AI Agenda

The Reality Check

Beyond the AI hype

ai

hope and hysteria

OLIVER WYMAN
A MARSH BUSINESS | FORUM

Contents

Foreword: It's time for a ***reality check*** on AI 02

Executive summary: AI is everywhere in our company ***except*** the P&L 03

The Reality Checks, ***at a glance*** 05

The ***return on intelligence*** (ROI) 06

Rewiring the enterprise for the ***bionic workforce*** 23

The market for ***machine attention*** 40

How ***guardrails*** can enable growth 56

The terms of ***dependence*** 70

Methodology 86

About the design: Clarity in an age of acceleration 88

Authors and contributors 89

About us 90

It's time for a *reality check* on AI

Four years into the generative AI era, everyone agrees the technology is advancing at extraordinary speed. No one seems quite sure where it has taken us, much less where it's headed next.

The AI giants have invested hundreds of billions of dollars in development and data centers, the US stock market has soared to record highs, and new AI models have leapfrogged one another with each new mind-blowing release.

A growing chorus of "AI doomers," meanwhile, is warning of internet-wide cyberattacks, an AI jobspocalypse, even the total annihilation of the human species. (The first two are no longer purely hypothetical: CEOs now rank cyberthreats among their top business risks, and AI has become a major cited reason for announced job cuts in the United States.)

And when it comes to *measurable business impact*, there is still little consensus among leaders on when (or whether) the returns will materialize. For all the hype, hope, and occasional hysteria surrounding AI, the gap between activity and results remains enormous — and the value that is emerging is highly uneven.

In this report we examine that disconnect in great detail. The research draws on more than 25,000 survey responses collected this year across 15 countries and territories, as well as the Oliver Wyman Forum and New York Stock Exchange CEO Agenda 2026 survey of 415 CEOs. It also incorporates insights from Quotient, AI by Oliver Wyman, client engagements on agentic AI, value realization, operating model change, and governance.

The research shows how AI is evolving from a story of technological innovation to a study in organizational decision-making. A new set of AI reality checks is coming into focus that will force leaders to make difficult choices in the next phase of AI.

THE REALITY IS...

Productivity without ownership isn't *ROI.* AI creates capacity, but companies need to decide who captures it

Workers won't embrace a future they believe excludes them

AI isn't only changing companies — it's also changing *customers*

Stronger governance can accelerate *growth* rather than inhibit it

AI adoption creates new *dependencies,* forcing companies to decide which ones are acceptable

Whether or not AI produces durable advantage will depend on what business leaders decide to do with that capability. We hope these pages make those decisions clearer.

John Romeo, CEO, Oliver Wyman Forum
Ana Kreacic, COO, Oliver Wyman Forum

AI is everywhere in our company *except* the P&L

That sums up the conundrum business leaders have faced since AI exploded into the popular consciousness four years ago. AI budgets are rising, pilots are multiplying, and tools are in employees' hands. But provable gains have often been delayed, diffuse, or dependent on changes outside the technology.

Across all five dimensions of this report — return on investment, workforce and culture, agentic commerce, governance and risk, and sovereignty — we found the same pattern time and again: AI capability is arriving much faster than the decisions required to turn it into value.

To succeed in the next phase of AI, organizations will need to make the difficult choices that many have deferred. Among them, leaders must:

Decide what return each AI investment is meant to produce, and who will capture it. We find that the ROI question is usually being framed too broadly, applied too narrowly, and asked too late. There is no single, magical ROI number to be discovered or defended; rather, targeted, workforce, foundational, and transformational investments all follow different paths to return and should be judged accordingly. Every major AI investment needs an ROI harvester: a senior leader accountable for turning the capacity AI creates into returns.

Decide what AI will do before asking the workforce to transform around it. We see a "Great Leadership Misalignment" that needs immediate correction. Leaders must agree on where AI is taking the enterprise and what role employees will play in that journey. If leadership doesn't write the AI workforce story, employees might assume they aren't part of it. Only 37% of executives and 11% of non-managers said their company's leadership has stated a clear AI ambition and is keeping them informed about progress. And many workers worry they will become the savings, with 68% saying they would give up part of their salary for a guarantee against AI replacement.

Decide how the business will win when machines help customers choose. AI isn't only changing companies — it's also changing the behavior of clients and consumers. Already, 48% of Generation Z consumers and 40% of millennials turn to AI first when looking for a new item without a specific brand in mind. When recommendations conflict, respondents are more likely to follow AI than any other source — ahead of friends and family, ratings and reviews, social media, or familiar brands. To reach these consumers, brands must scramble to master the developing science of what we call shopping agent optimization, or SAO.

Decide how to govern AI at the speed at which it operates. Stronger AI risk management and governance can accelerate growth rather than inhibit it. Where approved AI is easy to reach, clearly governed, and built into everyday software, 88% of workplace AI users report productivity gains, versus 52% where those conditions are largely absent. Sturdy guardrails reduce ambiguity by telling employees which use cases they can pursue, which ones require additional review, and who owns the decision. Better brakes allow racecars to move faster.

Decide which AI dependencies are acceptable, and which are worth paying to control. AI adoption has become AI dependency. No country or enterprise controls the full stack; the basic inputs — computing power, energy, models, data, identity, tools, and agent operating environments — create layers of interdependence that organizations need to manage. So far, only 35% of workers say their company has put the ability to switch AI providers in place. The practical question for companies, and nations, is where greater control over these factors justifies additional cost, and what that control looks like.

All of these reality checks boil down to the same management challenge: **How to harness a technology that is creating options faster than organizations can decide what to do with them**. The technology will keep improving whether or not leaders make those decisions correctly. The returns will not.

The Reality Checks, *at a glance*

1 ROI

AI creates capacity, not returns, which is why every major investment needs an ROI harvester

- The AI portfolio is strategy in action when investment flows to where the business must win, not where AI is easiest to deploy or returns easiest to prove
- AI often shifts the bottleneck rather than removing it, limiting value unless the broader process changes too

2 WORKFORCE

If leaders don't provide a clear workforce vision, employees may turn to their own, darker one

- AI is unbundling management, turning employees into AI managers and managers into "micro-CEOs"
- Enterprises are facing a crisis of apprenticeship as AI absorbs conventional entry-level work

3 AGENTIC COMMERCE

In the new era of shopping agent optimization, anything an agent can't read, evaluate, and buy will be unsellable

- The loyalty equation is being rewritten as AI weakens two of the oldest moats in retail: familiarity and inertia
- Sellers face a new trust question: Who authorized this agent — and to do what?

4 GOVERNANCE

A human in the loop cannot be the operating model for AI oversight

- To scale AI, risk management has to move upstream so business, tech, and risk start together
- The age of AI requires "living governance" oversight that moves at the speed of the technology

5 SOVEREIGNTY

The sovereignty premium should not be paid everywhere

- You don't need to own a frontier model to be sovereign
- Governments get more sovereignty by setting terms than by writing checks



The *return on intelligence* (ROI)

Why are measurable financial results so elusive for enterprises, and what does it take to earn them?

Authors
Aaron Fine, Ana Kreacic, Ashik Ardeshna, Edwin Anderson, Eric Czervionke, John Romeo, Simon Luong

Where is my *return* on AI?

That's what more leadership teams are asking, and it's the right question. The problem is that, with AI gains often delayed, diffuse, or dependent on changes outside of technology, the ROI discussion is often framed too broadly, applied too narrowly, and asked too late.

Too broadly: Some targeted applications of AI will have shorter and more directly attributable paths to return. Others will be foundational, with much of their AI-related impact emerging through the agents and applications built on top of them. Still others will depend on how much of a company's workforce adopts the technology, in good, bad, and often unpredictable ways. And most deployments that are truly transformational will require overcoming organizational obstacles that have built up over decades.

Too narrowly: Some AI benefits will build on others over time and are undercounted by traditional approaches. Equally important, many direct costs and risks are often left out. Workplace AI users tell us they spend 22 minutes of checking, correcting, or redoing AI's output for every hour of AI-assisted work. Leaders can overcome those issues with an expansive and flexible approach to measuring AI returns.

Too late: Counting returns after the fact is secondary, ironically. The primary goal must be to instill the discipline that gives the best chance to generate ROI, and to see early enough when it isn't so leaders can adjust course.

As difficult as the measuring issues are, they are only half the challenge. AI often moves bottlenecks before removing them, so the business must redesign around the new capability to capture the gain. And every major AI investment needs an ROI harvester: a senior leader accountable for turning the capacity AI creates into returns.

On top of all of this, there is a mix issue to grapple with. A company can make sensible AI investments one by one and still end up with the wrong portfolio. Quick-payback projects can crowd out vital foundational and transformational work, so the mix and sequence of investments should follow what the business needs to win and be adjusted as results come in.

Ultimately, the reality check that most differentiates good from bad is whether a company is making well-considered investments or undisciplined bets.

THE REALITY IS...

1. There is *no single ROI on AI*: Targeted, workforce, foundational, and transformational investments each have their own path to return
2. AI ROI is *applied too narrowly* when the calculation understates the full cost of change and mistakes the gain on the task for the return
3. The ROI question is *asked too late* when companies commit capital before testing and funding what has to be true for the return to materialize
4. AI often *shifts the bottleneck* rather than removing it, limiting value unless the broader process changes too
5. AI creates capacity, not returns, which is why every major investment needs an *ROI harvester*
6. The AI portfolio is *strategy in action* when investment flows to where the business must win, not where AI is easiest to deploy or returns are easiest to prove

By the numbers

Outsized AI returns are still rare

12%

of CEOs report AI-driven cost reduction or revenue uplift above 10%[1]

AI ROI leaders put AI to work on growth

4.4×

as likely as laggards to have deployed AI at scale in new products, services, and revenue streams (40% versus 9%)[2]

Workplace AI users spend substantial time checking output

22 min.

of checking, correcting, or redoing the output reported for every hour of AI-assisted work[3]

Redesigning roles and workflows can lead to stronger returns

Among CEOs already reporting returns above 10%,

52%

prioritized redesigning roles and workflows (compared with 33% for CEOs who cannot yet point to returns)[4]

Companies that scale AI report stronger AI ROI

Companies that have scaled AI in at least two areas are over

3×

as likely as other companies to say AI met or exceeded ROI expectations (49% versus 15%)[5]

Most AI-saved time has no enterprise destination

Only

8%

of workplace AI users report an organization-wide standard for what happens to AI-saved time[6]

Source: Oliver Wyman Forum x NYSE CEO Survey 2026 (footnotes 1, 2, 4 and 5); Oliver Wyman Forum Global Voices Survey, August 2026 (footnotes 3 and 6); Oliver Wyman Forum analysis

1. Survey questions: "What is the financial impact of your AI initiatives as of today as a percentage of your company's total cost savings?"; "What is the financial impact of your AI initiatives as of today as a percentage of your company's total revenue generation?" Base: N=415 (CEOs). Note: 12% is the 50 CEOs reporting AI-driven cost reduction and/or revenue uplift above 10% (11–20% or more than 20%).

2. Survey question: "What stage has your company reached today in deploying AI across the following areas? Select the highest stage reached for each area." Area: "New products, services, and revenue streams." Base: N=50 (AI ROI leaders); N=198 (AI ROI laggards). Note: "At scale" combines "AI-assisted tools at scale" and "Autonomous AI agents at scale." AI ROI leaders report AI-driven cost reduction and/or revenue uplift above 10%; laggards report no AI initiatives, say it is too early to assess the financial impact, or do not know. Descriptive comparison; it does not establish causation.

3. Survey question: "For every 60 minutes you spend completing work with AI's help, about how many minutes go to checking, correcting, or redoing AI output?" Base: N=7,845 (workplace AI users who gave an estimate from 0 to 60 minutes; "I cannot estimate this" excluded). Note: Mean. The survey does not establish how much of the checking burden is incremental to the pre-AI process.

4. Survey question: "Which of the following actions are you prioritizing in your workforce strategy in the next 1–2 years to drive high performance at your company? Select your top three priorities." Base: N=50 (CEOs reporting AI-driven cost reduction and/or revenue uplift above 10%); N=198 (CEOs reporting no AI initiatives, too early to assess the financial impact, or do not know). Note: Descriptive comparison; it does not establish causation.

5. Survey questions: "What stage has your company reached today in deploying AI across the following areas? Select the highest stage reached for each area."; "What is the financial impact of your AI initiatives as of today?" (item: "Return on investment"). Base: N=152 (CEOs whose companies have deployed AI-assisted tools or autonomous AI agents at scale in at least two use-case areas); N=263 (all other CEOs). Note: "Met or exceeded expectations" combines "Met expectations" and "Exceeded expectations." Descriptive comparison; it does not establish causation.

6. Survey question: "When AI saves time in your job, who usually decides what happens to that time?" Base: N=9,833 (workplace AI users). Note: 8% selected "My organization has a standard approach."

There is *no single ROI on AI*: Targeted, workforce, foundational, and transformational investments each have their own path to return

The same AI budget can fund very different kinds of investment. Each has a different path to return, a different time horizon, and a different set of things that must go right.

Targeted AI applications for a single use case are often quick wins. When AI upgrades a task or basic process the company already does, with very little else having to change, the results are usually positive. Think of the notes a contact-center agent writes up after every call. AI can draft them in seconds. The company already knows how many calls it handles and how long each writeup takes, so the saving is easy to see and count. Implemented well, the investment can pay for itself quickly. High-volume repetitive human tasks are areas where returns can be calculated easily and should inform prioritization of funding before deployment.

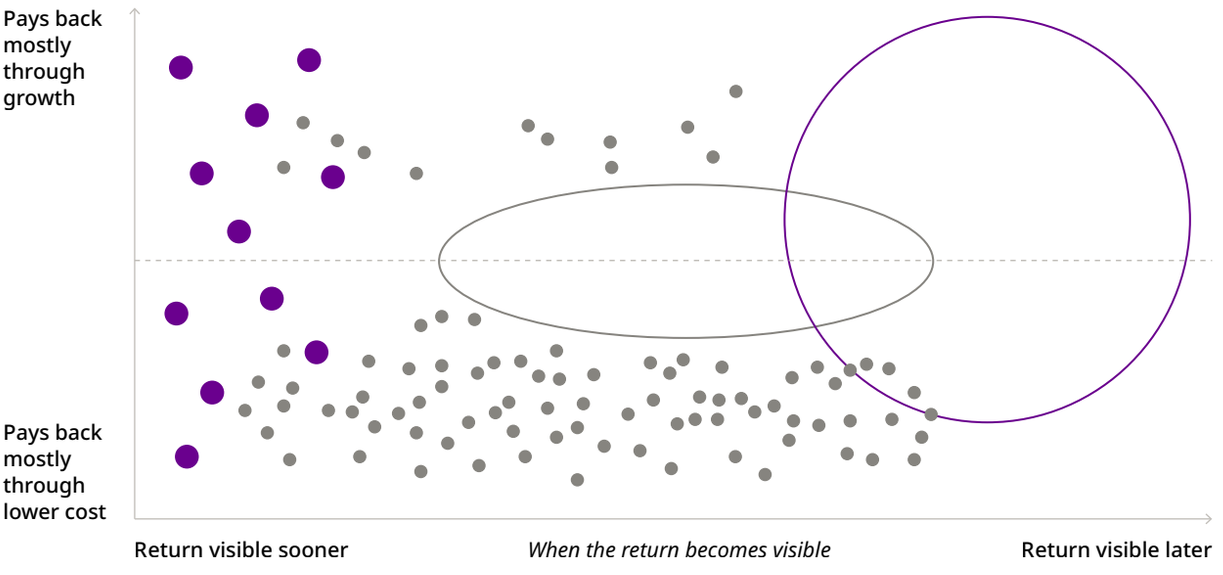
Workforce enablement is one of the hardest returns to pin down. The company gives employees AI tools and expects them to work better. This can range from engineering teams leveraging agentic coding harnesses to individual workers leveraging desktop AI tools. There is intangible value in employees learning the art of the possible as well as more tangible value in the productivity gains, if used properly. The perceived upside can be visible at companies reporting over 10% cost-reduction and/or revenue uplift from AI: 92% of employees report productivity gains, and 44% say AI enables them to do at least twice as much work (1.6 times the share among employees at other companies). But even real gains may be spread thin across thousands of people: a few minutes saved here, a better draft there. In some cases, AI expenses will outweigh the return, at least in the near term.

AI foundations are the data, processes, technology, and governance investments that other AI investments depend on. They cover things such as data capture and quality, development of enterprise approaches to deploy AI safely and cheaply, and reimagining the process required to make sure AI isn't solidifying poor practices and approaches. Many companies had identified these data and technology gaps long ago. Generative AI makes the cost of leaving them unresolved larger, and the potential payoff from fixing them greater. Much of the spending goes into repairing the underlying data and technology environment. The return shows up later through the AI applications and transformations those foundations make possible.

AI-enabled business transformation is the biggest investment and can lead to the biggest return. Transformation initiatives have the potential to produce the largest returns because they require the company to change how it operates and serves customers. This may include rethinking its offerings, how value is created, and how people and machines divide the work across a business line or the whole enterprise. That ambition also makes the return harder to deliver and slower to prove. Products, workflows, roles, technology, controls, incentives, and capacity may all have to change together. Building the AI capability can prove easier than making the organizational changes required to capture its value. The relevant question, therefore, is whether the company is making measurable progress toward its transformation objectives, not whether individual agents are being deployed. This involves retiring old processes, reallocating people and capacity, and ensuring that the potential long-term return justifies the full cost, time, and execution risk of the effort.

These four categories of investment aren't mutually exclusive. A company should be making several kinds of AI investment at once, and a single initiative can expand over time. For instance, a targeted application might reveal that the underlying data needs to be rebuilt while broad workforce adoption may eventually require roles and workflows to change. The complexity demands that leaders be clear about which type of investment the company is funding at each point, because each one should be judged against a different path to return. We explore this notion of strategic portfolio discipline later.

The same AI budget funds four overlapping kinds of investment, each with its own path to return



Targeted applications
Specific use cases with relatively direct, visible results

Workforce enablement
Broad-based productivity gains across many employees

Foundational capabilities
Shared investments that pay back through what they enable

Business transformation
End-to-end redesign and new growth with potentially larger, slower returns

Note: Positions are illustrative; overlaps show that the categories are not mutually exclusive.

Source: Oliver Wyman analysis

AI ROI is *applied too narrowly* when the calculation understates the full cost of change and mistakes the gain on the task for the return

The gain from AI is often easy to feel, but much of the cost is easy to miss. That combination can make the ROI case look cleaner than the economics really are. Consider a customer-service AI that resolves a larger share of inquiries without a human agent. In theory, this sounds promising, but if customers call back more often, escalations take longer, or remediation costs rise, some of that saving disappears. Automating more customer interactions creates value only if the actual cost and quality of resolution improve.

Many efforts to measure the return on AI fail to capture fully loaded costs associated not only with the one-time build but also the need to maintain and evolve the agentic infrastructure over time. Our employee data shows how material some of that surrounding work can be. Workplace AI users report spending an average of 22 minutes of checking, correcting, or redoing the output for every hour of AI-assisted work. This shows how the price of the AI tool alone can materially understate the cost of producing usable work.

Companies may overstate the benefit as easily as they understate the cost. Faster output creates less value if quality falls or the output doesn't lead to actual impact. If a financial case assumes AI will save employee time, the company needs to know what that time will be used to improve: output, cost, service levels, speed, or another measurable result. A credible AI ROI calculation counts the full cost of producing the result and measures the business outcome that actually changed.

The price of an AI tool leaves out the time spent checking its work

22 min.

of checking, correcting, or redoing the output for every hour of AI-assisted work

Note: Mean; the survey does not establish how much of the checking burden is incremental to the pre-AI process.

Source: Oliver Wyman Forum Global Voices Survey, August 2026; Oliver Wyman Forum analysis

Survey question: "For every 60 minutes you spend completing work with AI's help, about how many minutes go to checking, correcting, or redoing AI output?"

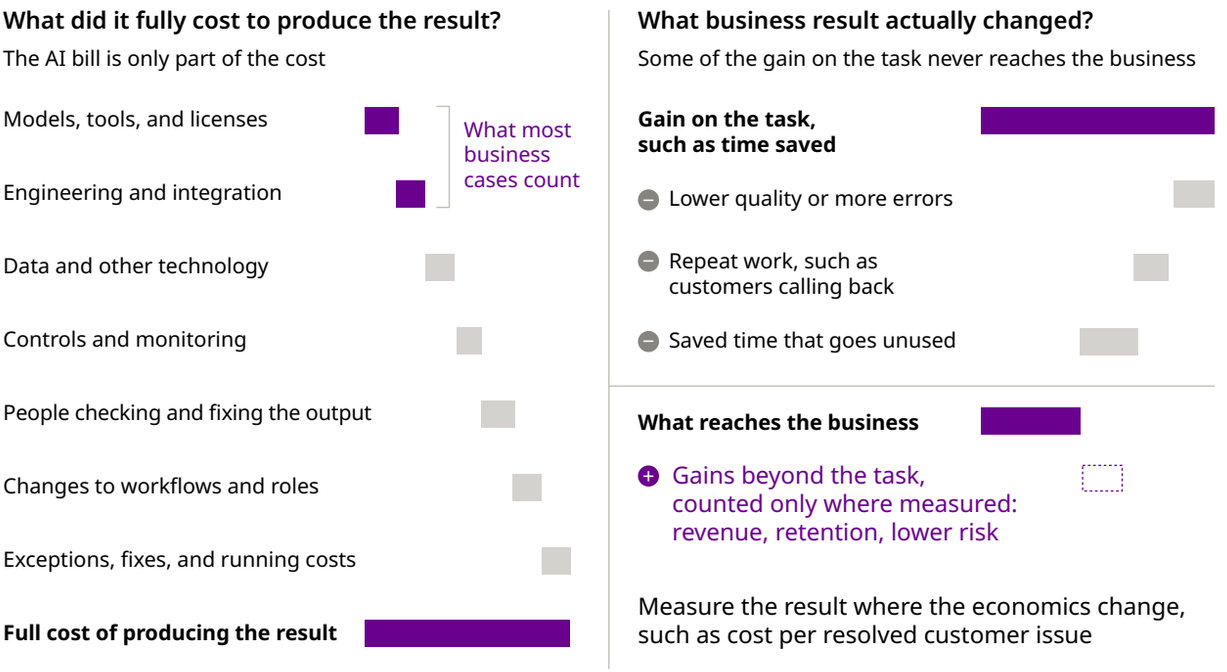
Base: N=7,845 (workplace AI users who provided an estimate from 0 to 60 minutes; "I cannot estimate this" excluded)

Financial, operational, strategic, and risk benefits don't need to be collapsed into a single number, but each *claimed* benefit needs a measure that matches the result the investment is supposed to produce. Customer service might be measured through cost per successfully resolved issue. Claims automation might be measured through loss per claim. A commercial use case might be measured through revenue or conversion. Without a defined business outcome and a way to measure it, the company cannot tell whether AI changed the economics of the business.

