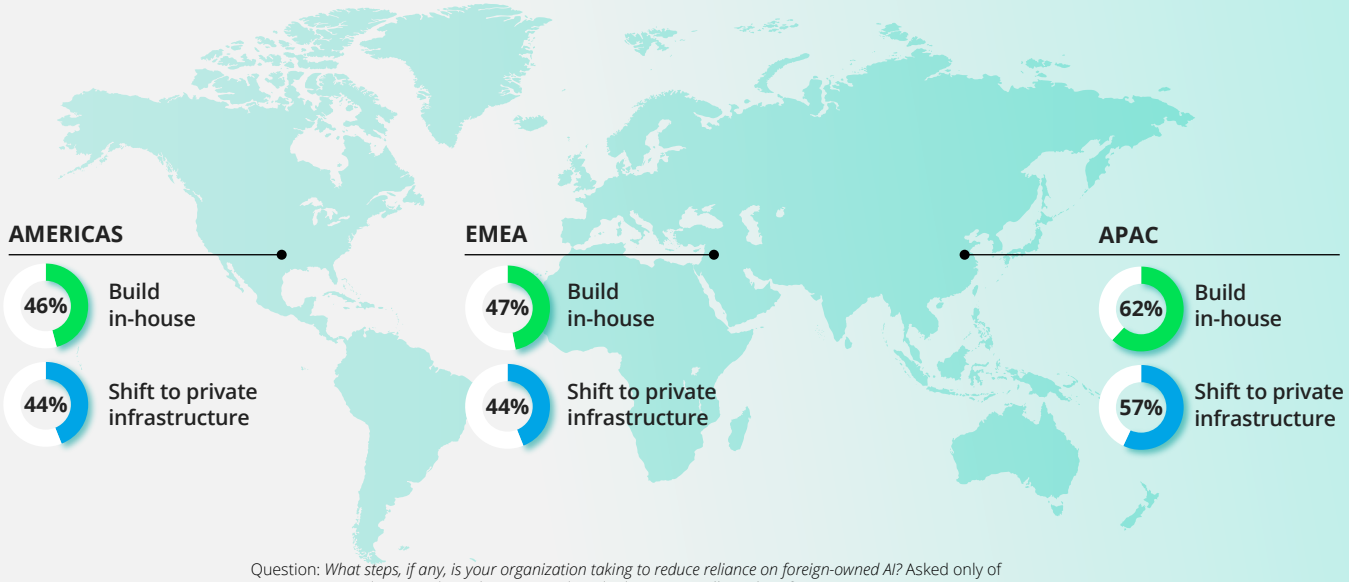
Base: all respondents (N=3,500)

* Illustrative framework based on executive interview themes, not survey results. Figures may not sum to 100% due to rounding.

Among companies concerned about foreign-vendor reliance, roughly half are prioritizing preferred-country vendors, developing in-house capability, or shifting to on-premises or private cloud, with supply chain audits and policy compliance close behind. No single approach dominates, and only 6% report taking no action at all. The mix varies by region, with Asia Pacific (APAC) companies notably more likely to build in-house capability (62%) and shift to private infrastructure (57%) than their counterparts in the Americas (46% and 44%) or EMEA (47% and 44%).

Acting on that concern isn't free, and leaders are candid about the trade-off. Depending on where they are located, some leaders are having to decide if they are willing to accept a suboptimal solution in favor of sovereignty. That's the calculation beneath the mitigation numbers: Prioritizing preferred-country vendors or building in-house capability often means giving something up in return, whether that's capability, speed, or cost, which is exactly why no single approach to sovereignty dominates the responses.

REGIONAL MITIGATION



Question: *What steps, if any, is your organization taking to reduce reliance on foreign-owned AI?* Asked only of companies at least moderately concerned; multiple answers allowed, so figures won't sum to 100%.

Base: companies at least moderately concerned (N=2,160)

Depending on where they are located, some leaders are having to decide if they are willing to accept a suboptimal solution in favor of sovereignty.

How to *build differently*



AI is advancing faster than organizations' ability to convert it into sustained business value. Closing that gap requires changes to both the operating model and the economic discipline around AI.

Framing it this way connects two challenges usually treated separately: the practical challenge of scaling AI by redesigning work, assigning decision rights, and building governance that keeps pace—and the economic challenge of managing it: Where investment creates value, how costs and ownership get allocated, and how returns get measured. The organizations that pull ahead will not necessarily be those that adopt AI the fastest. They will be the ones that pair technical progress with the management systems required to make AI productive, accountable, and economically sustainable.