People also often overlook the safety and risk concerns of AI deployment, which a singular financial metric may not adequately capture.

Every AI ROI calculation should therefore answer two straightforward questions: What did it fully cost to produce the intended result? And what business result *actually changed?*

The AI bill can understate the cost, and the gain on the task can overstate the return



Note: Illustrative, not to scale. Cost categories vary by use case; not every investment incurs every cost.

Source: Oliver Wyman analysis

The ROI question is *asked too late* when companies commit capital before testing and funding what has to be true for the return to materialize

Some of the ROI work that matters most has to happen before the investment is approved.

In practice, we see many companies racing to experiment and demonstrate that they are adapting without giving full thought to where the value will truly be generated over the short and longer term. Failing to do so in advance risks leadership engaging in the practice of retroactive storytelling later on.

Instead, management should start by asking a simple question: If this investment works, what return do we expect, and what has to be true to get it? The answer usually reaches well beyond the AI tool. A new AI pricing tool, for example, may pay off only if the company first cleans up its pricing data, switches off the old pricing system, and gets sales teams to act on the tool's recommendations. The expected return plus everything that has to be true to get it is the investment's return thesis. If nobody has a credible plan for one of those requirements, the expected return should come down before the money is committed.

A return thesis has to involve much more than the finance department because change is complicated and has many downstream dependencies. Technology may own the AI capability while operations changes the workflow, human resources changes roles or staffing, risk approves new controls, and the business at large acts on the resulting capacity or customer opportunity. The investment case is credible only when the people who control those decisions understand what the return depends on and are prepared to act.

As Chapter 2 on the bionic workforce shows, the majority of executives lack a clear, shared view of where their company's AI is headed, let alone what each return depends on.

A business case for ROI can be written before anyone knows the return precisely. Management can start with a range of likely returns, release money in stages, and agree in advance what result at each stage would justify scaling up, redesigning, pausing, or stopping. Each stage should test whether one of the requirements can be met. For the pricing tool, a first stage might test whether sales teams act on its recommendations before the company commits to switching off the old system.

An experiment also can create value through learning. The company should still specify what it needs to learn, which assumption the experiment is testing, and which investment decision will change if the evidence comes back differently. By the time the investment reaches scale, management should know what the return depends on, which assumptions remain uncertain, who controls the critical decisions, and what evidence will determine whether more capital follows. That is how the company improves the odds of creating the return before it has to measure it.

AI often *shifts the bottleneck* rather than removing it, limiting value unless the broader process changes too

Removing a bottleneck doesn't necessarily improve a process. An AI claims agent might process a routine claim in seconds, for example, but the overall cycle barely improves if every claim still waits in the same approval queue. Similarly, a development team could write code much faster with AI, but testing and release still decide how quickly a product reaches customers.

Most workflows were built around the limits of people: how many hours they have, what they know, and how hard it is for them to coordinate. Approval steps, roles, controls, and incentives grew up around those same limits. When AI removes one of them, the constraint moves. The surrounding workflow, decision rights, or downstream capacity becomes the new limit on performance.

Changing how a business runs was hard long before AI. Many large change programs stall because they require people to give up familiar routines and learn new ones under shareholder pressure for short-term results. Adding AI doesn't make a transformation more likely to succeed, and in some ways it makes it harder.

AI tools improve every few months, so a workflow designed around today's capabilities might need revision sooner than expected. Employees learn new tools

Example of the AI ROI dilemma

We recently asked a gathering of about 40 health system executives whether they would fund an AI tool that screens hundreds of inpatient charts a day to flag patients who might be eligible for hospice care. The participants moved quickly past the labor saved, roughly two full-time roles, which wasn't enough on its own. Most located the value in capacity since each timely transfer frees an acute-care bed, and sometimes scarce critical-care nursing, while filling hospice capacity that many hospitals contract for but underuse. Where those beds are in demand, that alone can make the investment pay off.

Pinning down precisely how much proved harder. Estimates swung wildly with assumptions about build and running costs, how many referrals were truly incremental, and whether the hospital was capacity-constrained at all. Participants also identified a further upside that is harder to put in a business case: Deaths in hospice fall outside a hospital's risk-adjusted mortality measure, so increasing the number of patients admitted can affect reimbursement, though that effect is hard to attribute. That is the problem many face with AI ROI: The AI did the screening but ultimately, the return depended on what the rest of the hospital could do with it.

while their jobs are redesigned around them. Some workers will be slower before they are faster, and some will reasonably wonder what these tools mean for their role. Plans should allow for that learning period and revisit the design as the tools improve.

Leaders should manage an AI transformation the way they would any large change program, taking account of the company's track record. A business that has struggled to change before should expect the same struggles now and plan around them from the start. A business that has changed well before can build on what worked.

An AI transformation plan should cover the whole flow of work around the AI step. Following a single claim from the moment it is filed to the moment it is paid shows where it will wait once the AI step is fast. Those waiting points are the natural place to start the redesign. Redesigning them is what turns a moved bottleneck into a removed one.

CEOs already seeing returns are more likely to make redesign a priority. Among CEOs reporting AI-driven cost reduction or revenue uplift above 10%, more than half (52%) are prioritizing the redesign of roles and workflows, compared with 33% of CEOs who cannot yet point to returns. Scale shows up alongside stronger returns as well. Companies that have scaled AI in at least two areas are about three times as likely as other companies to say returns met or exceeded their expectations.

Redesigning the business around what AI makes possible is typically the hardest work in an AI program. The companies that take it on are better positioned to capture the return.

AI returns cluster where companies scaled specific use cases and redesign work

3×

Companies that have scaled AI in at least two areas are about three times as likely as other companies to say AI met or exceeded their ROI expectations (49% versus 15%)

52%

of CEOs reporting AI-driven returns above 10% prioritize redesigning roles and workflows, versus 33% of CEOs who cannot yet point to returns

Note: Descriptive comparisons of different groups of CEOs; they do not establish causation.

Source: Oliver Wyman Forum x NYSE CEO Survey 2026; Oliver Wyman Forum analysis

Survey questions: "What stage has your company reached today in deploying AI across the following areas?"; "What is the financial impact of your AI initiatives as of today?" (item: "Return on investment"); "Which of the following actions are you prioritizing in your workforce strategy in the next 1–2 years to drive high performance at your company? Select your top three priorities"

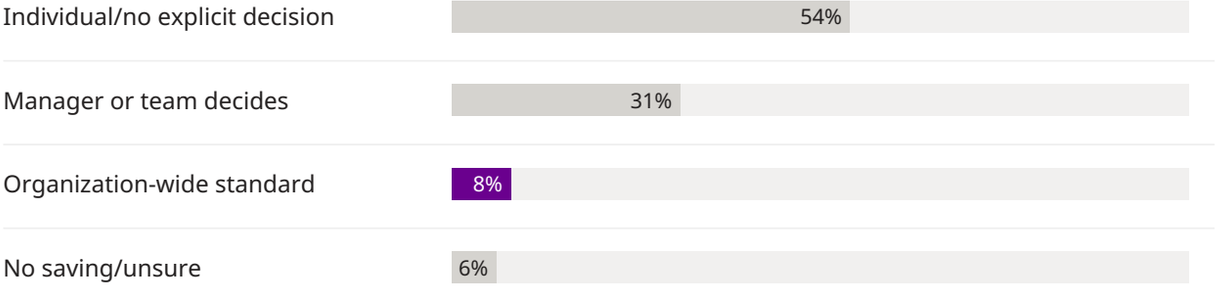
Base: Scale comparison: N=152 CEOs whose companies have deployed AI at scale in at least two use-case areas; N=263 all other CEOs. Redesign comparison: N=50 CEOs reporting AI-driven returns above 10%; N=198 CEOs reporting no AI initiatives, too early to assess, or do not know

AI creates capacity, not returns, which is why every major investment needs an *ROI harvester*

AI can create capacity without creating value. Turning one into the other requires an ROI harvester: a senior owner who is accountable from the outset for the intended business result and has end-to-end authority to make the business decisions to capture it.

Most AI-saved time has no enterprise destination

Who usually decides what happens when AI saves time, according to workplace AI users



Note: Results reflect employees' reports of how AI-saved time is allocated.
They are not organization-level incidence estimates.

Source: Oliver Wyman Forum Global Voices Survey, August 2026; Oliver Wyman Forum analysis

Survey question: "When AI saves time in your job, who usually decides what happens to that time?"

Base: N=9,833 (workplace AI users)

The harvester needs to understand where the gain is expected to come from, what the investment costs, and which management decisions will convert it. Most important, the harvester needs the authority and willingness to make those decisions when the benefit appears. The decisions might sit across several functions, but the result needs one owner attached to it.

In our imagination of the process, the harvester is involved from the start, with an active role in the design and then in the execution, because larger AI returns usually depend on a sequence of steps owned by different functions. Some steps pay off on their own, such as procurement exiting a supplier contract. Others pay off only once an earlier step is done, such as IT retiring an old system after operations has the new workflow running reliably. The harvester is accountable for the whole sequence. They act directly on the steps in their own area and work across functions for the rest to align on what is required, resolve trade-offs when functions compete, and escalate when a step stalls.

Such accountability is all too rare today. Only 8% of workplace AI users report an organization-wide standard for deciding what happens when AI saves time. The cost of AI tends to be managed centrally, while much of the capacity it creates is allocated person by person or left without a destination.

Each way of capturing the gain comes with difficult decisions, none of them new.

A **cost case** is hard because it often lands on people. Is the company prepared to slow hiring, cut contractor spend, or let people go? Can it do that effectively and humanely? A **reinvestment case** is hard because the payoff is indirect. Which higher-value work will take up the freed capacity, and who will check what it produces? **A growth case** is hard because it depends on customers. Are the business owners prepared to change the offer, the pricing, the sales process, or the customer experience?

Answering these questions, and acting on the answers, is the job of ROI harvesters. They need to make the management decisions that turn AI-created capacity into value. Measurement is the last step in the equation.

The AI portfolio is *strategy in action* when investment flows to where the business must win, not where AI is easiest to deploy or returns are easiest to prove

A company can correctly classify its AI investments, make a full cost-benefit analysis of each one, and assign the accountability for delivering returns, yet still end up with a disappointing result. What's needed is the discipline to prioritize investment in areas that are critical to the enterprise's long-term success.

We call this portfolio discipline, and it begins with leaders looking at their broad business strategy. Are we prioritizing for growth? Margin? Customer experience? Resilience? Risk reduction? All of these call for different combinations of targeted applications, workforce adoption, foundational and transformation investments. A company competing on margin might concentrate on targeted applications in its highest-volume operations, along with a process redesign to turn faster work into lower unit costs. One competing on growth is more likely to need strong customer data foundations and product and sales enhancements that bring new offers to market. A priority on resilience or risk reduction would shift weight toward governance, controls, and data quality.

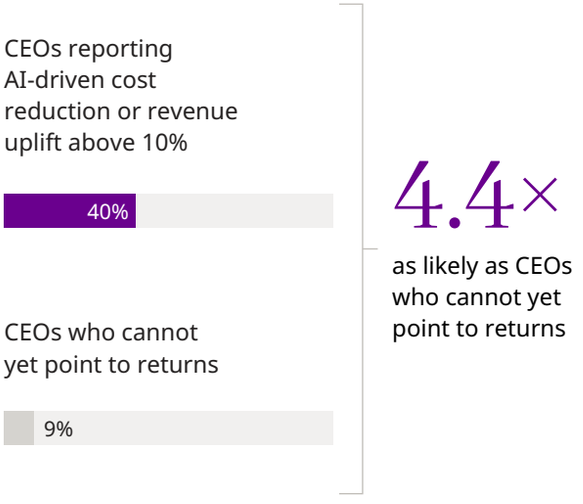
Without that strategic anchor, the portfolio tends to reflect the use cases that arrive first, have the strongest sponsor, or offer the easiest payback to demonstrate. The result is a sum of separate functional decisions rather than a coherent set of investments that compound advantage.

In practice, many companies begin with near-term use cases, and that can be helpful in mobilizing the organization toward consistent ROI. They can free up cash for later bets, build confidence, and show what the enterprise needs next. But they also can become distorting when their short payback becomes the portfolio rule. Applications with clear baselines and near-term savings can readily clear a conventional business-case hurdle. Data, governance, workforce capability, and operating-model redesign may struggle to clear that bar because their returns are more diffuse, arrive later, or show up in other teams' results. A short-term approach steers capital toward the returns that are easiest to prove quickly and away from the investments that the business strategy depends on.

Quick wins carry a secondary problem: Benefits from tools competitors can easily buy are unlikely to distinguish a company for long. Returns that create lasting advantage are more likely to come from growth, which is what our CEO data shows. Nearly four years after the launch of ChatGPT, only about one in eight CEOs surveyed report AI-driven cost reduction or revenue uplift above 10%. Yet those CEOs are 4.4 times as likely as CEOs who cannot yet point to returns to have deployed AI at scale in new products, services, and revenue streams. A portfolio of AI investments designed around easy-to-prove gains is unlikely to generate outsized returns.

CEOs reporting outsized AI returns are far more likely to have put AI to work on growth

Share of CEOs in each group whose companies have deployed AI at scale in new products, services, and revenue streams



Note: "At scale" combines "AI-assisted tools at scale" and "Autonomous AI agents at scale." Descriptive comparison; it does not establish causation.

Source: Oliver Wyman Forum x NYSE CEO Survey 2026; Oliver Wyman Forum analysis

Survey question: "What stage has your company reached today in deploying AI across the following areas? Select the highest stage reached for each area." Area: "New products, services, and revenue streams."

Base: N=50 CEOs reporting AI-driven cost reduction or revenue uplift above 10%; N=198 CEOs reporting no AI initiatives, too early to assess, or do not know

Putting AI to work at scale in new products and services rarely rests on a single investment. A new AI-enabled offer usually needs data that can support it, people who can sell and service it, and changes to how the business runs, so the four investment types have to be planned in unison. They can create significant value by working together. A targeted application can reveal a data gap that several other uses share. An investment in foundational capabilities can lower the cost and risk of many later deployments. Workforce enablement can expose where roles and decision rights need to change, and a transformation often combines all three and more. Treating these investments as independent understates the value some create by making others cheaper, faster, or possible at all. It also duplicates foundation spending as teams build their own versions, and produces pilots that work in isolation but cannot scale. Crediting this enabling value correctly is important.

Interdependence can look like a case for funding everything at once, since each investment seems to need the others. But doing so spreads capital thin, stretches the teams that have to absorb the change, and delays the evidence the portfolio needs for its next decision. Interdependence is better read as a reason for restraint and sequence: Work backward from the priority outcome to the dependencies that block it, and fund them in the order the evidence requires.

Take a company that wants to grow through AI-enabled offers. It might first build only the customer data that a few priority offers need, in a form later offers can reuse. That avoids two opposite failures: a platform built far ahead of demand, and the same data rebuilt for every use case. Targeted launches in one channel would then test whether revenue or conversion moves. Wider use across sales and service teams would follow once workflows and controls can absorb the

volume. Transformation capital to redesign the commercial model, for instance, would flow as the evidence strengthens. Each step draws on a different investment type and settles a different question: whether the data can support the offer, whether customers respond, whether the organization can absorb the change, and whether the larger prize is real. Real sequences are rarely this tidy; a competitor's move or a regulatory deadline can force the pace in one area. Even so, each stage should settle something that affects the next capital decision.

That is why AI portfolio construction is a strategic capital-allocation problem. Each investment needs a strategic objective, category, horizon, dependencies, owner, evidence milestones, and a date for the next scale, redesign, pause, or stop decision.

Reviews allow leaders to move capital as evidence changes. They can stop pilots that miss the next gate, accelerate enabling assets that several bets use, and protect longer-horizon investments while still demanding proof of progress. For those longer bets, proof usually consists of leading indicators agreed on in advance, such as data quality, adoption, or cycle time in the redesigned process. Agreeing on the reallocation rules before results come in also makes it easier to stop a well-sponsored project without reopening the debate.

The CEO and board set the outcomes, risk appetite, and time horizons. Management turns those choices into the mix, sequence, and reallocation rules, and ROI harvesters own the return on each investment the portfolio funds.

Ten steps to a credible return on AI

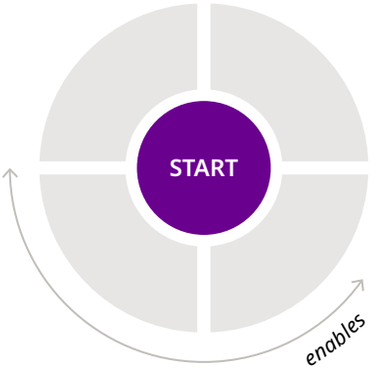
A guidebook to AI ROI

A. CONSTRUCT THE INVESTMENT PORTFOLIO

The investment portfolio

Targeted applications
Specific use cases with relatively direct, visible returns

Foundational capabilities
Shared investments that pay back through what they enable



Workforce enablement
Broad-based productivity gains across many employees

Business transformation
End-to-end redesign with potentially larger, slower returns

Start from where the business needs to win:
growth, margin, customer experience, or risk and resilience

1. Determine the portfolio mix
What mix of investments, and how much capital for each, best supports the business priorities?

2. Sequence investments by dependency
Which investments enable others or need to come first?

B. DESIGN EACH INVESTMENT

3. Assign end-to-end accountability
Who owns the result from investment through realization, with enough authority to drive the required changes?

4. Define what will work differently
Which end-to-end process, decision, or customer journey should work differently if this succeeds?

5. Define the business outcome
What value should that change create or enable, and how will it be measured?

6. Define what has to be true
What must AI do, and what else must change in data, process, technology, people, and controls for the value to materialize?

7. Define the return case
What return should this investment create or enable, over what time horizon, after full cost and risk?

C. PROVE, CAPTURE, AND DECIDE

8. Prove the economics
What business outcome changed, against what credible baseline, and what did it fully cost?

9. Capture the value
What management action converts the gain into an economic result (such as lower cost, more output, or growth)?

10. Make the funding decision
Should this investment scale, be redesigned, or stop so its capital can be reallocated?

Revisit AI investment portfolio decisions and mix

Source: Oliver Wyman analysis



2

Rewiring the enterprise for the *bionic workforce*

As humans and machines work side by side, a new architecture is required for leadership, management, and performance

Authors

Ana Kreacic, Helen Leis, Rupal Kantaria, Simon Luong, Terry Stone

The modern workforce has been rebuilt this extensively *only a few times before* – and rarely this quickly

The Industrial Revolution moved work from farms and workshops into steam-powered mills in the late 1700s. A century later, electrification and the rise of large corporations transformed work around assembly lines, defined jobs, and management hierarchies. Those shifts took decades to diffuse. AI is moving on a much faster cycle: Capabilities can advance in months and spread through an organization long before roles, skills, career paths, and management systems have had time to adjust.

The dawn of an AI world heralds a new and more urgent imperative for transformation, as machine intelligence becomes an active participant in the workforce. The technology is beginning to reshape the very nature of work: how it is broken down, who directs it, where judgment sits, what managers are for, and whether a task is performed by a human or a machine.

That makes AI an opportunity for organizational transformation as much as technological innovation. Leaders have to agree on where AI is taking the enterprise and what role employees will play on that journey. At the same time, management is being unbundled, judgment and ownership are moving up and down throughout the hierarchy, the career ladder that once built expertise is thinning, and a small group of emerging "AI superworkers" is showing what human-plus-AI performance can look like. This moment requires full organizational commitment, from top to bottom. There is no partial path to success.

THE REALITY IS...

1. The workforce transformation starts by fixing the *Great Leadership Misalignment*
2. If leaders don't provide a *clear workforce vision*, employees may turn to their own, darker one
3. AI is *unbundling management*, turning employees into AI managers and managers into "micro-CEOs"
4. Enterprises are facing a *crisis of apprenticeship* as AI absorbs conventional entry-level work
5. A new performance frontier is emerging: the *AI superworker*
6. The AI workforce advantage is cultural: making *exceptional performance repeatable*

By the numbers

Most executives say their company is ready for AI, yet most lack a shared direction

78%

of white-collar executives say their company is ready for AI at scale[7]

37%

report a clear AI ambition with progress shared[8]

Fear of replacement has employees hiding AI skills and eyeing job insurance

62%

of daily AI users have downplayed how good they are with AI at work, to avoid extra work, scrutiny, or looking replaceable[9]

68%

of employees would give up part of their salary for a guarantee never to be replaced by AI; 47% would give up 10% or more[10]

People frequently now learn at least as much from AI as from their manager

80%

of daily AI users say AI has taught them as much as or more than their manager in the past 12 months[11]

Junior learning work is moving to AI faster than employers are replacing it

Nearly 3 in 4

white-collar middle managers and above say AI now does tasks junior employees used to learn from[12]

1 in 4

of those who see this shift say their employer has fully rebuilt how junior employees learn the same skills[12]

Superworkers exist, and they cluster where employers support AI use

27%

of workplace AI users say AI at least doubles the work they can complete[13]

2×

as many superworkers emerge when five or six employer conditions are in place as when two or fewer are[14]

Companies seen as ahead show clearer direction and stronger intent to stay

10×

as many employees report a transparent, shared AI ambition versus those at companies seen as behind[15]

73%

intend to stay two years or more, versus 48% at companies seen as behind[15]

Source: Oliver Wyman Forum Global Voices Survey, April 2026 (footnotes 7, 8, and 13–15) and August 2026 (footnotes 9–12); Oliver Wyman Forum analysis

7. Survey statement: "I believe my employer is ready/equipped to adopt AI at a large scale." Base: N=659 (employed respondents who describe their job as white-collar and their level as executive management). Note: 78% is the share who agree.

8. Survey question: "To what extent has your company's leadership stated a clear ambition on the AI strategy?" Base: N=659 (as footnote 7). Note: 37% selected "Extensive; my company has provided a clear vision and has kept the employees informed on progress"; 31% selected "Moderate; my company has provided a clear vision but communication on progress has been minimal"; 31% selected "Limited; vague communications about the vision for AI" or "Not at all; no statement or vision on AI."

9. Survey question: "Have you ever downplayed how good you are with AI at work, to avoid extra work, scrutiny, or looking replaceable?" Base: N=3,429 (employed respondents who use AI in their job daily). Note: 62% is the share answering "Yes, often" or "Yes, occasionally." Among all workplace AI users (N=9,833) the share is 52%.

10. Survey question: "How much of your salary would you give up to work somewhere that guaranteed never to replace people with AI?" Base: N=12,279 (employed respondents). Note: 68% chose any amount above 0%; 47% chose 10% or more.

11. Survey question: "Who has taught you more useful things in the past 12 months?" Base: N=3,429 (daily workplace AI users). Note: 80% answered "AI" (47%) or "Both about equally" (33%). 11% answered "My manager."

12. Survey questions: "At your workplace, is AI now doing the kinds of tasks that junior employees used to learn on the job?"; "Where AI has taken over tasks that junior employees previously learned from, has your employer created other ways for them to build the same skills?" Base: N=4,686 (white-collar employed respondents at middle management level or above); N=3,446 (those answering "A lot" or "Some"). Note: 74% answered "A lot" or "Some"; 27% say their employer has fully done so.

13. Survey question: "With AI, how much more work can you get done compared to without it?" Base: N=9,836 (workplace AI users). Note: 27% answered "About double (2x)," "2–5x more," or "More than 5x."

14. Six conditions: leadership has stated a clear AI ambition and keeps employees informed; access to the best available AI tools; organization is quick to adopt new AI; clear rules on AI use; employer training has enhanced AI use; easy access to the AI tools employees want. Base: N=9,836 workplace AI users, of whom 3,778 report 0–2 conditions and 2,907 report 5–6. Comparisons are descriptive and do not establish causation.

15. Survey question: "How would you describe your company's progress in adopting AI compared to others?" Base: N=1,488 (employed respondents answering "Ahead of the competition"); N=1,793 ("Behind the competition"). Clear, shared AI ambition: share selecting "Extensive" for the question in footnote 8. Intend to stay: share scoring 6 or 7 on a 7-point agreement scale for "I intend to stay at my company for at least the next two years."

The workforce transformation starts by fixing the *Great Leadership Misalignment*

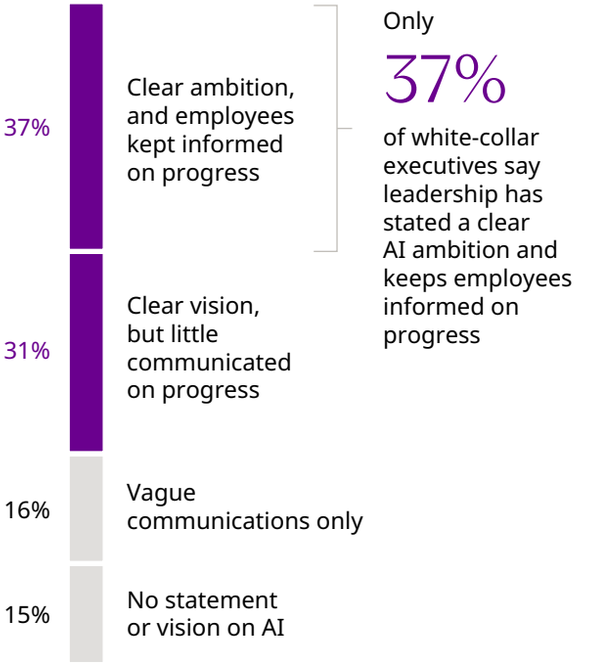
Every AI strategy needs to begin with a decision about what the technology is for. Is the ambition primarily growth and innovation, better customer outcomes, margin expansion, or doing the same work with fewer people? Many organizations have embraced AI before taking those decisions or making them explicit enough to guide the workforce. We call this the Great Leadership Misalignment, and it starts at the very top.

In our conversations with senior executives, the misalignment often originates inside the C-suite itself. A CEO and CFO might disagree on whether AI is primarily a cost lever or a growth lever, or on how fast to move. Closing this gap is one of the first imperatives of the AI era, because the workforce cannot move coherently toward a future that leadership hasn't clearly defined.

Even among the people closest to the strategy, a shared sense of direction is far from universal. Only 37% of white-collar executives say their company has stated a clear AI ambition and keeps employees informed about progress. An additional 31% see a clear vision but say little is communicated about progress. The remaining third describe the vision as vague or absent altogether.

Even at the top, a clear AI direction is a minority view

Characterizations by white-collar executives of their company's stated ambition on AI



Source: Oliver Wyman Forum Global Voices Survey (April 2026)

Survey question: "To what extent has your company's leadership stated a clear ambition on the AI strategy?"

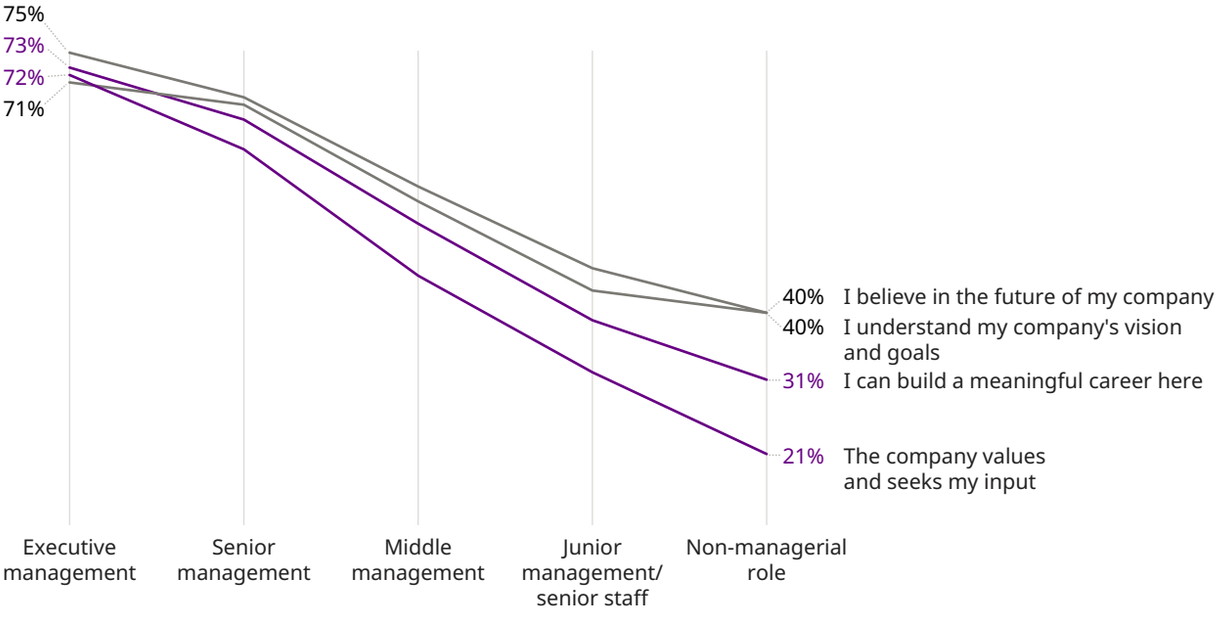
Base: N=659 (employed respondents who describe their job as white-collar and their level as executive management)

Until the strategic choices are sufficiently clear, thousands of local AI decisions can move in different directions, producing activity without transformation. The uncertainty at the top compounds as it works its way down through the rest of the enterprise. Only 11% of non-managers see a clear AI direction, roughly a third the level of their bosses.

The trickle-down confusion extends beyond AI. A transformation that asks every employee to change how they work depends on employees believing the strategic direction is real, understanding what objectives it's intended to achieve, and seeing where they fit in the big picture. When the leadership signal weakens before it reaches the people doing the work, employees begin drawing their own conclusions about where the enterprise is headed with AI and what the future holds for them.

Only 11%
of non-managers see a clear AI direction, roughly a third the level of their bosses

Confidence in the company falls sharply with every step down the hierarchy
White-collar employee sentiment on their workplace, by job level



Note: Respondents selecting "Other" job level are excluded.

Source: Oliver Wyman Forum Global Voices Survey (April 2026)

Survey question: "Please indicate your level of agreement with the following statements" (7-point scale). Strongly agree = 6 or 7

Base: N=8,038 white-collar employed respondents: executive N=659; senior N=1,163; middle N=2,759; junior management/senior staff N=1,583; non-managerial N=1,874

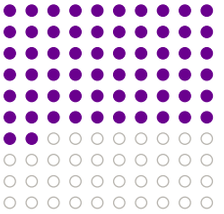
If leaders don't provide a clear workforce vision, employees may turn to their own, *darker* one

Leaders need to rally the troops. That means defining an objective, setting a strategy for achieving it, and making people feel like they have a real stake in the outcome. If business leaders fail to provide a clear vision of what AI means for staff and how they can share in the gains the technology may throw off, employees create their own and act in ways that don't help the enterprise. They hide capability and AI use, and assume productivity will translate into headcount reduction. The result can be organizational chaos as individually rational behaviors produce collectively dysfunctional outcomes.

Among daily workplace users, people who should be spearheading the organization's transformation, 62% have played down how good they are with AI to avoid extra work, scrutiny, or looking replaceable. And 68% of employees would give up part of their salary to work somewhere that guaranteed never to replace people with AI; nearly half would give up 10% or more. The employees companies most need to find are hiding.

Employees are already protecting themselves from AI by hiding their skills and offering to pay for job security

MOST DAILY AI USERS HAVE PLAYED DOWN THEIR AI SKILLS

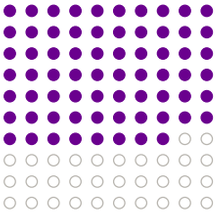


62%

of daily AI users have played down how good they are with AI at work, to avoid extra work, scrutiny, or looking replaceable

52% of all workplace AI users have done so

MOST EMPLOYEES WOULD TRADE PAY FOR PROTECTION FROM AI



68%

of employees would give up part of their salary to work somewhere that guaranteed never to replace people with AI

47% would give up 10% or more; the median employee would give up 5%

Source: Oliver Wyman Forum Global Voices Survey (August 2026)

Survey questions: "Have you ever downplayed how good you are with AI at work, to avoid extra work, scrutiny, or looking replaceable?"; "How much of your salary would you give up to work somewhere that guaranteed never to replace people with AI?"

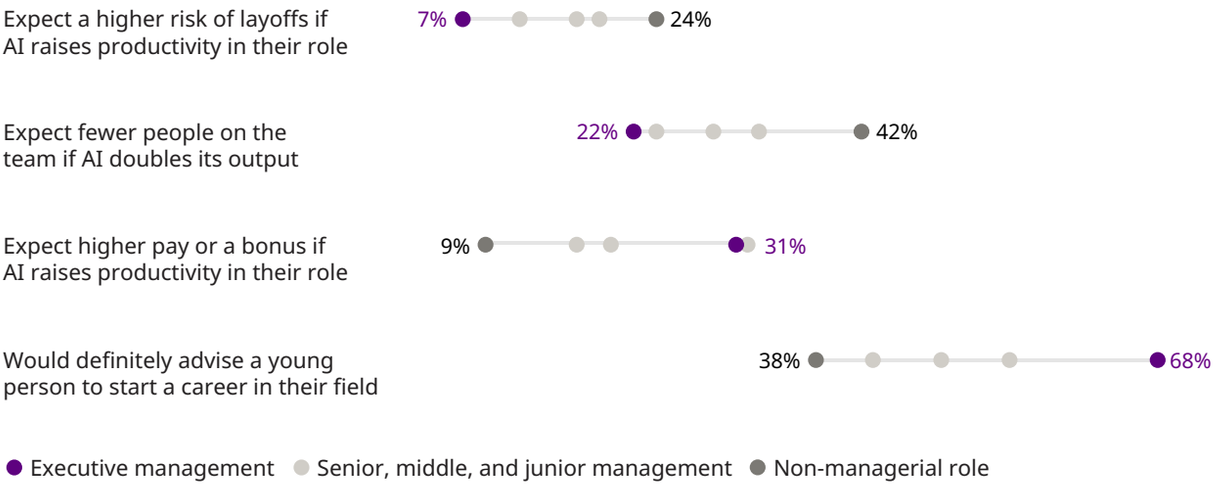
Base: N=3,429 daily workplace AI users for 62%; N=9,833 all workplace AI users for 52%; downplayed skills = "Yes, often" plus "Yes, occasionally"; N=12,279 employed respondents for the salary question; salary metrics: 68% chose any amount above 0%; 47% chose 10% or more; median response = 5%

Currently, executives and frontline employees see very different futures. Asked what will happen if AI raises productivity in their role, executives expect to be rewarded for it, with 31% anticipating higher pay or a bonus and only 7% a higher risk of layoffs. On the front line, the figures are reversed, at 9% and 24%. Same technology, same companies, opposite expectations.

Digging deeper into employee expectations reveals the extent of the challenge leaders face. Asked what should happen if AI were to double their team's output, employees rank a shorter week, higher pay, and time for training and new ideas at the top of the list, and fewer people on the team last. Asked what they expect will actually happen, the order inverts, with "fewer people" ranked first. Employees understand that the productivity dividend exists, and they are concluding it's unlikely to be shared with them.

Executives and employees expect different futures from the same technology

Anticipated effects of AI on the workforce cited by white-collar employees, by job level



Source: Oliver Wyman Forum Global Voices Survey, April 2026 (rows 1 and 3) and August 2026 (rows 2 and 4)

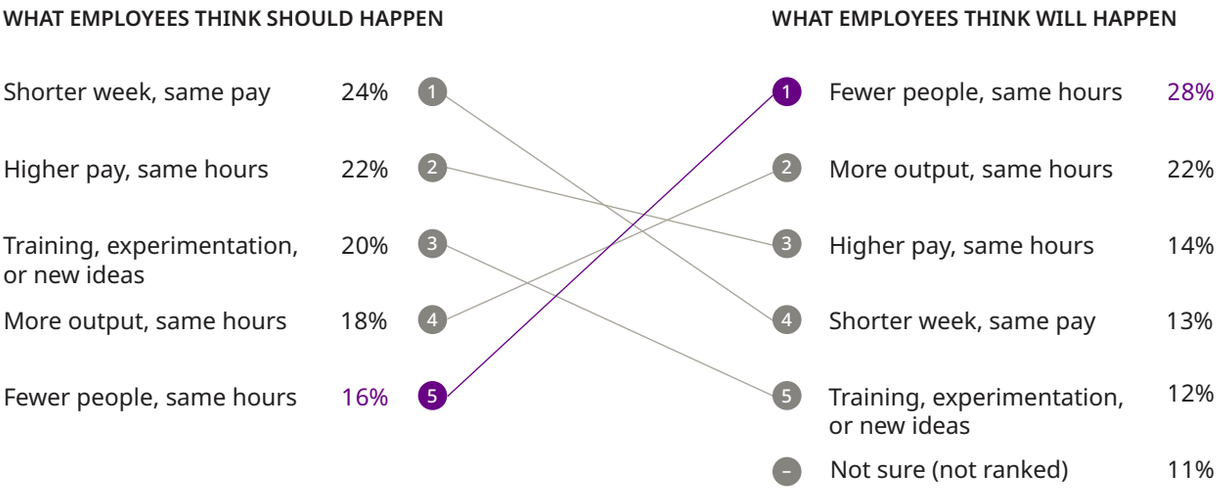
Survey questions: "If AI increases productivity in your role, what do you think will happen over the next year?" ("Higher risk of layoffs"; "Higher pay/bonus potential"); "If AI doubled your team's output, what do you think will actually happen?" ("The same hours, with fewer people on the team"); "Knowing what you know about AI, would you advise a young person to start a career in your field today?" ("Yes, definitely"); "Not sure" responses excluded

Base: White-collar employed respondents by job level; April N=7,079; August N=7,196 (young person) and N=7,139 (team output)

Employees may be reading their organizations accurately. In the United States, AI has been among the most frequently cited reasons for announced job cuts in 2026, and employees see those announcements whether or not they reflect real displacement, rather than relabeled cost-cutting.[16] The implication is that the first AI workforce intervention must belong to leaders. They have to agree on and promote a vision of what AI means for people. That story cannot be a messaging exercise. At its core, it is a bargain: what AI is for, and what employees can expect in return for actively participating in the transformation.

The outcome employees rank last is the one they expect first

Desired and anticipated effects of AI on the workplace cited by employed respondents, ranked



Note: Respondents selecting "Other" job level are excluded

Source: Oliver Wyman Forum Global Voices Survey (August 2026)

Survey questions: "If AI doubled your team's output, what should happen?"; "If AI doubled your team's output, what do you think will actually happen?"

Base: N=12,279 (employed respondents)

16. Source: Challenger, Gray & Christmas, "Challenger Report: June Layoffs Cool to 45,849, Down 53% From May; AI Leads Reasons for Fourth Consecutive Month," July 1, 2026

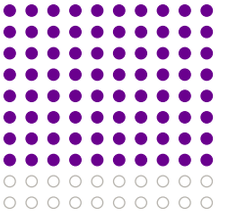
AI is *unbundling management*, turning employees into AI managers and managers into "micro-CEOs"

AI is upending the traditional division of labor, empowering rank-and-file employees and putting new responsibilities on managers.

At the bottom, a growing share of employees no longer look to their managers for instruction. Four out of five daily AI users say AI has taught them as much as or more useful things than their manager in the past 12 months. Managers, in turn, need to shift their focus to things AI cannot supply: human judgment about what matters, and empathetic development of junior talent. Those are harder tasks than the ones being taken away, and most managers were not prepared for them.

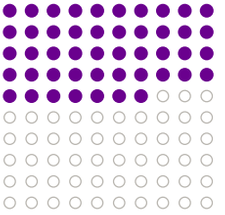
In this new scheme of things, basic components of management are becoming a growing part of junior employees' work. White-collar entry-level staff and non-managers who use AI frequently spend nearly a third of their working week, on average, managing AI tools; for more than a quarter of them, it's at least half the week. These employees are spending a growing share of their time directing work rather than doing it themselves, and doing so before they've ever had a direct report.

Is AI becoming a better teacher than managers?



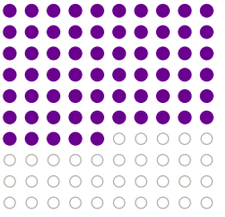
80%

of daily AI users say AI has taught them as much as or more useful things than their manager in the past 12 months



47%

say AI has taught them more useful things than their manager



65%

of daily AI users think AI could do their boss's job better today or within a few years

Source: Oliver Wyman Forum Global Voices Survey (August 2026)

Survey questions: "Who has taught you more useful things in the past 12 months?" ("AI"; "My manager"; "Both about equally"; "Neither"; "I don't have a manager"); "Could AI do your boss's job better than they do it?" ("Yes, today"; "Probably within a few years")

Base: N=3,429 (employed respondents who use AI in their job daily); boss question N=3,393 (excludes respondents without a boss)

We have often described AI as the world's smartest intern. An extraordinary intern produces extraordinary output only when someone knows how to direct it, and that responsibility is moving to every employee. All employees now increasingly direct, judge, and improve a digital workforce around them. The required skills include deciding what needs doing, specifying it clearly, judging whether the result is good enough, and owning the outcome. Organizations once developed these capabilities over a decade before entrusting someone with a team.

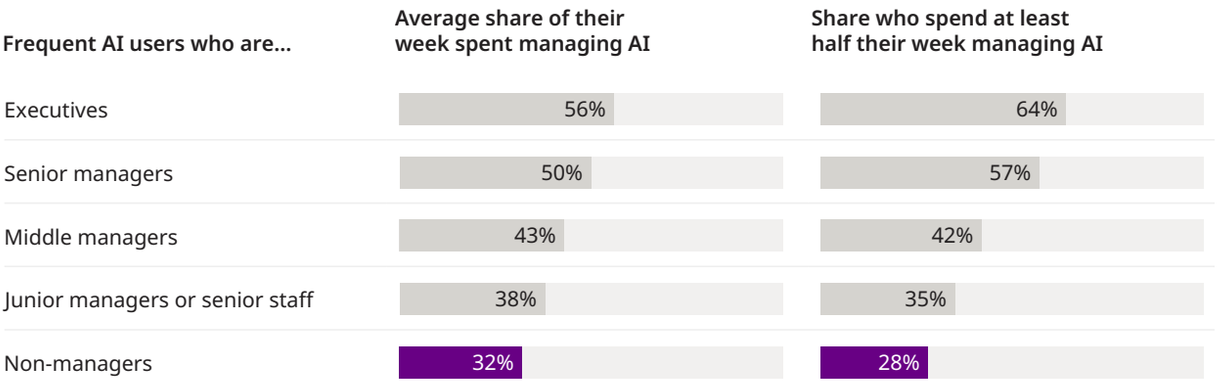
What does that imply for managers? They should begin to operate like "micro-CEOs" of agentic-powered human teams. As employees increasingly direct machine capacity and take decisions without waiting for supervision, the manager's role also moves upward: aligning human-plus-machine capacity around a common objective, resolving competing priorities, building judgment across the team, and carrying the organizational consequences of the choices that team makes. AI can help someone learn or draft feedback but it can't form human bonds, sponsor a career, or repair the morale of damaged teams.

Unbundling management also changes accountability. An employee who can direct the output of a small team can also create the risk footprint of one, from an unchecked error repeated at scale to sensitive data or intellectual property going where it should not. Organizations need explicit rules for what people can delegate to AI, what they must verify, and when they must escalate.

Few organizations have noticed this unbundling. Only one in five non-managers describes their role as having shifted toward identifying what needs to be done, and 41% say they are simply doing more of the same work. Most learned to use AI on their own: 65% of non-managerial AI users taught themselves while only 18% received formal training from their employer.

The unbundling of management therefore creates a double obligation: Redefine the manager upward, and rebuild, deliberately, the way everyone below learns to direct.

Managing AI is now part of every job, and it already fills a third of the week for non-managers
Proportion of time white-collar employees spend managing AI, by job level



Note: Respondents selecting "Other" job level are excluded.