Here are six key action items that can help enterprises break through the AI bottleneck and build differently:

01

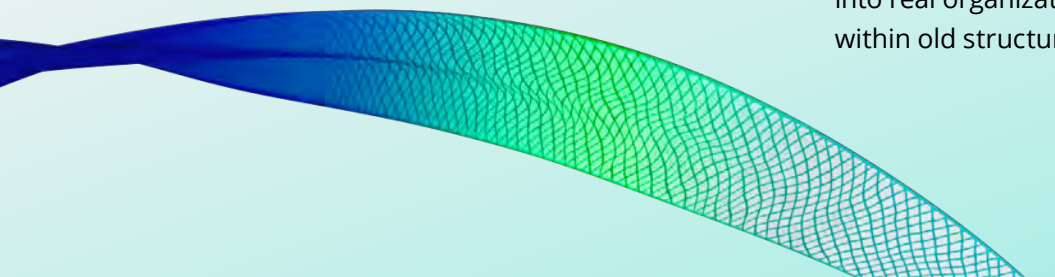
Redesign the work, not just the workflow

Pair that redesign with clear decision authority. AI can make existing processes faster, but capturing its full value requires rethinking work end to end: eliminating unnecessary steps, handoffs, and approvals; moving decisions closer to the work; and defining where human judgment remains essential. Start where AI is already producing meaningful efficiency gains, then ensure every redesigned role or team has explicit authority to act with the appropriate level of escalation. Leaders should also clarify which decisions AI can recommend, which actions people may take directly, and when approval is still required, so new capabilities translate into real organizational agility, not simply faster answers within old structures.

02

Trace cost to the token, not just the total bill

Visibility into total AI spend has arrived quickly, but visibility into what drives that spend has not. This gap will become more critical as agentic workloads scale, because long reasoning chains, repeated model calls, retries, tool use, and interactions among agents can multiply the AI compute usage behind a single business task. Organizations should be able to trace AI cost from the enterprise bill to individual workflows, models, calls, and ultimately token consumption and unit cost. That requires consistent metering and attribution of usage to the business activity creating it. It also requires a named owner for the economics of each major AI workload. With that foundation, companies can compare models, optimize prompts and workflows, route tasks more efficiently, and understand whether falling unit prices are actually translating into lower costs at scale.



03
Connect AI initiatives to the commercial decisions that produce revenue

Efficiency and cost gains have moved quickly, while revenue is following a longer timeline. Closing the gap requires connecting AI initiatives directly to the commercial decisions that create growth. For each initiative that is expected to generate revenue, leaders should be explicit about the mechanism: Is AI improving pricing? Helping to identify demand? Increasing conversion? Strengthening retention? Accelerating product development? Enabling a new product or service? Just as important, those revenue-focused initiatives need business owners with the authority to act on what AI reveals. Revenue gains are already appearing, typically where AI intersects directly with customers, products, and commercial activity. The next step is to make that connection intentional, so AI investment is tied to measurable changes in customer behavior and financial performance rather than a generalized expectation that revenue will eventually follow.

04
Match governance to the autonomy AI is actually given

AI is being given more room to act directly, but organizations are lagging behind in building the governance structures to match. One way to reduce that lag is by explicitly tying governance requirements to the authority of the system itself. Organizations should distinguish between AI that informs a decision, AI that recommends an action, AI that executes within defined limits, and AI that can act more autonomously. They can then establish the appropriate approval, monitoring, escalation, and accountability for each level. As agentic systems become more common, their controls increasingly need to operate at the point of execution through clear identities, permissions, policy enforcement, monitoring, and auditability—rather than relying solely on policies and periodic reviews. The idea is not to create more governance everywhere; it is to create governance that becomes more rigorous as autonomy and consequences increase.

05
Build Physical AI expertise ahead of the deployment curve

Physical AI adoption is on a path to expand well beyond where it stands today, but the expertise needed to deploy it effectively remains scarce. Because Physical AI combines software intelligence with sensors, robotics, control systems, safety requirements, and real-world operating environments, companies can't assume that expertise developed for Generative or Agentic AI will transfer automatically. Before deployment accelerates, organizations should begin building multidisciplinary capabilities in use-case selection, hardware and systems integration, simulation, safety engineering, reliability, and operations. They should also improve their ability to distinguish between AI capabilities that are ready for reliable operational use and those that are still being demonstrated primarily in controlled settings. Companies that develop this expertise early will be better positioned to identify where Physical AI is economically viable, integrate it into operations safely, and scale up as the technology matures.

06

Treat sovereignty as an ongoing strategic input, not a one-time fix

No single approach to managing reliance on foreign-owned AI technology and infrastructure has emerged as the clear default, because the trade-offs among capability, cost, speed, resilience, and control are real. The task is therefore less about choosing a sovereign architecture once and more about building sovereignty considerations into AI decisions over time. Which workloads and data require greater geographic or operational control? Where do critical dependencies sit across models, compute, and providers? What alternatives exist if regulations, geopolitics, supplier economics, or technology leadership change? The answers may vary for different workloads, pointing to different approaches that range from global providers or trusted regional partners to private infrastructure or in-house capabilities. The ultimate goal is not to maximize independence, but to retain sufficient control, flexibility, and optionality so the architecture can adapt as circumstances evolve.

Taken together, the findings point to an addressable bottleneck. Companies that make costs and returns visible, redesign work and decision rights together, and build governance, expertise, and control commensurate with AI's expanding role will be better positioned to convert its momentum into lasting advantage.



Acknowledgments

Authors

Business leadership



Jim Rowan
US Head of AI
Deloitte Consulting LLP
jimrowan@deloitte.com



Beena Ammanath
Executive Director
Global Deloitte AI Institute
Deloitte LLP
bammanath@deloitte.com



Nitin Mittal
Principal
Global AI Leader
Deloitte Consulting LLP
nmittal@deloitte.com



Costi Perricos
Global GenAI
Business Leader
Deloitte LLP
cperricos@deloitte.co.uk

Global AI subject matter expert network

The authors would like to thank the Deloitte subject matter specialists who contributed to the development of the survey and report:

Javier Echaniz Alonso, Marcus Goetz, Cliff Goss, Scott Holcomb, Jas Jaaj, Saurabh Kumar, Anjani Kumar Sripada, Patrick Laurent, Chris Lewin, Parth Patwari, Antonio Russo, Prakul Sharma, Rohit Tandon, Ed Van Buren, and Greg Vert.

Special thanks

Technology and research team



Vivek Kulkarni
Managing Director
US AI Transformation Leader
Deloitte Services LLP
vivkulkarni@deloitte.com



Lisa Hohener
Associate Director
Deloitte AI Institute
Deloitte Consulting LLP
lhohener@deloitte.com



Christina Clark
Manager
Deloitte AI Institute
Deloitte Consulting LLP
christinclark@deloitte.com



Takiyah Wilder
Manager
Deloitte AI Institute
Deloitte Consulting LLP
twilder@deloitte.com

Additional contributors

The authors would also like to thank the team of professionals who brought this report and campaign to life, including Vyacheslav Bogrov, Madeline Brennen, Jenny Brown, Laura Moreno Campos, Carlos Alonso Corona Gonzalez, Aaron Dunser, Artem Dzhezhela, Samantha Dullea, Victoria Estacio, Dave Faucher, Brendan Flaherty, Eric Fox, Rob Gabriel, Dorothy Geehern, Colin Gillespie, Rainiero Gonzalez Dominguez, Volodymyr Kostiuchenko, Aakarsh Kulkarni, Avani Lakhotia, Joshua Masters, Francine Umana Matarrita, Adriana Mendez, Corey Newman, Jamie Palmeroni, Sergii Pinalov, Jose Porras, Lauren Saltzman, Sakshi Saxena, Kate Schmidt, and Ximena Yglesias Herrera.