Source: Oliver Wyman Forum Global Voices Survey (April 2026)

Survey question: "How much of your working week do you spend managing AI tools? (e.g., writing prompts, reviewing AI output, selecting which tool to use)"

Base: N=4,339 (white-collar employed respondents using AI in their job daily or 3–4 times a week): executive N=490; senior N=836; middle N=1,615; junior management/senior staff N=798; non-managerial N=600

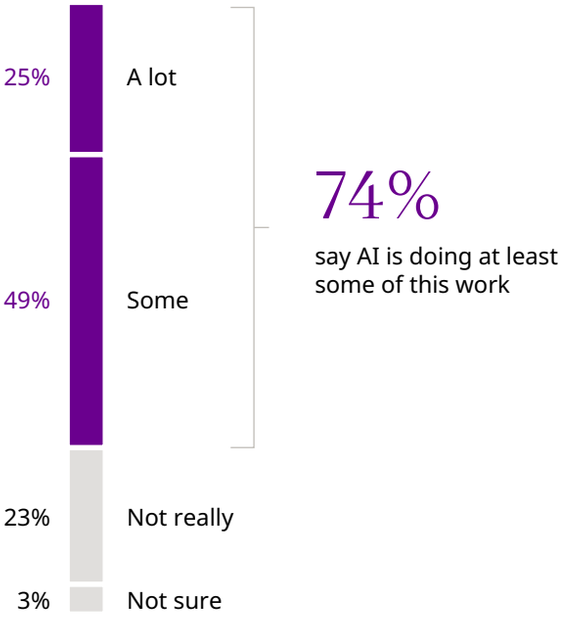
Enterprises are facing a *crisis of apprenticeship* as AI absorbs conventional entry-level work

Junior-level work has always served two purposes at once: It produces output and it produces the next generation of experienced people. Research, first drafts, coding, and analysis are among the tasks easiest to hand to AI. They are also the tasks that traditionally have taught less experienced employees how a business works.

Among white-collar respondents at middle management level or above, 74% say AI is doing at least some of the tasks junior employees used to learn from. Of this group, barely a quarter say their employer has fully rebuilt the pathways through which juniors develop the same skills. Leaders' intentions suggest a reason for that last point: 43% of CEOs who responded to our survey with the New York Stock Exchange plan to reduce junior roles in the next two years, up from 17% a year earlier. Actual reductions have so far lagged these plans, but early signs of change are visible in payroll data. Employment of 22- to 25-year-olds in the most AI-exposed occupations now stands 19% below what it would have been had it kept pace with peers in less-exposed fields.[17]

Junior learning work is moving to AI faster than employers are rebuilding how people learn

Extent to which AI is seen as performing junior white-collar employees' tasks



Source: Oliver Wyman Forum Global Voices Survey (August 2026)

Survey question: "At your workplace, is AI now doing the kinds of tasks that junior employees used to learn on the job?"; "Where AI has taken over tasks that junior employees previously learned from, has your employer created other ways for them to build the same skills?"

Base: N=4,686 (white-collar employed respondents at middle management level or above); N=3,446 (those answering "a lot" or "some" to the first question)

17. Source: Brynjolfsson, E., Chandar, B. & Chen, R., Stanford Digital Economy Lab, Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence, revised August 12, 2026. Base: ADP administrative payroll data covering millions of U.S. workers through June 2026. Note: Employment among workers ages 22–25 in highly AI-exposed occupations was approximately 19% below where it would have been had it kept pace with employment among similarly aged workers in less-exposed occupations.

Companies are trading away learning for an uncertain return. The savings from automating junior tasks are rarely captured deliberately. The cost of a weaker apprenticeship appears later, when a cohort that had fewer chances to build judgment is expected to exercise it. Entry-level roles are becoming more demanding at the same time. AI raises the skill floor while removing some of the work through which people used to gain experience. Job postings increasingly ask junior candidates for judgment and leadership once expected of senior staff.

Leadership needs to tackle the apprenticeship issue head on. Before turning work over to AI, leaders should know what the work produced and what it was teaching. That can include everything from pattern recognition and context to exception handling and judgment to maintaining relationships and accountability. Then they can build new development opportunities through simulated cases, wider scenario exposure, feedback, and faster practice. The learning architecture needs a deliberate design and a clear owner.

The challenge is whether organizations can keep producing experienced humans while machines absorb more of the work that once made humans experienced.

27%

of those who see junior learning work moving to AI say their employer has *fully rebuilt* how those employees build the same skills. 54% say partly; 15% say it has not

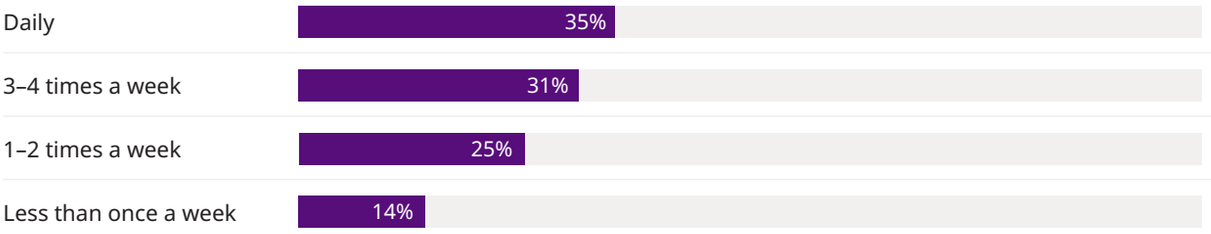
A new performance frontier is emerging: the *AI superworker*

Some people have already learned how to operate in this new division of labor, and their performance looks radically different. Twenty-seven percent of workplace AI users say the technology at least doubles the amount of work they can complete. These people are increasingly described as AI superworkers. More than a quarter of everyone using AI at work, and more than a third of those who use it daily, are superworkers. These are self-reported figures, and Chapter 1 shows how easily perceived productivity gains can run ahead of proven value. But the biggest finding is how differently they work from other employees.

27%
of workplace AI users say AI at least doubles the amount of work they can complete

One in four AI users says AI at least doubles their output

Share of workplace AI users who say AI at least doubles the amount of work they can complete, by frequency of AI use at work



Source: Oliver Wyman Forum Global Voices Survey (April 2026)

Survey question: "With AI, how much more work can you get done compared to without it?"
At least double = "About double (2x)", "2–5x more", or "More than 5x"

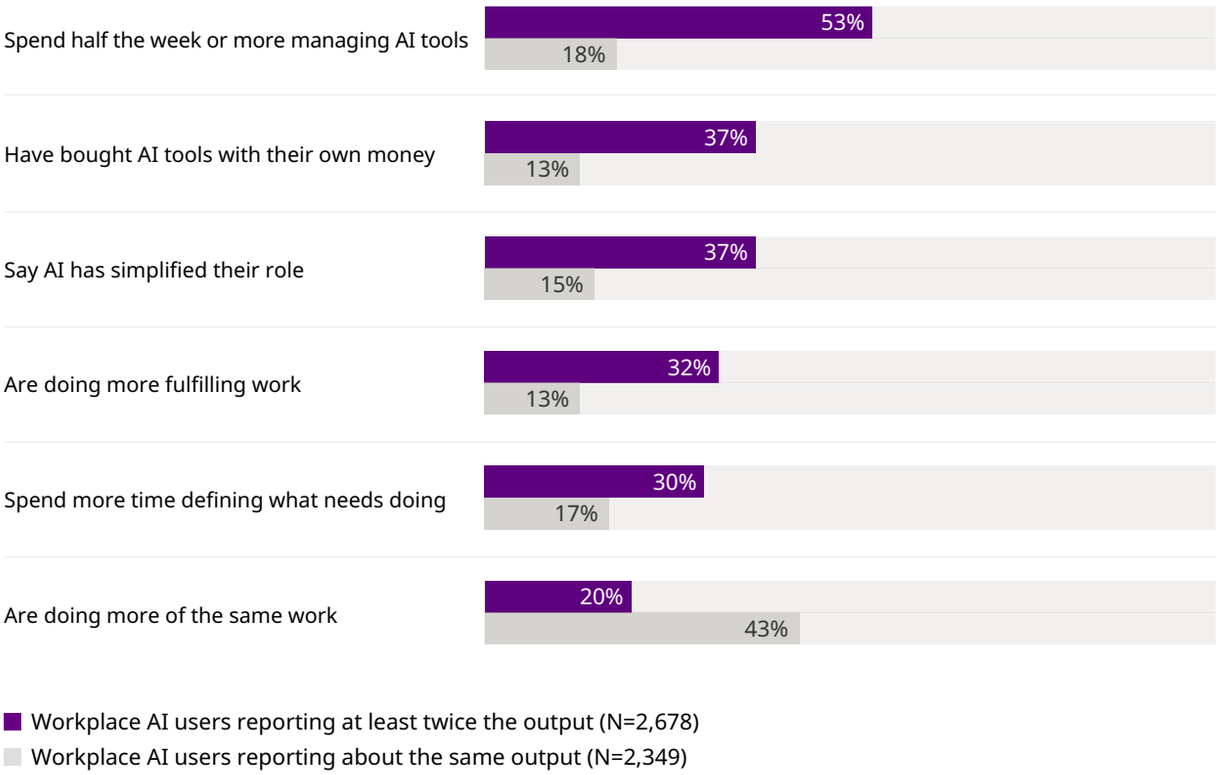
Base: N=9,836 (employed respondents who use AI in their job at least occasionally):
daily N=2,827; 3–4 times a week N=2,835; 1–2 times a week N=2,308; less than once a week N=1,866

Superworkers spend the week directing rather than repeating. More than half spend at least 50% of their working week managing AI tools, versus 18% of users who report no productivity gain. They are nearly twice as likely to say they spend more time identifying what needs to be done and less time figuring out how to do it. And they are half as likely to say they are doing more of the same work than before.

Superworker performance also creates a new retention risk. These workers are nearly twice as likely as low-gain users to say AI has fueled expectations that they do more work or take on more tasks, and they are more likely to have considered leaving over how their organization handles AI. Organizations that treat extraordinary output as an invitation to simply load on more work may quickly exhaust their most valuable AI talent.

Superworkers spend the week directing work; others spend it repeating work

Prevalence of workplace behavior among respondents who say AI at least doubles their output and those who say AI leaves it about the same



Source: Oliver Wyman Forum Global Voices Survey (April 2026)

Survey questions: "With AI, how much more work can you get done compared to without it?"; "How much of your working week do you spend managing AI tools?"; "Have you spent your own money on AI tools to use for work?"; "How has AI shifted your role or mindset within your organization?"; "Select statements that you agree with regarding AI's impact on your work-life balance and/or productivity"

Base: Workplace AI users reporting at least 2x output (N=2,678) or about the same output (N=2,349). Users reporting "up to 50% more" or "not sure" are not shown

The AI workforce advantage is cultural: making *exceptional performance repeatable*

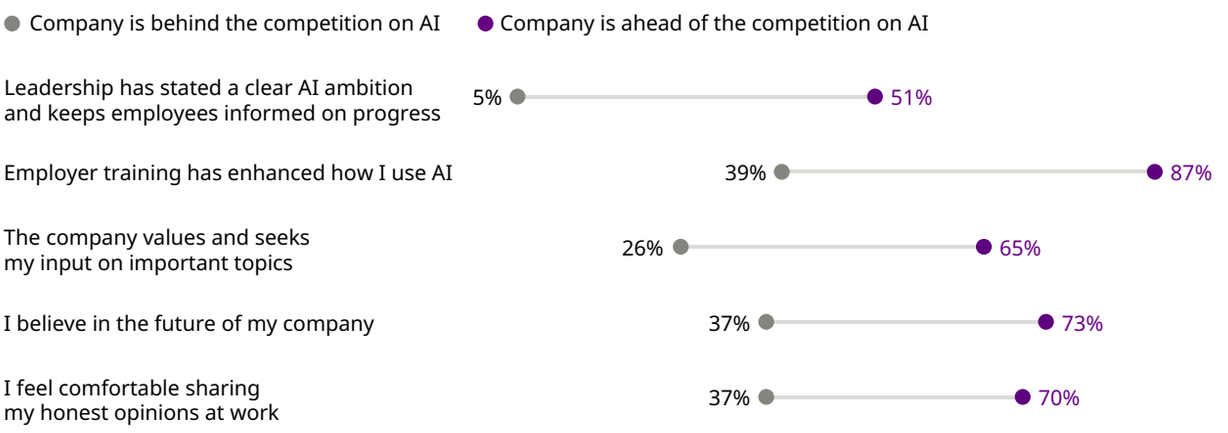
Superworkers are not simply a small group of unusually capable people who happened to figure AI out first. They are more common in organizations that give people both the capability to work differently and a reason to do so openly. These workers feel they've been given a clear direction, training that changes how they work, rules they can follow, and enough trust from leadership to experiment and say what they think.

Twelve percent of employed respondents say their company is ahead of the competition on AI; 15% say it is behind. The two groups describe very different places to work. At companies employees regard as being ahead, 51% say leadership has stated a clear AI ambition and shares progress; at companies seen as behind, only 5% do. That's a tenfold gap in perception, the widest of any trait we measured.

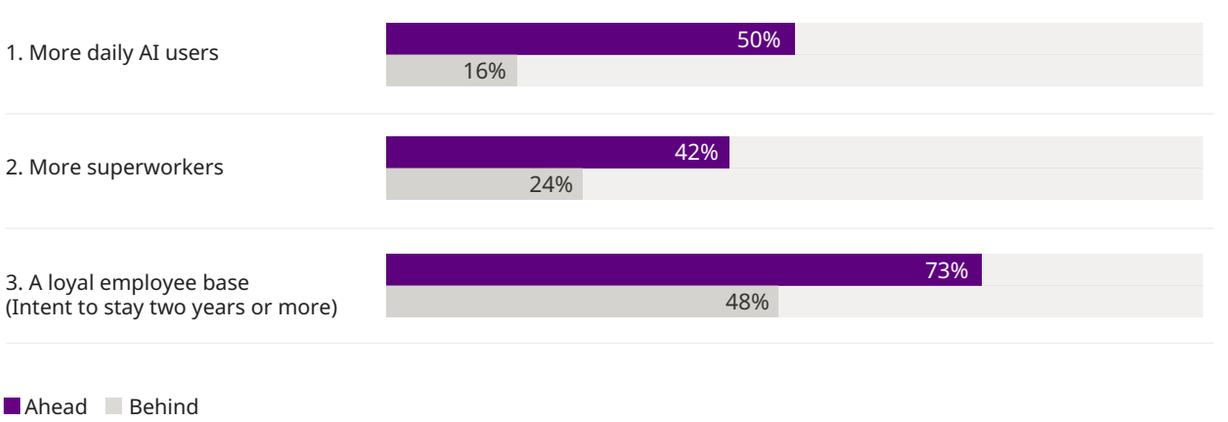
The picture is not entirely comfortable: Employees at leading companies are also more likely to say AI has already reduced headcount and to worry about their own potential redundancy. At the same time, they are more likely to expect to share in AI productivity gains, and more likely to believe in the future of the company. That is what organizational alignment looks like from the employee side: People can see the disruption coming, but they are more willing to lean into it when they understand where the company is going and believe there is a place for them in that future.

Employees who see their company as ahead on AI describe a different culture

Respondent sentiment on employer AI approach, by how employees rate their company's AI progress versus competitors



WHAT THAT CULTURE BUYS



Source: Oliver Wyman Forum Global Voices Survey (April 2026)

Survey questions: "How would you describe your company's progress in adopting AI compared to others?" ("Ahead of the competition"; "Behind the competition"); "To what extent has your company's leadership stated a clear ambition on the AI strategy?" ("Extensive; my company has provided a clear vision and has kept employees informed on progress"); "My employer's training and resources have enhanced my use of AI" ("I agree with this statement"); 7-point agreement, with scores of 6 or 7, with "My company values and seeks my input on important topics," "I believe in the future of my company," and "I feel comfortable sharing my honest opinions at work"

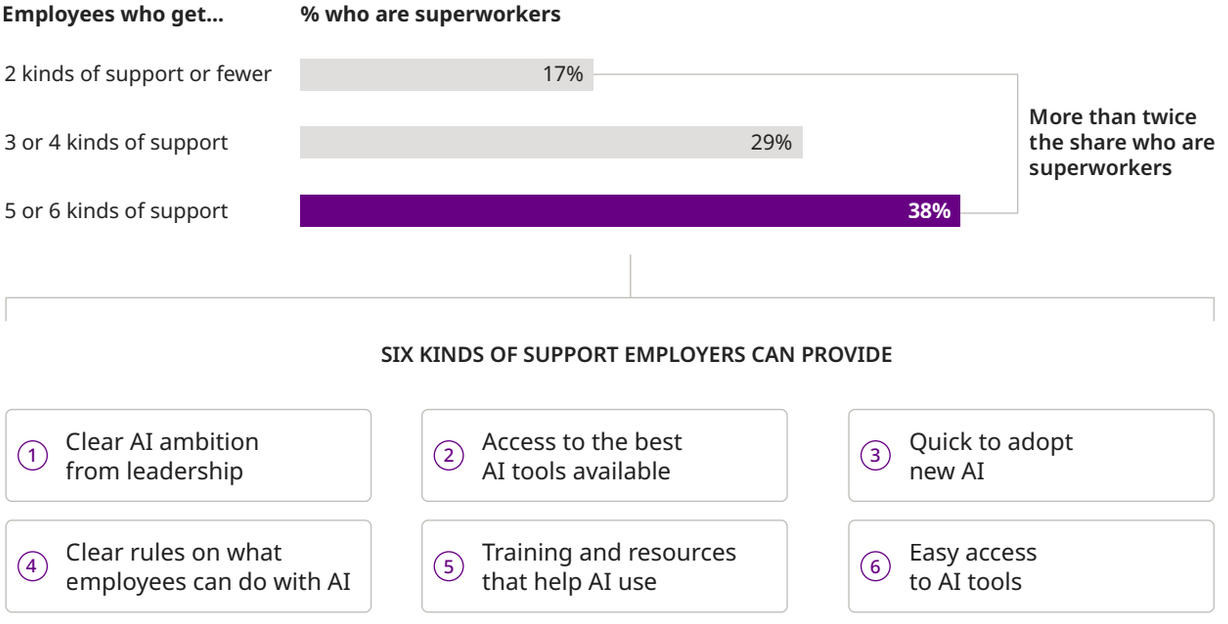
Base: N=1,488 (ahead) and N=1,793 (behind) employed respondents; output comparison for superworkers filtered to workplace AI users, N=1,450 and N=1,511

Six conditions that employers can influence are associated with more superworkers: a clear AI ambition with progress shared, access to frontier models, quick adoption, clear rules, training that improves use, and easy access to specific AI apps/tools. Where five or six are present, 38% of AI users are superworkers, compared with 17% where two or fewer conditions prevail. To stress how consistent this pattern is, the higher superworker rate associated with the above conditions holds across every white-collar job level, broad industry group, and market surveyed.

One large financial-services firm is rolling out Digital Twins across its 5,000-plus workforce. Each is designed to learn how an employee works and help with research, meeting preparation, and email. The model offers a glimpse of the bionic workforce: one employee surrounded by persistent machine capacity that becomes more useful as it learns how they work. It also raises a harder question about trust. A system that learns how someone works can become a record of how they work, so employees need clear consent, access rules, and confidence that it is there to extend their capacity rather than monitor performance or make them easier to replace.

The superworker matters less as a new employee category than as evidence that the performance frontier has moved. The companies that pull ahead will be the ones that make exceptional human-plus-AI performance progressively less exceptional.

Superworkers are more than twice as common where employers provide the right support
Employees who use AI at work, grouped by how many kinds of support they get



Source: Oliver Wyman Forum Global Voices Survey (April 2026)

Survey questions: "With AI, how much more work can you get done compared to without it?" (superworkers = "About double (2x)", "2–5x more", or "More than 5x"); "To what extent has your company's leadership stated a clear ambition on the AI strategy?" ("Extensive"); agreement with "My organization provides access to the best AI tools available", "My organization has clear rules on what employees can and cannot do with AI", "My organization is quick to adopt new AI tools compared to other companies", "My employer's training and resources have enhanced my use of AI" (agree); "At your workplace, how easy is it to get access to AI tools you want to use?" ("Very easy" or "Somewhat easy")

Base: N=9,836 workplace AI users; 2 kinds of support or fewer N=3,778; 3 or 4 N=3,151; 5 or 6 N=2,907



The market for *machine attention*

When the buyer is no longer human, what does winning the customer even mean?

Authors
Adrian Oest, Dan Jones, Emma Stall, Joshua Gilbert, Michael Zeltkevic, Muriel Schtickzelle, Rupal Kantaria, Simon Luong

Never before in the history of commerce has a shopper *not been human*

AI is changing that. People will still have the need, the money, and, when they want, the final say over purchases. But AI will increasingly search, compare, recommend, and eventually buy on their behalf. By 2030, some of the most significant consumer and business buying decisions may no longer be made by humans alone. Already, people are beginning to trust AI with choices they once made themselves or with the help of others.

Companies need to start preparing for this future now. AI is poised to influence every stage of commerce: how companies market, which channels matter, how products get found, how people compare and choose, what loyalty means, and how a transaction gets completed. Strategies honed for decades to optimize around human attention, judgment, and action need to be reformulated to take machines into account.

THE REALITY IS...

1. AI is changing shopping behavior *faster* than many businesses realize
2. Brands can't simply *buy their way* to the top of an AI recommendation
3. In the new era of *shopping agent optimization*, anything an agent can't read, evaluate, and buy will be unsellable
4. The *loyalty equation* is being rewritten as AI weakens two of the oldest moats in retail: familiarity and inertia
5. Sellers face a *new trust question*: Who authorized this agent — and to do what?
6. Everyone will want to be *the customer's agent*, but no one will be the only one

By the numbers

AI is already influencing consumer shopping more than most people realize

71%

of consumers have used AI to find or purchase new products[18]

AI recommendations increasingly shape the final choice

87%

of AI shoppers end up choosing a product AI recommends at least some of the time[19]

By 2030, the vast majority of e-commerce could be AI-influenced

82%

of global e-commerce could involve AI in product discovery, comparison, or selection by 2030, up from 40% in 2026[20]

AI ranks second only to search as a starting point for product discovery

33%

of consumers start their shopping journey with AI, ranked second only to the standard search engine[21]

8×

more customers start their shopping journey with AI than in a physical store (33% versus 4%)[21]

AI is now the No. 1 trusted source for shopping advice, above friends, family, and brands

32%

of consumers would follow AI when recommendations conflict, versus 23% who would follow friends and family[22]

Only 11%

would follow a brand's advice, the lowest of any source[22]

Consumers are ready for AI to manage some of their spending

2 in 3

consumers (65%) would let an AI agent manage some of their household spending within rules they set[23]

Only 20%

wouldn't let AI manage any of their spending[23]

Source: Oliver Wyman Forum Global Voices Survey, August 2026 (footnotes 18, 19, and 23) and April 2026 (footnotes 21 and 22); Oliver Wyman Forum analysis (footnote 20)

18. Survey question: "For which of the following purposes have you already interacted with AI?" Base: N=12,837

19. Survey question: "When using AI to shop for a product, how often do you end up choosing one of the products that AI recommends?" Base: N=9,072 (AI shoppers)

20. Oliver Wyman Forum directional forecast based on OWF consumer adoption data and global e-commerce market growth assumptions. Note: AI-influenced commerce includes journeys in which AI plays a role in discovery, comparison, recommendation, or product selection. It does not necessarily mean that an AI agent completes the transaction.