Acknowledgments

Continue the conversation

Global Deloitte AI Institute leaders



Wessel Oosthuizen
Africa
weoosthuizen@deloitte.co.za



Chris Lewin
Asia Pacific
chrislewin@deloitte.com



Dr. Elea Wurth
Australia
ewurth@deloitte.com.au



Geert Hallemeesch
Belgium
ghallemeesch@deloitte.com



Jefferson Denti
Brazil
jedenti@deloitte.com



Aisha Greene
Canada
aigreene@deloitte.ca



Jan Hejtmánek
Central Europe
jhejtmanek@deloittece.com



Martín Cabrera
Chile
mcabreraa@deloitte.com



Roman Fan
China
rfan@deloitte.com.cn



Richard Eudes
France
reudes@deloitte.fr



Björn Bringmann
Germany
bbringmann@deloitte.de



Prashanth Kaddi
India
kaddip@deloitte.com



Alfredo Maria Garibaldi
Italy
agaribaldi@deloitte.it



Tomotake Kozu
Japan
tomotake.kozu@tohmatsu.co.jp



Nicolas Griedlich
Luxembourg
ngriedlich@deloitte.lu



Sultanbek Khunkaev
Middle East
sukhunkaev@deloitte.com



Naser Bakhshi
Netherlands
nbakhshi@deloitte.nl



Tiago Durão
Portugal
tdurao@deloitte.pt



Tomas Meca Figueras
Spain
tomecafigueras@deloitte.es



Carlos Labanda
Spanish-Latin America
clabanda@deloitte.com



Antonio Russo
Switzerland
antorusso@deloitte.ch



Sulabh Soral
United Kingdom
ssoral@deloitte.co.uk



Dr. Amanda Williamson
New Zealand
amawilliamson@deloitte.co.nz

About the Deloitte AI Institute

The Deloitte AI Institute™ helps companies connect all the different dimensions of the robust, highly dynamic, and rapidly evolving artificial intelligence ecosystem. The AI Institute leads conversations on applied AI innovation across industries, with cutting-edge insights, to promote human-machine collaboration in the "Age of With™." The Deloitte AI Institute aims to promote the dialogue and development of AI, stimulate innovation, and examine challenges to AI implementation and ways to address them. The AI Institute collaborates with an ecosystem composed of academic research groups, startups, entrepreneurs, innovators, mature AI product leaders, and AI visionaries to explore key areas of artificial intelligence including risks, policies, ethics, the future of work and talent, and applied AI use cases. Combined with Deloitte's deep knowledge and experience in artificial intelligence applications, the Institute helps make sense of this complex ecosystem, and as a result, delivers impactful perspectives to help companies succeed by making informed AI decisions.

No matter what stage of the AI journey you are in: whether you are a board member or a C-suite leader driving strategy for your company—or a hands-on data scientist bringing an AI strategy to life—the Deloitte AI Institute can help you learn more about how enterprises across the world are leveraging AI for a competitive advantage. Visit us at the Deloitte AI Institute for a full body of our work, subscribe to our podcasts and newsletter, and join us at our meetups and live events.

Let's explore the future of AI together. →

Methodology

To obtain a global view of how AI is being adopted by companies on the leading edge of AI, Deloitte surveyed 3,500 leaders between May and June 2026. Respondents were senior leaders in their companies and included board and C-suite members, and those at the president, vice president, and director levels. The survey sample was split equally between IT and line-of-business leaders. Twenty-five countries were represented: United States (n=1,200), Canada (n=175), Brazil (n=115), Mexico (n=100), India (n=200), Australia (n=100), Singapore (n=75), Japan (n=100), United Kingdom (n=200), Germany (n=150), France (n=130), Netherlands (n=50), Italy (n=75), Spain (n=100), Poland (n=40), Czech Republic (n=30), Austria (n=50), Portugal (n=15), Ireland (n=85), Saudi Arabia (n=185), United Arab Emirates (n=180), Kuwait (n=45), Qatar (n=50), Oman (n=25), Egypt (n=25). All participating companies have one or more working implementations of AI in daily use. Respondents were required to play a role in their company's AI and data science strategy, investments, implementation, or value measurement. They either influenced decisions, participated in decision-making, served as the final decision-maker, or managed AI technology implementations. The survey data was supplemented with case studies and qualitative findings derived from 15 interviews with executives and AI and data science leaders at large companies across a range of industries and countries.

Deloitte applied its human-led, AI-accelerated approach to select parts of this report. AI tools, including Deloitte's internal platform Sidekick and other enterprise-grade large language models, supported early theme exploration, survey-question drafting, market research, and style and grammar review. Reviewers checked and validated all AI-supported outputs before use. Report authors made every final decision on direction, survey questions, and interpretation of findings, and final approvals followed Deloitte's brand and risk compliance process.

Endnotes

1. Deloitte, "Physical AI: The moment of acceleration," March 18, 2026.
2. Vivek Kulkarni, Maya Bodan, and Jamie Kilgour, "The work chart vs. org chart: AI-driven organizational delayering and cross-functional role fusion," *Deloitte Insights*, June 16, 2026.
3. Nicholas Merizzi et al., "The pivot to tokenomics: Navigating AI's new spend dynamics," Deloitte, January 2026.



Deloitte.

As used in this document, “Deloitte” means Deloitte Consulting LLP, a subsidiary of Deloitte LLP. Please see www.deloitte.com/us/about for a detailed description of our legal structure. Certain services may not be available to attest clients under the rules and regulations of public accounting.

This publication contains general information only and Deloitte is not, by means of this publication, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This publication is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your business. Before making any decision or taking any action that may affect your business, you should consult a qualified professional advisor. Deloitte shall not be responsible for any loss sustained by any person who relies on this publication.

Copyright © 2026 Deloitte Development LLC. All rights reserved.