21. Survey question: "When you need to find a new product or service and don't have a specific brand or provider in mind, where do you go first?" Base: N=12,883. Note: 43.6% reported search engine, 33.3% reported AI, 9.2% reported online marketplace, 5.8% reported social media, 4.4% reported physical stores, and 3.8% reported friends.

22. Survey question: "When shopping for products, you might get recommendations from different sources. If the following sources gave you conflicting recommendations, whose advice would you be most likely to follow?" Base: N=12,883. Note: 31.9% reported AI, 22.6% reported friends and family, 21.8% reported ratings and reviews, 13.2% reported social media, and 10.6% reported a familiar brand.

23. Survey question: "What percentage of your annual household spending would you allow an AI shopping agent to manage within rules you had set?" Base: N=12,837. Note: 65% reported a non-zero amount, 21.3% reported zero, and 14.1% were unsure.

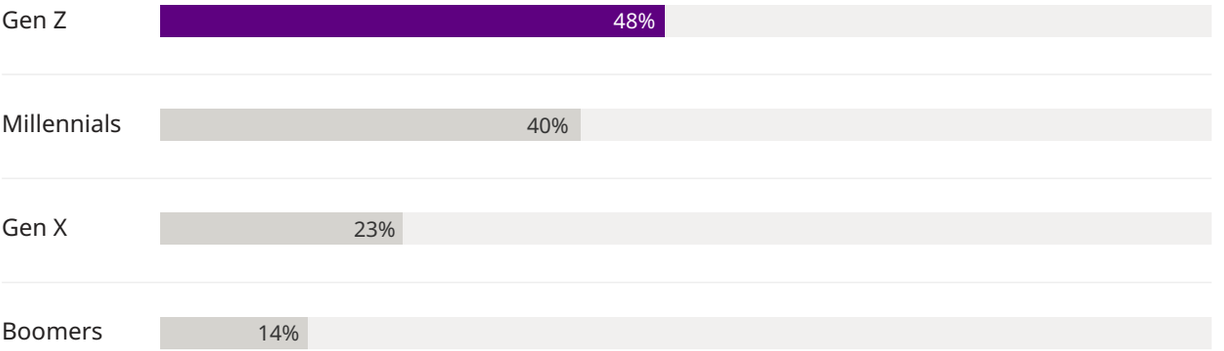
AI is changing shopping behavior *faster* than many businesses realize

AI is already a mainstream starting point for shopping, especially among Generation Z and millennial consumers. It now outranks social media, friends and family, and physical stores as the first stop for product discovery.

Customers increasingly can skip major parts of the journey that businesses have spent decades learning how to influence. A shopper can brief AI on what they're looking for — and let the machine search, compare, and narrow the field before visiting a brand or marketplace.

Nearly half of Gen Z start shopping with AI when looking for a new item

Share of each generation who go to AI first when looking for a new product or service



Source: Oliver Wyman Forum Global Voices Survey (April 2026)

Survey question: "When you need to find a new product or service and don't have a specific brand or provider in mind, where do you go first?"

Base: N=12,883 (total April 2026 survey across 15 markets)

Meta's successful launch of Muse in September showed how eager customers are to try these tools. Brands and retailers are scrambling to catch up.

The extent of AI's influence depends in good part on whether a shopper begins their search with a need or a brand. In our survey, 42% of people who use AI for shopping say they usually start by describing a need and asking for recommendations; 33% begin with a brand already in mind. When the need comes first, the agent can shape which brands enter consideration before the shopper names one.

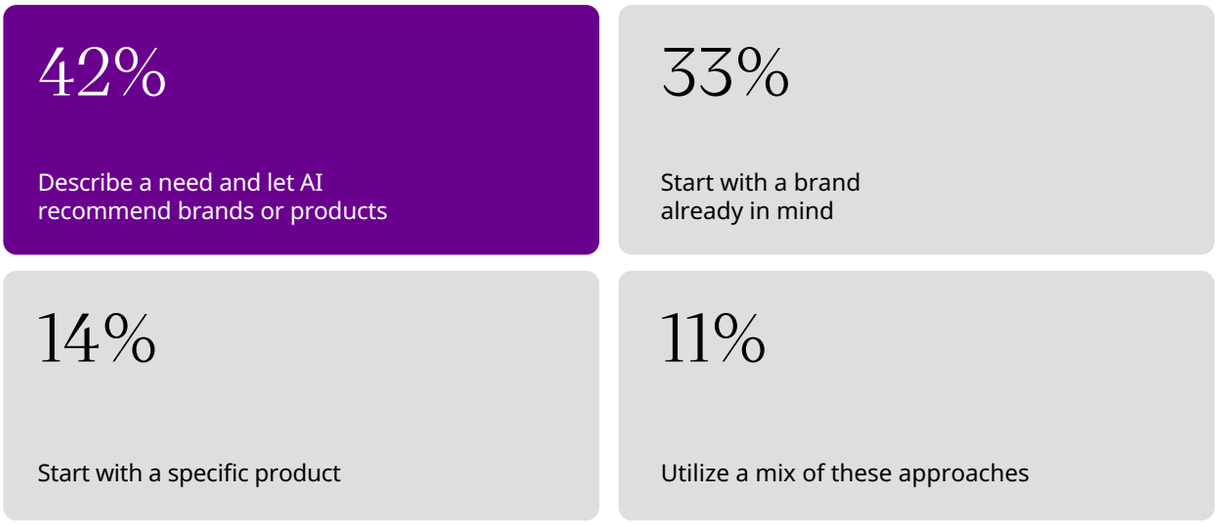
When shopping recommendations conflict, AI tends to be decisive. Respondents were more likely to follow AI than their friends and family or ratings and reviews in such cases. AI is beginning to compete with, and beat, sources shoppers have relied on for years.

Brands have two routes into AI-assisted shopping: be named by the customer before the search begins or be chosen by the agent during it. The first still depends on marketing to make the brand, or the qualities associated with it, memorable enough that customers carry them into the brief they give AI. Once the agent begins comparing options, fit and evidence determine which brands stay in contention.

AI ranks first among sources consumers say they would follow, ahead of friends and family and ratings and reviews

Among shoppers who use AI, more start with a need than with a brand

Approaches AI shoppers use when starting their AI shopping interactions



Source: Oliver Wyman Forum Global Voices Survey (August 2026)
Survey question: "When using AI to shop, which best describes how you usually start?"
Base: N=9,072 (respondents who have previously used AI for finding/purchasing products)

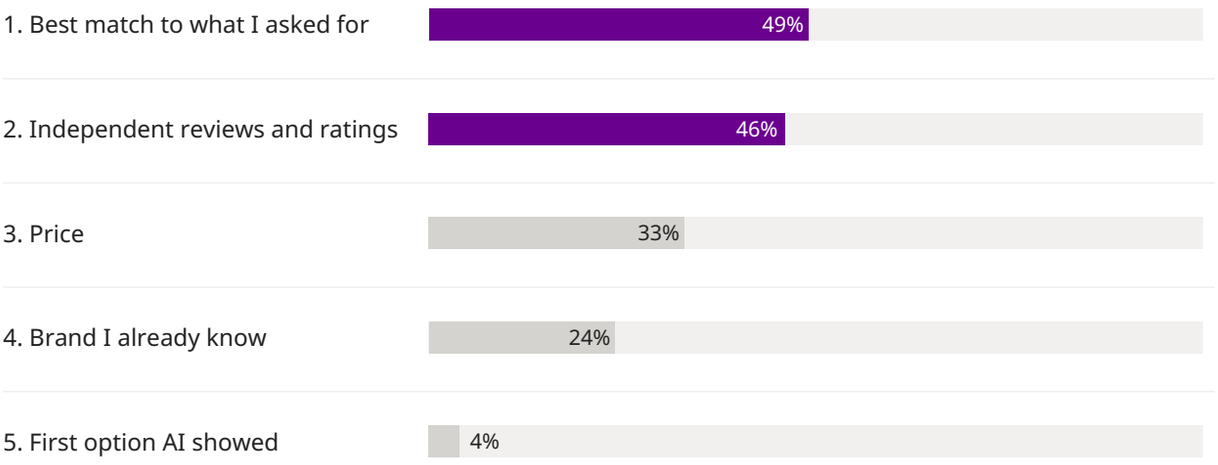
Brands can't simply *buy their way* to the top of an AI recommendation

Brands pay supermarkets to display their goods at eye level. They pay search engines to appear at the top of the page. Now AI is creating a third shelf. An AI answer to a consumer query may surface only a handful of options from a much wider field, selected according to how well they fit what the customer asked for and the evidence available to support them.

That's exactly what many consumers are looking for. When choosing among products suggested by AI, 49% say they care most about what best matches what they asked for, while 46% cite independent reviews and ratings. The AI shelf is harder to buy because there is no single shelf: The shortlist is generated to fit an individual shopper's brief, and that brief can change with the customer and the occasion.

Fit and independent reviews matter more than price or brand familiarity

Factors taken into account when choosing among AI suggestions, ranked



Note: Multi-select question; answers do not add up to 100%.

Source: Oliver Wyman Forum Global Voices Survey (August 2026)

Survey question: "When you are choosing between products an AI has suggested, what matters most to you?"
Base: N=12,837

As AI providers explore advertising, paid placement could weaken the trust that makes AI recommendations valuable. Two-thirds of consumers say they would have less trust in an AI tool if brands paid to appear in its recommendations, including 40% who say they would trust it much less. Paid recommendations may still have a place, but they will need clear labels and a visible distinction from non-sponsored results.

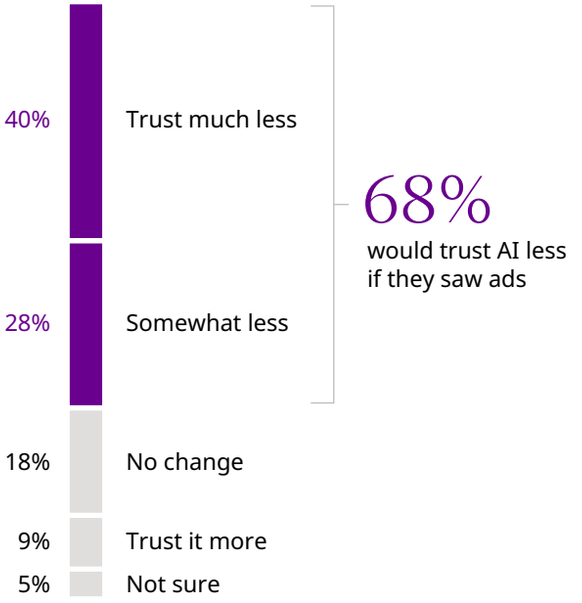
A brand controls only part of the evidence that AI models use. In our AI-shopping simulation, 88% of cited sources were independent of the brands being recommended. Reviews, marketplaces, publishers, and other third parties can shape the answer alongside a company's own product information.

That shifts the basis of influence from persuasion to easily identifiable proof that both agents (and humans) can find. While brand building still matters for shaping what consumers know, value, and ask for, once AI begins comparing options, claims need hard facts behind them. The qualities a company wants to be known for — be it price, performance, reliability, sustainability, service, or something else — increasingly need to be supported by credible evidence.

The same shift is emerging in business buying, where an agent can narrow dozens of suppliers to a shortlist before the buyer starts comparing them. That provides challengers with a new route to being considered and gives incumbents one less reason to assume familiarity will earn the business. Winning the AI shelf requires a close fit to the brief, product facts an agent can read, and proof that extends beyond the brand's own claims.

Most consumers say ads in AI recommendations would reduce their trust in the tool

Extent to which paid recommendations would affect trust in an AI tool



Source: Oliver Wyman Forum Global Voices Survey (August 2026)
Survey question: "If you found out an AI tool was recommending products because brands paid for that placement, how would it affect your trust in its recommendations?"
Base: N=12,837

88%
of the sources that AI cited in our simulation were independent of the recommended brands

In the new era of *shopping agent optimization*, anything an agent can't read, evaluate, and buy will be unsellable

Companies have spent years mapping every step of the customer journey, from initial research through product ownership. Now they need to do the same for AI agents. The new practice that we call "shopping agent optimization," or SAO, requires brands to make product information, supporting evidence, and transaction flows easy for authorized AI agents to understand and act on.[24]

Existing e-commerce websites can fail this test well before an agent makes a recommendation. Product information written for people may be scattered across pages, embedded in images, inconsistent across channels, or missing facts the agent needs. In our audit, 70% of the product pages tested did not provide a complete, purchase-ready set of facts.

Making the information readable is only the first step. Existing checkout processes, including checkout forms, carts, options, identity checks, payment controls, and confirmation flows, were built around a person clicking through screens. Many merchants also deliberately use bot-detection and anti-fraud tools to identify nonhuman interaction patterns and block them. Agentic commerce requires a more precise distinction, one that enables authorized agents to get to and through the checkout while still blocking malicious automation. Without this, an agent may know exactly what to buy and be unable to complete the transaction.

70%
of product pages tested did not provide a complete, purchase-ready set of facts

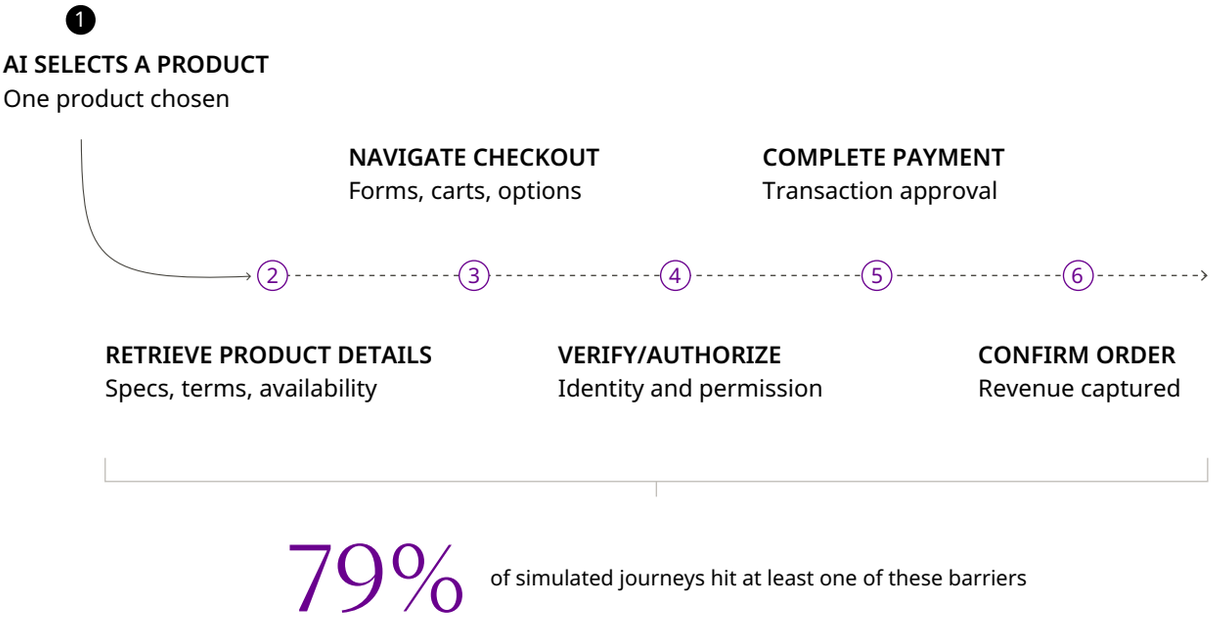
24. A shopping agent is an AI system that can act on a shopper's behalf, for example by searching, comparing, and buying, as well as giving advice. Shopping agent optimization (SAO) means making product information, evidence, and transaction flows easy for authorized AI agents to understand and act on, from discovery through checkout and service.

In our simulated journeys, 79% encountered at least one barrier after the recommendation. Some barriers are essential, especially verification and authorization. But if a legitimate agent cannot prove its authority, understand the required step, or proceed, that's a failure.

Companies should redesign the agent journey to remove unnecessary friction while preserving the controls needed to establish identity, authority, and acceptable risk. Product, commerce, identity, payments, operations, and risk teams should map where legitimate agents are blocked, determine what evidence an agent must present at each stage, and agree on how that evidence should be evaluated. The last thing a brand wants is to win the recommendation only to lose the sale several steps later.

Most agent-led shopping journeys hit a barrier after a product is recommended

Where an AI shopping agent must keep moving after a product is recommended



Source: Oliver Wyman Forum analysis of 62 retailer product journeys across five product categories using simulated public-web AI shopping journeys

The *loyalty equation* is being rewritten as AI weakens two of the oldest moats in retail: familiarity and inertia

Shoppers habitually default to names they know and products they've bought before because comparing every option for every purchase takes more time than most people are willing to spend. Familiarity narrows the field; inertia rules out fresh exploration. Agents can be a game-changer, though. They can compare price, fit, reviews, availability, and alternatives every time the desire for a purchase arises.

Fifty-four percent of respondents say they would consider a brand they had never heard of if AI said it best matched their needs. A place on the shortlist is not the same as a sale, but a strong fit can earn consideration before familiarity does.

54%
of consumers would consider a brand they had never heard of if AI said it best matched their needs

Agents also can disrupt established buying habits. Six in 10 respondents say they would consider switching from their usual brand if an AI agent found an alternative better suited to their needs. The ease of comparison makes it easier for consumers to try something new.

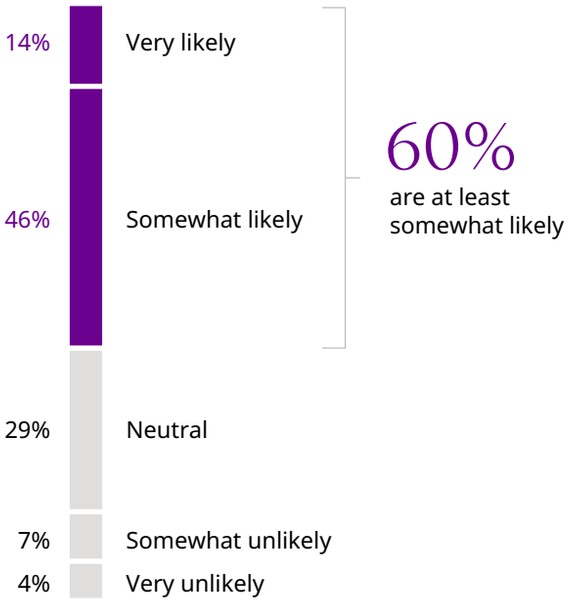
AI's impact on brand value may vary sharply by category. Where products are easy to compare and switching costs are low, agents can make incumbency less valuable by reopening the choice every time. Insurance renewals are an obvious example: An agent could continuously compare price, coverage, and service rather than defaulting to the existing provider. In other categories, AI may reinforce the value of distinctive brands. In luxury, for example, iconic brands may benefit when heritage, authenticity, scarcity, and resale value are part of the brief, while less differentiated aspirational brands may face more pressure. In travel and service-oriented industries, some elements may become easier to compare and commoditize, while reputation, trust, experience, and service quality become more important precisely because they are harder to reduce to price alone.

Companies should ask which parts of their advantage an agent can compare directly, which it can verify through evidence, and which still depend on trust, experience, or identity.

Loyalty will increasingly have to be earned anew at every repeat purchase — but what wins that competition will differ by category.

Six in 10 would consider switching from their usual brand if AI found a better fit

Likelihood of considering a switch if AI identified a better alternative



Source: Oliver Wyman Forum Global Voices Survey (August 2026)

Survey question: "If an AI assistant identified an alternative product or service that it determined was better suited to your needs than the one you currently buy, how likely would you be to consider switching?"

Base: N=12,837 (total August 2026 survey across 15 markets)

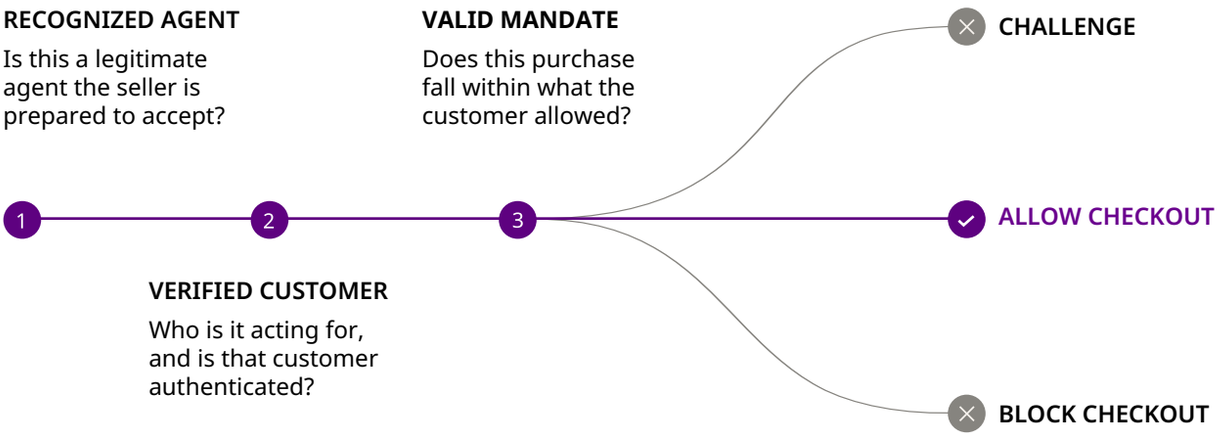
Sellers face a *new trust question:* Who authorized this agent — and to do what?

For 30 years, e-commerce websites have checked whether visitors behaved like humans to defend against fraud. Agents may fail that test by design. A machine-like visitor may now be a legitimate representative of a real customer. Detection tools can help merchants distinguish recognized agents from malicious bots, but recognizing the agent does not prove who is behind it or what that person authorized.

Consider the possibilities facing a seller today. A recognized agent may be acting for a verified customer within a valid mandate. The same recognized agent could also be compromised, used by a fraudster, or instructed to act beyond what the customer permitted. And an unrecognized or malicious bot may attempt to imitate legitimate agent activity. "Human or bot?" is too crude a question to pose. The seller needs to establish three things: which agent is acting, which customer stands behind it, and whether the purchase falls within that customer's mandate.

Sellers need to verify the agent, the customer, and the mandate

Merchant-side questions that determine whether an agent-led checkout may proceed



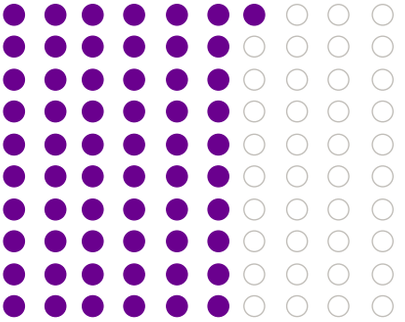
Source: Oliver Wyman Forum analysis

In the agent-led checkout models now taking shape, these three trust questions converge at the merchant's front door. Emerging payment protocols are beginning to turn that mandate into portable, cryptographically signed evidence of the customer's intent, including the purchase or spending boundaries the customer authorized. If unauthorized activity gets through, the seller may still bear operational consequences such as chargebacks, disputes, and fraud losses. That evidence therefore needs to travel with the transaction and remain available if someone later disputes what the agent was allowed to do.

The exposure runs well past payment authorization into account opening, personal data, returns, subscriptions, and obligations written on the assumption that a person was present. Consumers are giving companies deploying AI agents little room for error: 61% say they would either leave the company or stop using its AI features after one serious AI mistake involving their money or personal data.

Proof of authority will become an essential part of the customer experience. Sellers need ways to accept agent-led business when the agent is recognized, the customer is verified, and the mandate is valid, without lowering the bar against fraud, misuse, or unauthorized action. Payment providers also need enough of that evidence to evaluate the transaction, and the record needs to remain available if a later dispute tests what the customer permitted. Sellers that can establish, transmit, and preserve that proof may accept business their competitors have to refuse.

After one serious AI mistake with money or personal data, 61% would leave or stop using the company's AI

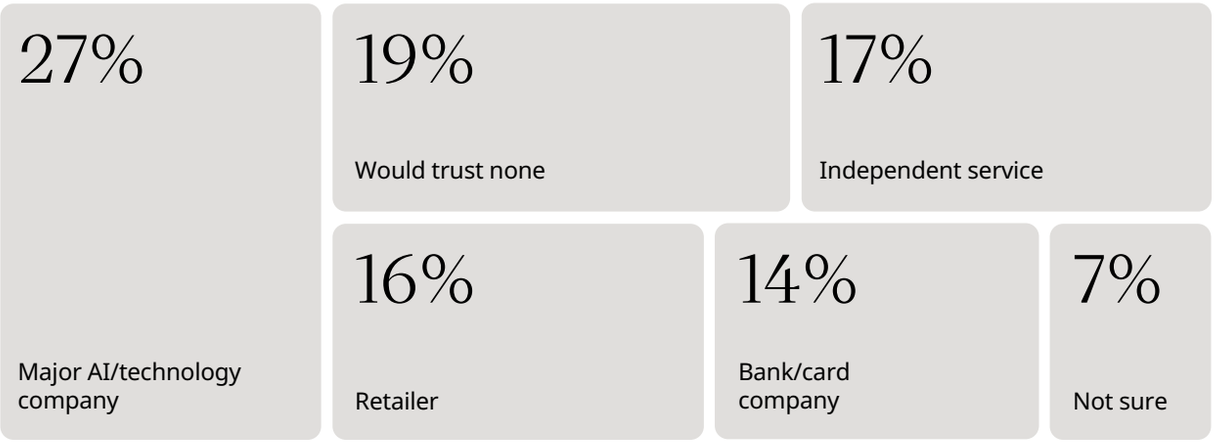


Everyone will want to be *the customer's agent,* but no one will be the only one

Retailers, banks, technology firms, and independent services are all building shopping agents, but it's an open question which will win consumers' trust. A major AI or technology company currently leads, with 27% of respondents saying they would trust it to provide a shopping agent, while others are cited by fewer than 20% of respondents. Nearly one in five say they wouldn't trust any of the providers listed.

Several providers may coexist because consumers want different degrees of delegation for different purchases. People express the greatest willingness to hand authority over to AI agents for routine purchases such as groceries and food delivery. They want to retain more control over discretionary items such as air travel and consumer electronics, and are least willing to delegate major decisions such as buying a car or taking out a loan or mortgage. Nearly half are already willing to delegate some spending authority: 49% would let an AI agent spend at least some amount on their behalf without checking before each purchase. That authority can increasingly be encoded into the payment itself. Emerging agentic payment credentials can restrict an agent by spending limit, merchant, purpose, or time window, allowing consumers to give an agent a defined "license to spend" rather than unrestricted access to their money.

No type of shopping agent provider wins the trust of even 30% of consumers
Sources consumers would most trust to provide an AI shopping agent

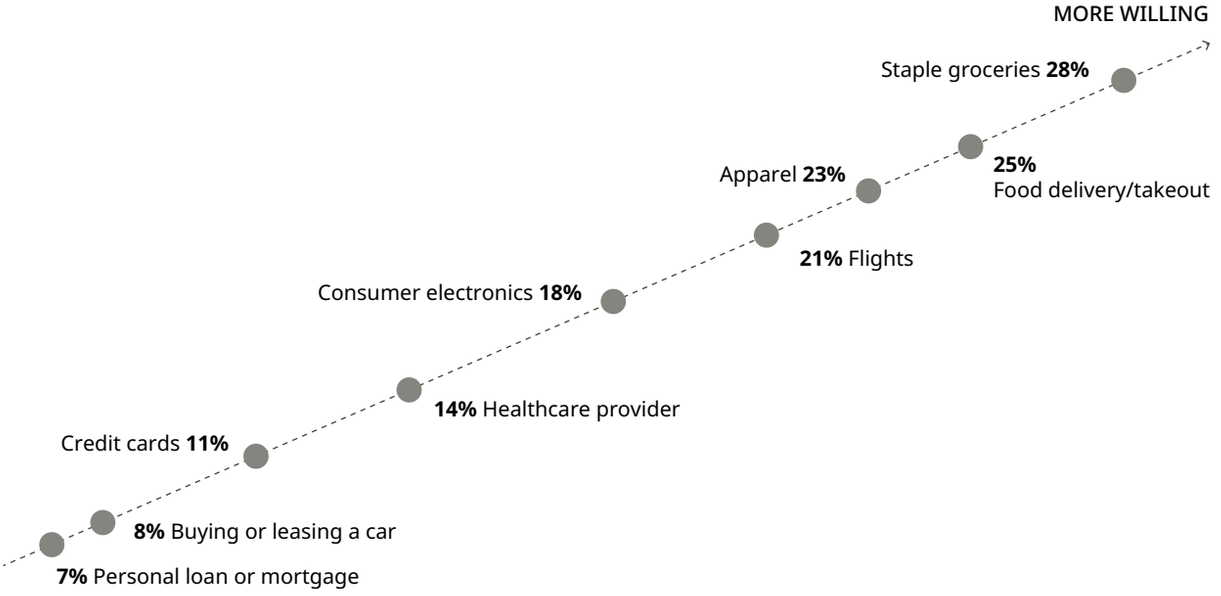


Source: Oliver Wyman Forum Global Voices Survey (August 2026)
Survey question: "If you used an AI shopping agent, who would you most trust to provide it?"
Base: N=12,837

Customers may use different agents depending on the type and cost of purchase, which provider they already trust, or how much control they want to retain.

Consumers also care about whose interests the agent serves. Only 24% believe an AI shopping assistant is really working for them when it gives a recommendation. Separately, one in four say they would hold back from delegating some purchases because they worry AI may put brands' interests ahead of their own. A branded agent can still earn trust; its incentives and mandate will shape how much authority customers are willing to give it.

Consumers are more willing to delegate routine purchases to AI than consequential ones
Consumer willingness to let an AI agent make a purchase on their behalf, by item category



Source: Oliver Wyman Forum Global Voices Survey (August 2026)
Survey question: "For which of the following would you be happy for an AI agent to find and buy on your behalf, within rules you had set in advance?"
Base: N=12,837 (total August 2026 survey across 15 markets)

Sellers need to prepare for both sides of this market: serve customers through their own AI experiences and make products accessible to the third-party agents customers choose. That access cannot be assumed. In September 2026, Amazon blocked Muse from shopping on Amazon.com, telling users that access by an "unauthorized AI agent" violated its terms. The episode points to another layer of competition in agentic commerce: Merchants will decide which agents they recognize, what credentials they require, and what recognized agents are allowed to do. Emerging protocols already contemplate different permissions for activities such as browsing and payment.

The market may therefore develop something resembling a system of digital IDs for agents, mandates, and access rules. An agent can prove who it is; the customer can prove what it is allowed to do; and the seller can decide whether those credentials are sufficient for entry. More commerce may eventually happen through interactions between a customer's agent and a seller's agent, including negotiation over service, delivery, loyalty benefits, contract terms, or price within preset guardrails.

A critical strategic question is how to participate in a multi-agent market. Companies need to decide where they want to own the customer interface, which outside agents they will permit or support, what those agents can access, and which commercial and risk rules govern the transaction. Those choices will determine how much of the customer relationship they control as machine buyers become a meaningful route to market.

1 in 4 are held back from delegating purchases by concern that AI may serve brands' interests over their own

4

How *guardrails* can enable growth

Enterprises that lean into AI risk management create conditions for lasting adoption and innovation



Authors
Daniel Cope, Elena Belov, Frédéric Thomas-Dupuis, Nick Dykstra, Ted Moynihan, Til Schuermann

Good risk management and governance enable the enterprise to go faster with AI

AI is expanding the potential of what companies can do. How well they manage the risk determines how much of that potential they can capture.

Every new model pushes the boundary of what's possible, often in ways that still manage to astonish. But possibility and enterprise value are not the same thing. The useful frontier is the part a company can reach repeatedly, safely, economically, and with enough confidence to put real decisions, data, and customers behind it.

Formula 1 teams understand that distinction instinctively. Their cars run near the physical limit because the controls are designed for high performance. The track sets boundaries, telemetry reveals what the driver cannot see, and the team rehearses relentlessly to respond to failure. The car can go fast precisely because the organization knows how to manage the consequences.

Managing AI risk needs the same dynamic approach. The governance provides the pathways, controls, accountabilities, and feedback loops that let an enterprise take risk deliberately rather than avoid it by default. AI's behavior is probabilistic, responding to more context. A helpful assistant can quickly become part of a workflow, eventually touching sensitive data or taking actions on the company's behalf. Governance needs to be equally agile.

Companies that can manage AI's uncertainty, speed, reach, and scale can pursue more consequential uses, and recover faster when failures occur. When competitors can buy the same models, that operating discipline can become the source of an incredibly powerful advantage.

THE REALITY IS...

1. Good governance has to feel both *safe and easy* to use
2. Some AI risks become visible only in experimentation — that is what a *risk laboratory* is for
3. To scale AI, risk management has to *move upstream* so business, tech, and risk start together
4. A *human in the loop* cannot be the operating model for AI oversight
5. The age of AI requires *living governance* oversight that moves at the speed of the technology

By the numbers

One serious AI mistake can shatter customer trust

39%

of people would stop using a company entirely if its AI made one serious mistake with their money or personal data; another 23% would avoid its AI[25]

Clear, embedded AI is associated with more reported value

88%

of workplace AI users with embedded AI and clear guidance report a productivity gain, versus 52% with neither[26]

The public holds companies that use AI accountable

52%

would hold an organization that used AI, or the person who approved it, responsible for a proven loss; 24% would hold the developer responsible[27]

Overly strict AI governance spawns shadow AI

35%

of employees using nonapproved AI say the approved tools were worse; 27% say approval takes too long[28]

63%

use nonapproved tools where the approach feels too strict, versus 38% where it feels about right[28]

Greater AI productivity comes with greater data stakes

90%

of users applying AI across four or more distinct work activities report a productivity gain, versus 45% using it for one activity[29]

63%

of those same users have entered sensitive work information into an AI tool, versus 33% using AI for one activity[29]

Human review is expensive and often imperfect

22 min.

of checking, correcting, or redoing the output for every hour of AI-assisted work[30]

Only 17%

of users always verify important AI-generated information elsewhere[30]

Source: Oliver Wyman Forum Global Voices Survey, August 2026 (footnotes 25, 27, 28, and 30) and April 2026 (footnotes 26, 27, and 30); Oliver Wyman Forum analysis

25. Survey question: "Imagine a company's AI made one serious mistake with your money or personal data. What would you most likely do?" Base: N=12,837 (all respondents). Note: 39% selected "Stop using that company entirely"; 23% selected "Keep using the company but avoid its AI features."; 29% selected "Keep using the company as before, if they fixed the mistake"; 10% selected "Not sure."

26. Survey questions: "Which best describes your workplace guidance on AI use?" ("Clear policy and I understand it"); "Which AI tools or features are available to you at work?" ("AI features embedded into workplace software I already use"); "How has AI impacted your productivity?" ("AI has improved it"). Base: N=2,298 (both) and N=3,233 (neither), workplace AI users

27. Survey question: "An AI system used by a company causes a proven financial or personal loss. Who should carry primary responsibility for putting it right?" Base: N=12,336 (all respondents, excluding "not sure"). Note: 52% combines "The organization that used the AI" (37%) and "The person who approved or oversaw the decision" (15%).

28. Survey questions: "If you have used nonapproved AI tools at work, what was the main reason?" Base: N=3,567 (employed respondents reporting nonapproved use who gave an applicable reason). "Which better describes your company's approach to AI?"; "Do you use AI tools at work that your employer has not approved or does not know about?" Base: N=815 (too strict), N=6,160 (about right)

29. Survey questions: "How have you used AI at your job in the last 3 months?"; "How has AI impacted your productivity?"; "When using AI for work, which of the following have you entered or uploaded into an AI tool?" Base: N=3,024 (four or more work activities) and N=2,003 (one work activity), workplace AI users. Note: The eight work activities were generating material, customer interaction, data analysis, writing assistance, screening or summarizing, search, work planning, and training. Sensitive information includes customer/client, financial/legal, source code or proprietary technical, or employee/HR information.

30. Survey questions: "For every 60 minutes you spend completing work with AI's help, about how many minutes go to checking, correcting, or redoing AI output?" Base: N=7,845 workplace AI users ("I cannot estimate this" excluded); mean. "When you use AI for important information, how often do you verify its output elsewhere?" Base: N=11,731 respondents

Good governance has to feel both *safe and easy* to use

The test of AI governance lies in the route employees are expected to use. A policy can be sound on paper and still fail in the trenches. An approved AI pathway has to be effective for both the organization and the employee.

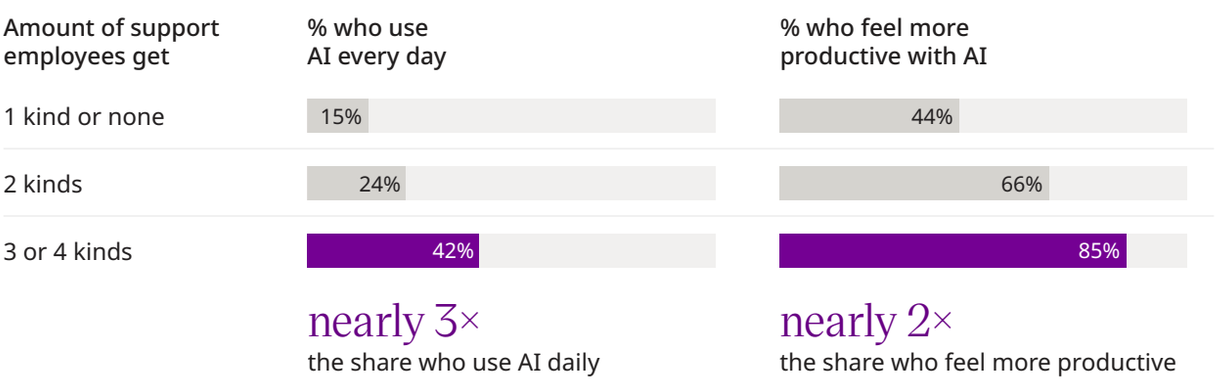
It must control risk without making tools hard to access, difficult to understand, disconnected from daily work, or materially less capable than the alternatives. Otherwise, adoption slows or activity moves outside the governed environment.

The more support employers provide, the more employees use AI and the more productive they feel

Four kinds of support employers can provide for everyday AI use

Easy access to AI tools	Access to the best AI tools available	A clear AI policy that employees understand	AI built into the software employees already use

Employees who use AI at work, grouped by how many kinds of support they get



Note: Comparisons are descriptive and do not establish causation; the four kinds of support also tend to come together.

Source: Oliver Wyman Forum Global Voices Survey (April 2026); Oliver Wyman Forum analysis

Survey questions: "At your workplace, how easy is it to get access to AI tools you want to use?" (easy access = "Very easy" or "Somewhat easy"); "My organization provides access to the best AI tools available" (best tools = agree or strongly agree); "Which best describes your workplace guidance on AI use?" (clear policy = "Clear policy and I understand it"); "Which AI tools or features are available to you at work?" (built in = "AI features embedded into workplace software I already use"); "How often are you using AI in your current job?" ("Daily"); "How has AI impacted your productivity?"(feel more productive = "AI has improved it")

Base: N=9,836 employees who use AI at work: 1 kind of support or none, N=3,036; 2 kinds, N=2,637; 3 or 4 kinds, N=4,163

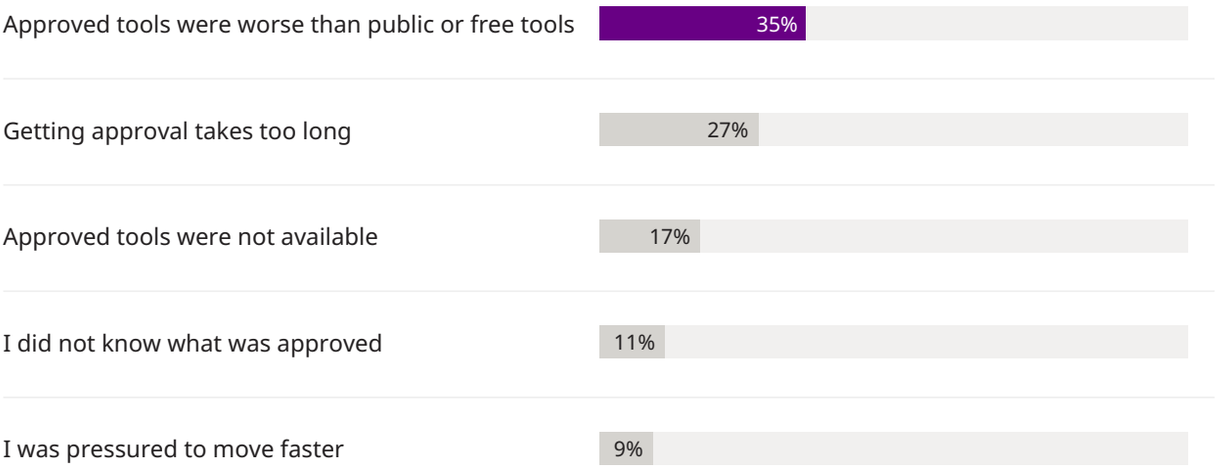
Four conditions are essential for combining safety with greater adoption and productivity: easy access to the AI people want, clear guidance, AI embedded in everyday software, and access to the best available tools. Where three or four are present, 85% of workplace AI users report a productivity improvement and 42% use AI daily. Where zero or one is present, those figures fall to 44% and 15%, respectively.

Where the approved route fails, employees improvise. Among users of nonapproved AI, 35% say the approved tools were worse, 27% say approval took too long, 17% say the tools were unavailable, and 11% did not know what was approved. Where governance feels too strict, 63% use nonapproved tools and 53% have used personal devices or accounts to bypass restrictions; where it feels about right, the figures fall to 38% and 28%.

Good governance makes the approved route the obvious choice. It puts capable tools in the systems employees already use, with rules they can understand quickly. If that route cannot do the job, employees will use AI less or move the work out of sight.

When approved AI falls short, employees go around it

Reasons cited by employees for using nonapproved AI tools at work



63% versus 38%

of employees use nonapproved AI tools where their company's approach to AI feels too strict versus about right

Source: Oliver Wyman Forum Global Voices Survey (August 2026); Oliver Wyman Forum analysis

Survey questions: "Do you use AI tools at work that your employer has not approved or does not know about?"; "If you have used nonapproved AI tools at work, what was the main reason?"; "Which better describes your company's approach to AI?"; "Have AI restrictions at work ever led you or your colleagues to use personal devices or personal AI accounts to get around them?"

Base: Reasons: N=3,567 (employed respondents reporting nonapproved use who gave an applicable reason). Approach: N=815 (too strict), N=6,160 (about right), N=1,086 (too loose); "no real approach" and "not sure" excluded

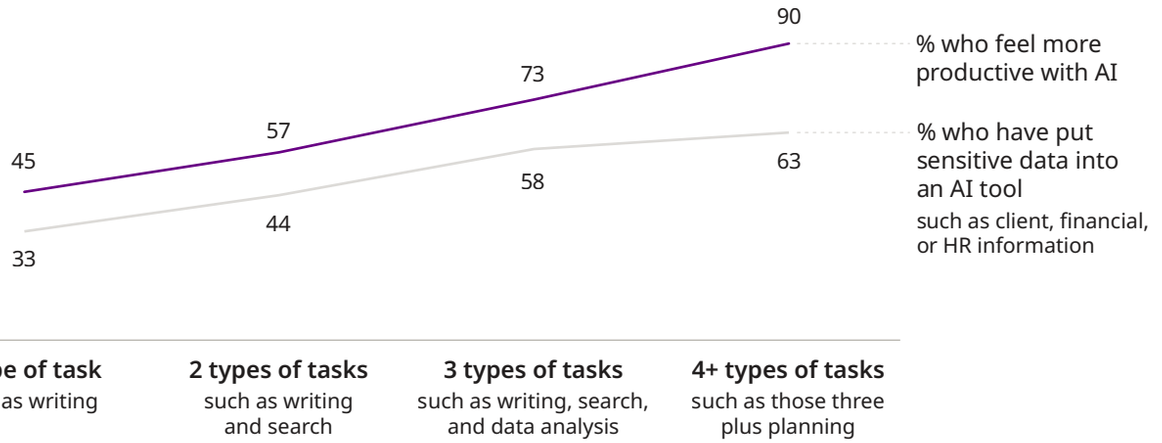
Some AI risks become visible only in experimentation — that is what a *risk laboratory* is for

Standard governance pathways should cover routine, well-understood uses of AI. They cannot cover every high-value application. The uses with the greatest potential often require proprietary data, deeper workflow integration, more autonomy, or a task the organization has never governed before. Those same features make the risk harder to judge in advance.

As employees apply AI across more parts of their job, reported productivity rises sharply — and so does the use of sensitive work information. Nearly half of employees (45%) using AI for one work activity report improved productivity, compared with 90% among those using it across four or more. Across the same groups, the share entering customer or other proprietary information rises from 33% to 63%. These are descriptive relationships, but they expose the governance challenge: AI becomes more useful as it moves deeper into real work, and that real work brings more consequential data and risk.

As employees use AI for more types of work, including work involving sensitive data, they report greater productivity

Employees who use AI at work grouped by how many types of tasks they use it for



Note: Descriptive; breadth of use is also associated with frequency of use.

Source: Oliver Wyman Forum Global Voices Survey (April 2026); Oliver Wyman Forum analysis

Survey questions: "How have you used AI at your job in the last 3 months?" (eight types of task: generating material, customer interaction, data analysis, writing assistance, screening or summarizing, search, work planning, and training; "Other" excluded); "How has AI impacted your productivity?" (feel more productive = "AI has improved it"); "When using AI for work, which of the following have you entered or uploaded into an AI tool?" (sensitive = customer/client, financial/legal, source code or proprietary technical, or employee/HR information)

Base: N=9,582 employees who use AI at work and selected at least one type of task: 1 N=2,003; 2 N=2,239; 3 N=2,316; 4 or more N=3,024

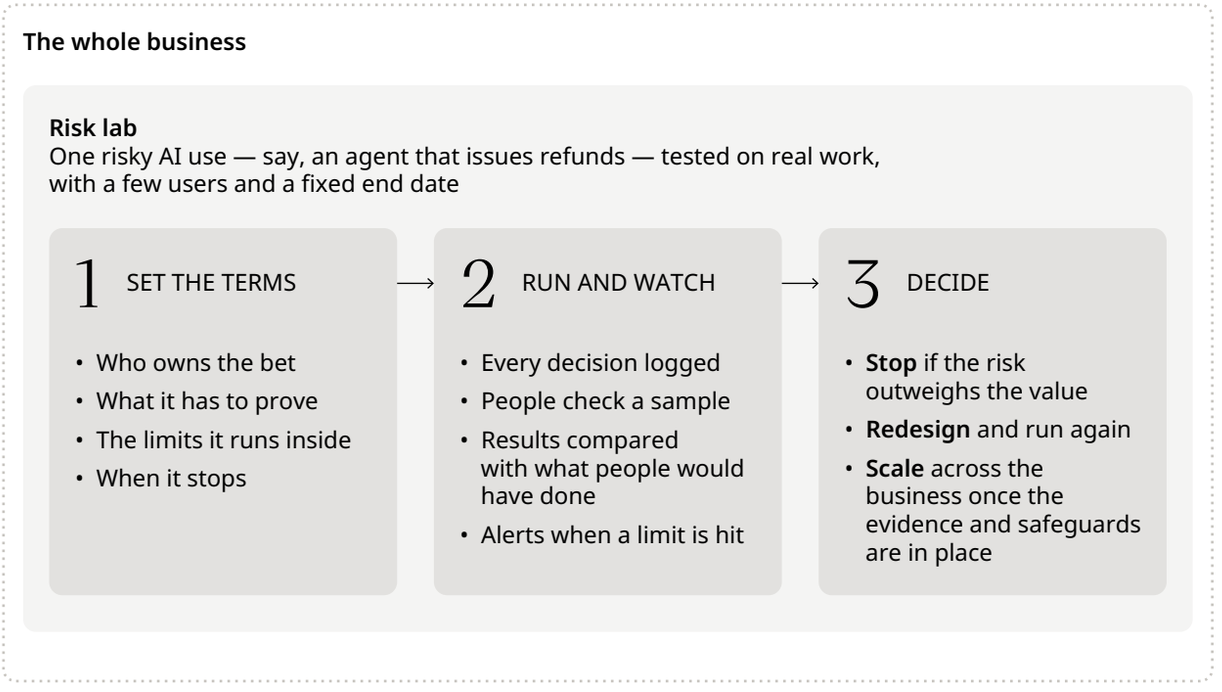
Half of workplace AI users have already entered at least one category of sensitive work information into an AI tool. Employees with clear guidance are more than twice as likely to have done so as those with no policy (59% versus 26%). That does not make every upload appropriate, but it does show that useful AI often needs real work context. Governance has to define how that context can be used safely rather than assume it can simply be kept out of AI.

Standing policies don't cover everything; some risks become visible only when the application is in use. A company still needs a way to test them without granting full production access. A risk laboratory provides a middle ground: a bounded, reversible environment with explicit limits, active monitoring, an accountable executive owner, and predefined conditions for stopping, changing, or scaling the use.

Consider a customer-service agent with authority to resolve complaints, including refunds and goodwill credits, without a person approving each decision. The value is obvious: faster resolution, lower cost to serve, and fewer customers lost while they wait. So is the exposure. The agent spends money, reads account histories, and makes promises on the company's behalf, all at a volume no team could check case by case. No preapproved pattern covers the potential harm.

A risk laboratory can test it, though. It can monitor one product line and one market, with credit limits per case and per day, over a four-week run. Every decision can be logged and a daily sample reviewed, with a reconciliation of what it paid out against what a human team would have paid, and exit conditions set in advance, such as resolution times, leakage, and any policy breach.

Fencing off one risky AI use lets a company accept risks it could never accept across the whole business



Source: Oliver Wyman Forum analysis

A senior executive owns the bet and the reasons the potential value justifies the risk. Because the test is small and reversible, the company can tolerate more uncertainty in the lab than it ever could in production. It also leaves the lab knowing what goes wrong, which safeguards matter, and what has to change before the agent's authority expands.

Moving from a laboratory into production requires a second judgment: Can the company control the risk without stripping away the value? Limits, monitoring, fallbacks, and recovery processes make a system safer, but they also reduce capability and raise cost.

Risk management therefore has to track both sides of the cost-benefit equation. The goal is to achieve enough control to keep risk acceptable while permitting enough AI capability to make the use worthwhile.

AI needs multiple governance routes: Routine uses might follow preapproved patterns, while new uses enter a bounded laboratory and higher-impact uses scale in stages, with explicit thresholds. Each route should make clear what is allowed, what evidence is required, and what triggers deeper review. The purpose is to turn uncertainty into something governable, so the enterprise can take bolder risks deliberately.

2×

Employees with clear guidance on entering sensitive work information into AI tools are more than twice as likely to have done so than those with no policy (59% versus 26%)

To scale AI, risk management has to *move upstream* so business, tech, and risk start together

In many companies, risk management enters the equation only after the business has defined the use and technology has built it. That sequence breaks down with AI, where the model, data, autonomy, and workflow all shape value and risk at the same time. A late control requirement can do more than delay launch; it can change what the system is able to do and whether the use is still worth pursuing.

As a result, high-impact uses need business, technology, operations, legal, and risk involved at the design stage. The pathway should match the consequence. An internal drafting tool does not need the same evidence burden as a system that influences credit, health, employment, or safety. In regulated industries, AI governance should connect to existing disciplines such as model validation, privacy, third-party risk, record keeping, and operational resilience rather than create a parallel AI-only process.

Starting together should not mean a larger committee for every use case. The point is to turn good decisions into reusable infrastructure. Once the company has established a sound framework, permitted data, standard controls, and evidence requirements for a use, later teams should inherit them. Governance then focuses on what is genuinely new: new data, new tools, more autonomy, or higher consequences.

Companies should be able to safely implement subsequent use cases more quickly because previous ones left reusable pathways behind.

Cloud and model platforms can supply much of the technical scaffolding, such as identity controls, logging, tracing, policy configuration, guardrails, evaluations, and monitoring. They cannot decide which errors matter in a claims process, what data an agent may use, how much autonomy is acceptable, or when a result must be escalated. Those choices remain specific to the enterprise and the use case.

The public expects accountability to stay with the organization using the AI. A majority (52%) say the organization using the AI is responsible for any loss it creates, compared with 24% who assign primary responsibility to the developer.

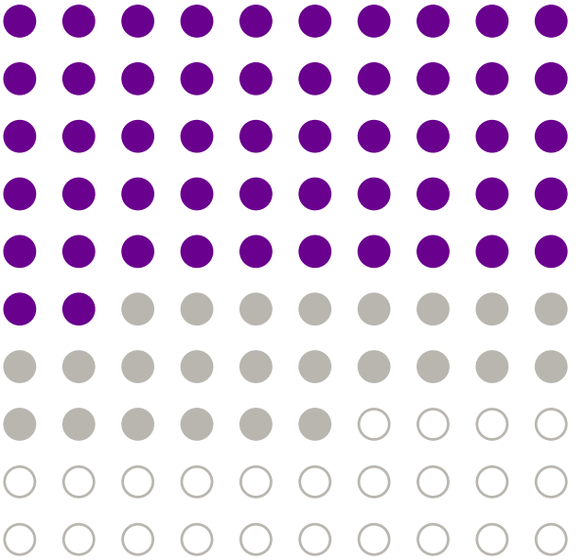
64%

of white-collar executives say their company has a clear AI policy, compared with 37% of non-managerial employees

Reusable pathways also have to reach the people doing the work. Among white-collar executives, 64% say their company has a clear policy they understand, compared with 37% of junior white-collar employees in non-managerial roles. A governance system is scalable only if employees can see which route applies and use it without reopening settled decisions.

The governance challenge becomes more complex when AI begins executing high-volume workflows, operating continuously, or taking actions on the company's behalf. Policies and approval routes are no longer enough; oversight has to be built into the system itself. We explore that next.

The public holds companies that use AI responsible for their mistakes



52%

would hold an organization that used AI, or the person who approved or oversaw the decision, primarily responsible for putting right a proven loss

24%

would hold the company that developed the AI primarily responsible. 19% say responsibility should be shared

Source: Oliver Wyman Forum Global Voices Survey (August 2026); Oliver Wyman Forum analysis

Survey question: "An AI system used by a company causes a proven financial or personal loss. Who should carry primary responsibility for putting it right?"

Base: N=12,336 (all respondents, excluding "not sure")

A *human in the loop* cannot be the operating model for AI oversight

As AI moves from drafting individual outputs to executing high-volume workflows, no organization can put a person in front of every action. Review becomes a bottleneck, plausible errors slip through, and one mistake can repeat across thousands of decisions before anyone sees it. Humans do not operate at machine speed.

These limitations are already becoming visible. Employees spend an average of 22 minutes of checking, correcting, or redoing the output for every hour of AI-assisted work. Yet only 17% of employees always verify important AI-generated information elsewhere, and 42% say wrong AI information caused a real problem in the past year. Review extends the length of AI-assisted work processes by more than a third, consuming a chunk of the supposed productivity gains, but it is still not producing reliable assurance.

A human in the loop is already costly, and it still lets errors through

Employees who use AI at work

22 min.	17%	42%
of checking, correcting, or redoing the output for every hour of AI-assisted work	of employees always check important AI answers against another source	of employees have had wrong AI information cause a real problem in the past 12 months

Source: Oliver Wyman Forum Global Voices Survey, August 2026 (minutes) and April 2026 (verification, wrong information); Oliver Wyman Forum analysis

Survey questions: “For every 60 minutes you spend completing work with AI’s help, about how many minutes go to checking, correcting, or redoing AI output?” (mean); “When you use AI for important information, how often do you verify its output elsewhere?” (“Always”); “Have you personally experienced a situation in the past 12 months where AI gave you information that turned out to be wrong and caused a real problem?” (“Yes, and it had significant consequences” plus “Yes, but it was minor”)

Base: Minutes N=7,845 employees who use AI at work (“I cannot estimate this” excluded); verification N=11,312 (employed respondents, “not applicable” excluded); wrong information N=11,565 (employed respondents)

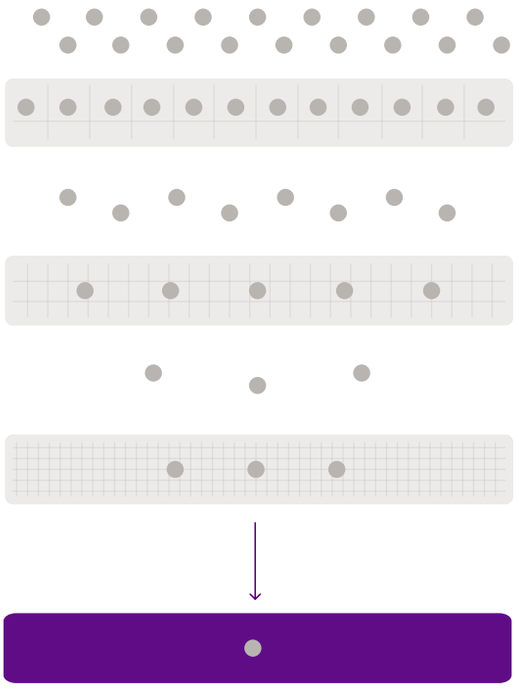
The answer is an automated control system that does more of the routine oversight before a person ever sees a case. Rules limiting the data, tools, and actions available to the AI should be in place before first use. During use, the system is monitored and tested to make sure that what it does continues to meet the required standard. Fixed checks catch known violations. Where quality requires judgment, a second model can help score outputs at scale — often called LLM-as-a-judge — but it must be calibrated against human examples and paired with hard checks.

Automated responses constrain or stop the system when thresholds are crossed. Only ambiguous, novel, or high-consequence cases should reach people.

Human judgment still matters, but where and how it is exercised is changing. People should review cases that require context, discretion, or accountability while the system handles repeatable controls at machine speed. Good risk management employs human judgment where it matters most.

Layered controls save scarce human judgment for the few cases that need it

Every case the AI handles
Without these layers, a person would have to check every one



1. Boundaries set in advance
Settles what can be prevented: limits on what the AI can access, checks on what goes in and out, caps on what it can do

2. Automatic monitoring
Settles what can be handled automatically: a record of what the AI did, tests against the required standard, a response when it drifts

3. Escalation rules
Passes on only the cases that are unclear, new, or high-stakes

Human review
People decide only the cases no layer above could settle

Note: Dots are illustrative.
Source: Oliver Wyman Forum analysis

The age of AI requires *living governance* oversight that moves at the speed of the technology

Traditional enterprise IT changed at visible moments: a patch, a new version, a new integration. AI can change more rapidly and subtly, such as when a provider updates the model, new data or prompts alter behavior, a tool is connected, or a use spreads beyond its approval. Failure patterns that were invisible in a pilot become material at scale. The risk can move even when the organization has not formally changed the system.

Living governance adjusts guardrails to AI's evolving capabilities and use
Operational evidence continuously updates the governance decision



What can reopen the governance decision

Model change provider updates the model	**Data or prompt change** different inputs, different behavior	**Tool or integration change** new reach for the system	**Use-case change** the use stretches beyond approval
Scale change failure patterns appear at volume	**Autonomy change** the system acts more, proposes less	**Consequence change** a mistake matters more	

Source: Oliver Wyman Forum analysis

Living governance is the response. More than continuous monitoring, it is a frequent, change-aware approach in which approval begins the oversight cycle rather than ending it. The company sets permissions and controls at the outset, observes how the system behaves in operation, and tests performance against agreed standards. That evidence flows back into the governance decision. Permissions widen or narrow. Controls tighten or relax. Escalation changes. Material shifts trigger reassessment. Approval is renewed, rescoped, paused, or withdrawn.

Review triggers, therefore, extend beyond scheduled cycles and major releases. New models, data, prompts, tools, uses, scale, autonomy, or consequence can reopen the decision. The same evidence should show whether the system remains reliable, valuable, within its risk boundaries, and economic to operate, giving business, technology, operations, and risk one view of the system.

Scheduled reviews improve the system over time, but they cannot handle failures that unfold in hours. Fast-moving problems need automatic responses designed in from the start — the AI equivalent of the brakes and telemetry that let a racing car run at the limit. Governance becomes a recurring cycle of permission, observation, learning, adjustment, and renewal, moving at the speed of the technology it governs.

That is what effective AI risk management looks like in practice. Competitors can buy the same models. The advantage lies in the system around them: an approved path employees will use, a deliberate way to test risk, reusable decisions, oversight built into production, and a feedback loop that changes the rules as the system changes. Good governance makes that discipline scalable. It allows companies to pursue more consequential uses, recover faster when something goes wrong, and protect customer trust while continuing to expand what AI can do.

5



The terms of *dependence*

For countries and companies, AI sovereignty means keeping enough control over the AI they depend on without cutting themselves off from the global frontier

Authors
Abhishek Sharma, Darren Goldie, Jad Haddad, Jose Lopez, Mario Rizk, Matthieu De Clercq, Xavier Boileau

No one is a *sovereign island* in the world of AI

The dictionary definition of sovereignty is freedom from external control. Such freedom is an illusion when it comes to AI. Full ownership of the stack — from energy and chip production, through frontier models and the talent behind them, to proprietary data and software — is out of reach for any single country or company. Even the leading economies cannot make an advanced AI chip alone: The required tools, factories, and refined minerals are divided among a few countries and suppliers.

The realistic goal is controlled interdependence: having enough control over AI capabilities and dependencies to keep critical work running and governable, and a credible way to replace any of those dependencies.

Countries and companies share that goal. Countries are focused on maximizing economic competitiveness, security, freedom of action, and the resilience of critical services. Companies need continuity of operations, predictable costs, and control of the business. Some disruption will be deliberate — suppliers changing terms, states pursuing leverage, criminals seeking gain — and designed to defeat the backups.

In practice, AI sovereignty means choosing the terms of dependence: knowing what to rely on, how hard it would be to replace, and how to react if a critical dependency were compromised.

THE REALITY IS...

1. Every organization now has to *price in AI sovereignty*
2. The *sovereignty premium* should not be paid everywhere
3. You don't need to own a *frontier model* to be sovereign
4. Governments get more sovereignty by *setting terms* than by writing checks
5. You can *partner your way* to greater sovereignty

By the numbers

More than half of companies expect disruption if their main provider goes offline[31]

55%

expect at least moderate disruption if their company's main AI provider becomes unavailable

65%

say their company has not put the ability to switch AI providers in place

Swapping models and the systems built around them presents challenges[32]

Only 0.1%

of customers fully switch cloud providers each year

Up to 60

engineer-days needed to switch primary foundation models

AI sovereignty risk is being recognized faster than it's being managed[33]

46%

of CEOs say supplier and model concentration is formally registered as a board or enterprise risk

10%

of technology leaders have a signed statement defining what data and assets they retain if an AI supplier raises prices or cuts off access

Leaders expect advantage from what they own, not what they rent[34]

74%

of CEOs expect a lasting advantage from redesigned processes and operating models; 64% from proprietary data and customer context

4%

say the same of privileged access to a leading AI model or supplier

Geopolitical and cyber risks are high on the CEO agenda[35]

41%

of CEOs rank geopolitics, trade policy, and industrial policy as a top threat — second only to macroeconomic conditions

32%

cite cybersecurity, data protection, and information integrity, including attacks on critical infrastructure, as a top threat

Global AI chokepoints sit in a small group of countries[36]

17%

of the world's 100 largest economies own or host a hard-to-replace part of the AI supply chain

7

countries own a true chokepoint; another 10 own a narrower critical input or host hard-to-replace capacity

Source: Oliver Wyman Forum Global Voices Survey, August 2026 (footnote 31); UK Competition and Markets Authority (footnote 32); Oliver Wyman 2026 Parallel Surveys of CEOs and Tech Leaders on Agentic AI and IT (footnotes 32, 33, and 34); Oliver Wyman Forum x NYSE CEO Survey 2026 (footnote 35); Oliver Wyman Forum Sovereign AI Index, Q3 2026 snapshot (footnote 36)

31. Survey questions: "How much visibility do you have into what your company spends on AI?"; "If your company's main AI provider became unavailable for one month, how disruptive would it be?"; "Which of the following has your company put in place to reduce its dependence on any single AI provider?" Base: N=5,789 (employed respondents with visibility into their organization's AI spending, August 2026). Note: "At least moderate disruption" combines moderate, major, and critical. Questions were multi-select; 65% is the share who did not select "the ability to switch AI providers." Non-selection does not establish tested switching capability.

32. 0.1%: UK Competition and Markets Authority, Cloud services market investigation, final decision (July 2025), Appendix H. Note: Full switching between the three largest cloud providers, measured year on year from 2020 to 2023; about 0.1% of customers switched in each year. 60 days: Oliver Wyman 2026 Parallel Surveys of CEOs and Tech Leaders on Agentic AI and IT, fielded July to August 2026; N=200. Survey question: "How many engineer-days would it take to switch your primary foundation model on your main use case?"; 21 to 60 engineer-days is the typical impact estimated by technology leaders.

33. Survey questions: "Is supplier and model concentration formally registered as an enterprise or board-level risk in your organization?" and "If your primary model supplier doubled prices or withdrew your model version, could you show in a signed written statement which artifacts stay under your control?" Base: Oliver Wyman 2026 Parallel Surveys of CEOs and Tech Leaders on Agentic AI and IT, fielded July to August 2026; N=200. Note: 46% is the share of CEO respondents; 10% is the share of technology-leader respondents.

34. Survey question: "As frontier models converge in capability and fall in price, where do you think your durable competitive advantage will sit?" Base: CEO respondents, as footnote 33. Note: 4% is the share selecting privileged access to a leading AI model or supplier as a lasting source of advantage; 74% and 64% are the shares selecting redesigned processes and operating models, and proprietary data and customer context, respectively. Respondents could select up to two responses.

35. Survey question: "What are the top three threats for your business over the next three years?" Base: N=415 (total CEO survey). Note: 41% and 32% are the shares ranking each item as a top threat to their business; geopolitics, trade policy, and industrial policy ranked second after macroeconomic conditions.

36. Note: The index covers the 100 largest economies. The 17 economies (17%) include true chokepoint owners, narrower critical-input owners, and hosts. "No position identified" is an evidence boundary, not proof of no relevant activity.

Every organization now has to *price in AI sovereignty*

For a company or a country, relying on one model, one cloud, or one provider is often the cheapest and simplest starting point. But as an organization builds its workflows around that provider, switching becomes harder and more expensive. Protecting the ability to change providers is one of the costs of AI sovereignty, and that cost rises the longer a single provider stays embedded.

In cloud computing, the option to switch is rarely exercised; the UK competition authority recently found that only about one in a thousand cloud customers switched providers in a typical year. Companies are applying the lesson from cloud at the model layer, where swapping models still presents challenges. But dependency is forming around the model: in the agent memory, project context, orchestration software, and automation that accumulate as critical workflows are built, often across several suppliers. Once that constellation of providers is lined up, it is hard to reshuffle. Half of technology leaders estimate that switching just the primary foundation model on their main use case takes 21 to 60 engineer-days; the agent memory, context, and orchestration built around it add to that bill. Platforms that route work across several providers' models hedge only the layer that was already easier to switch.

Dependency also can be shaped by commercial arrangements. Discounts tied to multiyear spending commitments, licensing terms, and bundled services can make an alternative uneconomic even when switching is technically possible. The decision to exit boils down to two questions: Can the organization make a move, and can it afford to move?

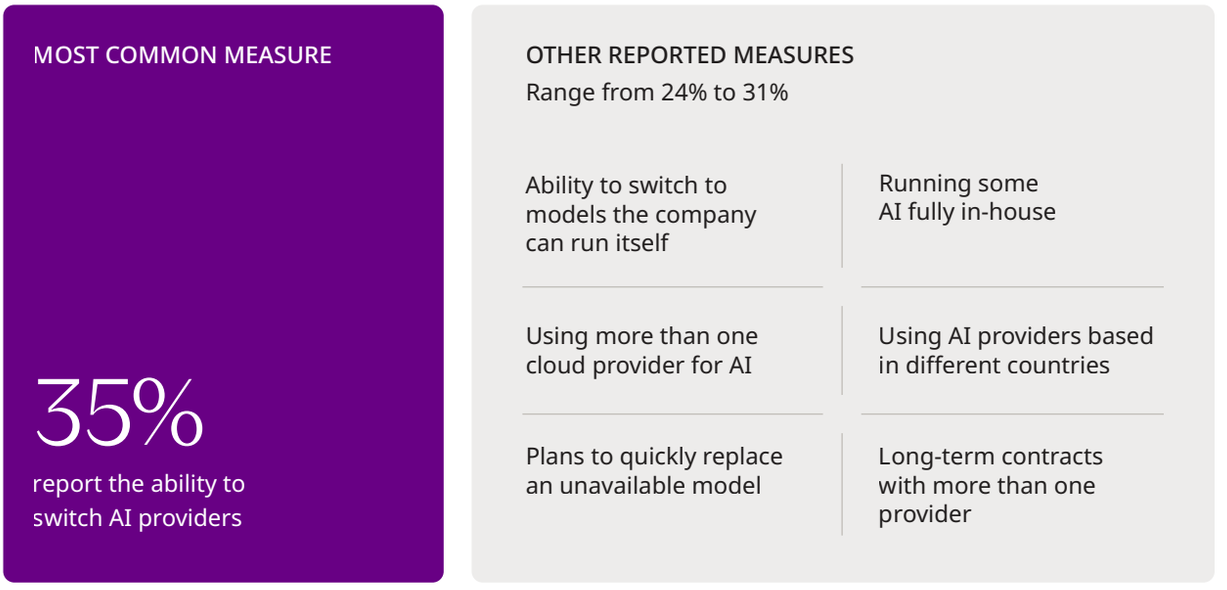
Regulators reached the same conclusion about cloud services, treating provider concentration as a resilience risk and requiring firms to maintain tested exit plans for critical providers — but only after nearly two decades of accumulated dependence. AI-enabled systems are being embedded into core operations on a much shorter timeline, so firms cannot assume supervision will arrive first.

21 to 60
engineer-days needed to switch the primary foundation model

Preparation lags awareness of the risk. Forty-six percent of CEOs say supplier and model concentration is formally registered as a board or enterprise risk, and 55% of companies expect at least moderate disruption if their main AI provider becomes unavailable. Yet only 35% have set up the ability to switch providers, and just 10% of technology leaders have a written statement, signed by their executive sponsor, of which data and assets stay under their control if a key supplier raised prices or cut off access. These are self-assessments; in a real shock, fallback plans could fare worse than expected. For a critical dependency, a fallback plan that has not been funded and tested is not a reliable option.

Organizations are hedging AI provider dependence in multiple ways, but no safeguard is yet standard

Measures in place to reduce dependence on any single AI provider reported by respondents with visibility into company AI spending



Note: Respondents could select multiple measures.

Source: Oliver Wyman Forum Global Voices Survey (August 2026)

Survey question: “Which of the following has your company put in place to reduce its dependence on any single AI provider?”

Base: N=5,789 (respondents with visibility into company AI spending)

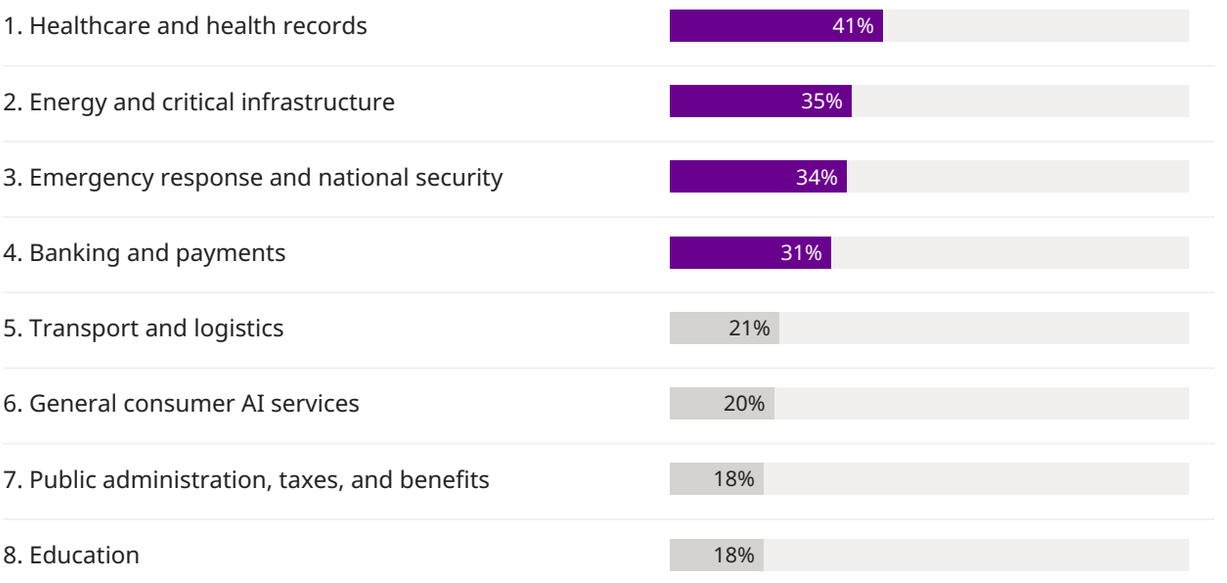
The *sovereignty premium* should not be paid everywhere

Not every AI service or workflow needs the same protections against disruption. A consumer-facing app can tolerate delay or substitution, but interruptions of AI in health records, energy systems, and payments can cause wider harm. One standard for every service could waste money where the stakes are low while leaving critical services vulnerable.

The cost of sovereignty, whether it means paying more or accepting less capable AI, should rise with the potential harm and fall for more replaceable dependencies. Asked which services should have the most resilient AI access, people in 12 of 15 markets put healthcare first, ahead of energy, national security, and payments. Public pressure will favor the services people see, but the sovereignty premium should also follow the chain of dependence down to the energy and critical infrastructure those services run on.

People prioritize national AI sovereignty in essential services and industries

Services that respondents say they want their country to maintain if there's a prolonged loss of access to foreign AI providers, ranked



Note: Respondents could select up to three services.

Source: Oliver Wyman Forum Global Voices Survey (August 2026)

Survey question: "Which AI-enabled services should your country be able to keep operating during a prolonged loss of access to foreign AI providers?"

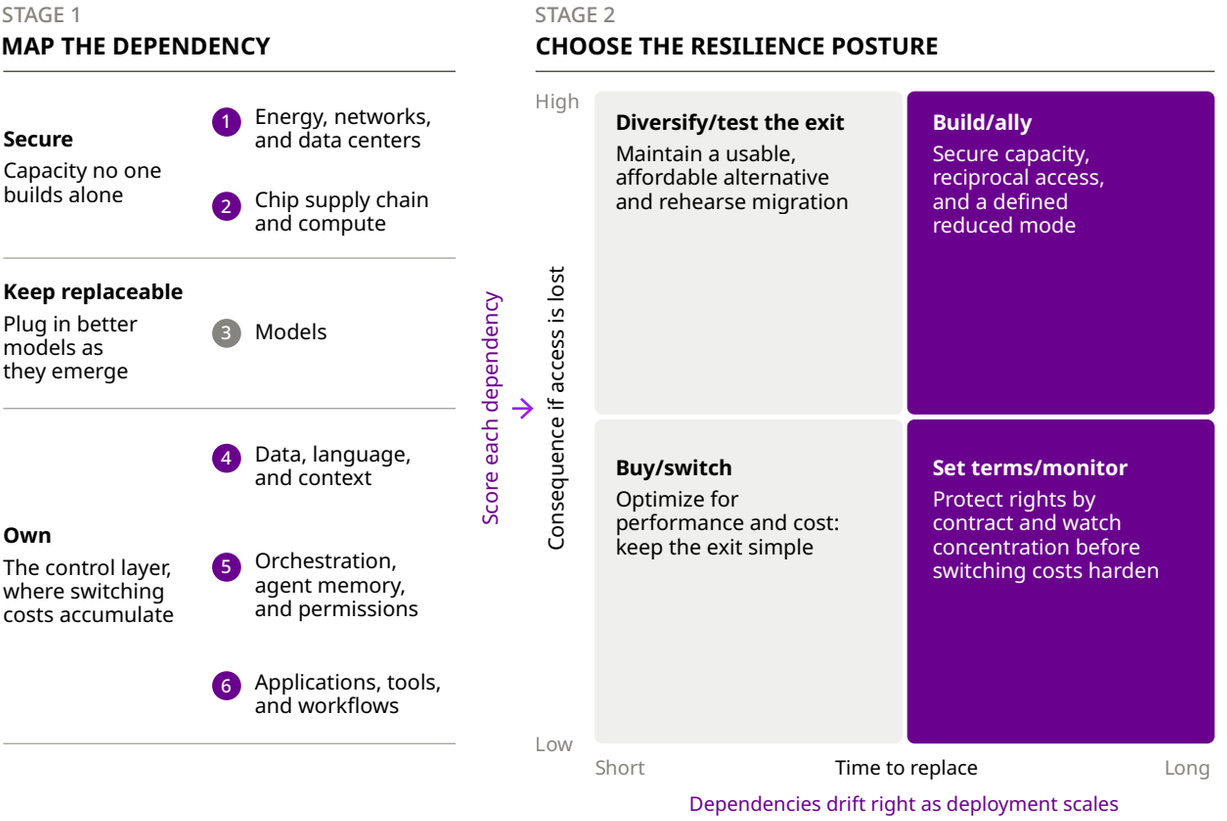
Base: N=12,837 (all respondents)

Prioritization is only the first step. Any extra cost paid for greater control and resilience should buy a specific protection against a defined failure, regardless of the cause. Local hosting addresses jurisdiction and access risk, while redundant capacity, portable data, contractual rights, and trusted suppliers address different failures. None is interchangeable, and even several together can leave a shared dependency hidden.

The sovereignty premium should target the environment around the model: the computing power, energy, and connectivity an alternative needs to run, and the data architecture, identity, evaluation, tools, and workflows that put the AI to work. Whether the model, software, and integrations come from one supplier or several, an exit starts to resemble a rebuild rather than a switch.

Where substitution is slow or costly, firms should have plans for implementing a smaller model, narrower functionality, or more human handling to keep essential functions running until a full solution is in place.

AI sovereignty stack: The AI model is often easier to replace than everything built around it



CROSS-CUTTING CAPABILITIES

Talent and capital

The expertise and funding to build, evaluate, integrate, operate, and replace critical capabilities

Security

The ability to protect infrastructure, models, data, identity, tools, and access

Operating capability

The processes and institutional capacity to procure, integrate, monitor, recover, and switch in practice

Note: Purple shading marks dependencies that are slow or costly to replace.

Source: Oliver Wyman Forum analysis

You don't need to own a *frontier model* to be sovereign

Leaders do not expect advantage to come from the AI model itself: Only 4% of CEOs say privileged access to a leading model or supplier will be a lasting advantage, while 74% point to redesigned processes and operating models and 64% to proprietary data and customer context. AI sovereignty follows the same logic. Critical work needs enough capability under the company's control to keep running, and that is often possible with less than frontier performance. For most workflows, reliable access, organizational context, security, and integration matter more. Frontier capability is worth its added dependency only where the performance gap changes crucial outcomes.

Open-weight models have trailed the frontier by a matter of months, and the computing power needed for a given level of performance has halved roughly every eight months. At that pace, today's frontier capability becomes cheap and controllable within a year. And unless the demands of routine work rise as fast as capability does, tomorrow's frontier models will be overkill for most of that work.

Even employees who prefer company AI still use public tools regularly

Among respondents with access to custom or in-house AI

56%

would keep the company's AI if they could retain only one tool

65%

of those who would keep company AI still move tasks between company and public AI at least weekly

Source: Oliver Wyman Forum Global Voices Survey (August 2026)

Survey questions: "If you could only keep one for your work, which would you choose?"; "How often do you manually move information or work between a public AI tool and company AI tool because each is better at a different part of the task?"

Base: N=2,278 (respondents with access to custom or in-house AI); N=1,268 (respondents who would keep company AI)

Inside companies, controlled and frontier AI are already complements rather than substitutes. Among respondents with access to custom or in-house AI, 56% would choose the company tool if they could keep only one, yet 65% of that group still move tasks between company and public AI at least weekly. Company tools, which trade capability for control, win on internal context and data protection. Public tools, which offer the frontier directly, win on speed, quality, and breadth.

Organizations can maximize their AI sovereignty by owning the control layer while keeping the frontier replaceable. The control layer is the part of the model's environment that accumulates through use: the data, context, and evaluation that critical work depends on, as well as the memory, permissions, tool connections, and workflow state that AI agents build up inside orchestration software. By keeping those portable, a company can preserve the ability to plug in better models as they emerge. Frontier access then becomes more of a managed option than a dependency.

Control also can mean proximity. Where critical work must continue without a reliable connection to centralized infrastructure — say, in a vehicle, facility, clinic, or deployed unit — smaller models running at the edge preserve essential functions.

Control is also cultural. A model shapes what users see. One trained largely outside a country's language, institutions, and legal tradition carries blind spots and biases. Companies inherit them. Anchoring models in users' context is a sovereignty premium worth paying.

Only 4%

of CEOs say privileged access to a leading model or AI supplier will be a lasting advantage

Governments get more sovereignty by *setting terms* than by writing checks

Regulation and procurement redirect demand, investment, and corporate behavior at a scale that public budgets rarely reach. Procurement rules create anchor customers. Where national security is involved, government is often the dominant buyer and can define the market outright. With hundreds of billions of dollars flowing into AI globally each year, the most effective sovereign AI investment may be redirecting private capital through the levers of government rather than spending public funds.

That leverage runs in both directions. A well-calibrated government requirement gives more suppliers enough demand to reach scale, and each additional supplier increases optionality for the private sector. But a requirement set at a level only the largest providers can meet entrenches the supplier concentration it was meant to reduce. Localization rules split the same way: They attract real capacity where the economics work, but where they don't, such rules simply add cost and a local address to the same foreign stack. Government policies should be judged by whether they increase the number of providers a country or company can actually switch to.

Some constraints still call for direct investment or partnership. Regulation alone rarely produces energy, computing power, or specialist talent at the speed required. Public investment can accelerate all three. But government investment in national AI capability faces political and fiscal constraints. Across 15 markets, three-quarters of residents support major public spending to build AI capability, yet only 19% support it if it requires higher taxes or spending cuts elsewhere. This constraint is tightest in Europe and loosest in Asia and the Gulf. AI may also squeeze the budgets that would fund it: If it displaces taxable work and adds public spending on retraining and income support, governments will have less to spend. Few governments will find the fiscal room or political leeway to fund national AI capability at the scale their ambitions imply. That leaves most of the work to the instruments that do not require public money.

Only 19%

of residents across 15 markets support major public spending to build AI capability if it requires higher taxes or spending cuts elsewhere

Support for national AI investment is broad, but how conditional it is varies sharply by market

Conditional and unconditional support for AI investment, by market

Market	Conditional	+	Unconditional	=	Total support	Conditional support as share of total support
Italy	56%		7%		63%	88%
Spain	60%		10%		70%	86%
France	47%		11%		58%	82%
Canada	49%		11%		60%	81%
Mexico	62%		15%		77%	81%
Germany	49%		13%		62%	80%
Brazil	62%		16%		79%	79%
South Africa	62%		17%		79%	79%
United Kingdom	53%		14%		67%	79%
United States	44%		17%		62%	72%
UAE	62%		27%		89%	69%
Hong Kong	54%		25%		79%	68%
Saudi Arabia	57%		33%		91%	63%
China	56%		37%		93%	61%
India	53%		39%		92%	57%

Global: 74%

Note: Conditional = support only if investment does not raise taxes or reduce other spending; unconditional = support even if taxes/cuts are required. Conditional support share is calculated from unrounded respondent counts as conditional support divided by total support. Markets are ranked by this normalized measure.

Source: Oliver Wyman Forum Global Voices Survey (August 2026)

Survey question: “Would you support your government making a major public investment to build national AI capabilities?”

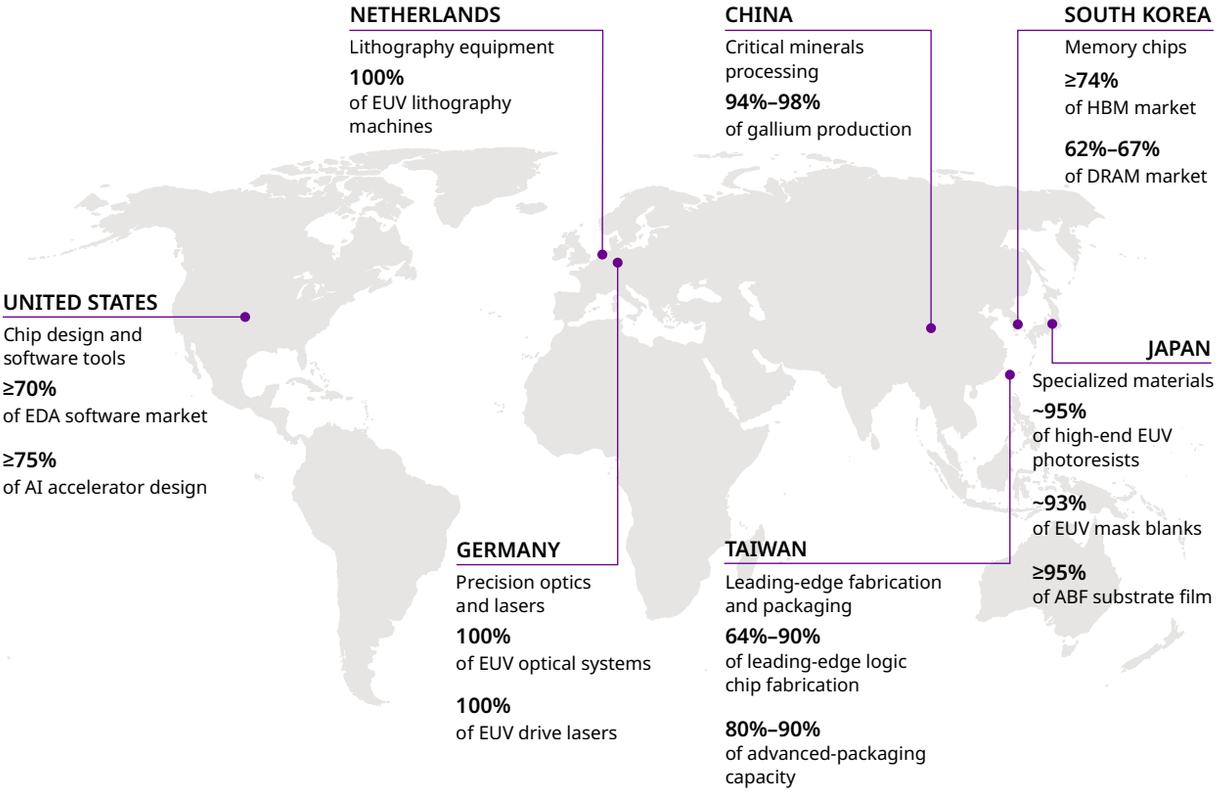
Base: N=12,837 (all respondents); market Ns=503–936

You can *partner your way* to greater sovereignty

No country holds every capability that matters for sovereign AI. Few have the scale on their own to influence the terms of global suppliers. This makes alliances and partnerships crucial. Our Sovereign AI Index makes the constraint concrete: Hard-to-replace positions in the AI supply chain sit in just 17 of the world's 100 largest economies, and only seven hold a true chokepoint. For many countries and companies, sovereignty means credible access to more than one technology ecosystem and the freedom to combine elements of each. Partnerships can add capacity, bargaining power, and routes around a disrupted relationship.

Chokepoint examples: No single economy controls the full AI chip supply chain

Critical capabilities for advanced AI chips are concentrated in different economies, creating mutual dependencies



Note: Percentages are estimated shares only; the fabrication range reflects differing definitions of the leading edge.
Source: Oliver Wyman Forum analysis

A partnership strengthens sovereignty only when it widens the choices available. It should add a supplier, an architecture, a route to capacity, or expertise that the partners lacked on their own. Many arrangements fail that test on technical grounds. Unless systems interoperate so that data, identity, and tools move between them, a second provider adds cost without adding choice. Companies build choice through interoperable architectures, joint procurement, and shared evaluation; countries do so through reciprocal access, shared capacity, and research and procurement commitments. Where scarcity is structural, as with critical minerals, semiconductors, energy, and specialist talent, the arrangement has to be durable because replacement takes years.

Supplier diversity is not dependency diversity. Two separate providers may still rely on the same chipmaker, power grid, undersea cable, or export-control regime. A partnership that leaves everyone dependent on the same underlying provider or stack doesn't strengthen anyone's sovereignty.

17

of the world's 100 largest economies have hard-to-replace positions in the AI supply chain, and only seven hold a true chokepoint

No economy leads every capability that sovereign AI requires

Strengths in models, research, compute, supply chain chokepoints, private capital, and government readiness are distributed across economies, and even the economy that leads on most counts, including chokepoints, cannot make an advanced AI chip on its own. That is what makes interdependence unavoidable.

China
1. AI research workforce
2. Model leadership
3. Verified onshore AI compute

France
2. Government AI readiness
3. Model leadership
7. AI research workforce

Germany
4. AI supply chain chokepoints
5. AI research workforce
6. Recorded private AI capital

India
3. AI research workforce
5. Model leadership

Italy
5. Recorded private AI capital
8. AI research workforce
10. Model leadership

Japan
5. AI supply chain chokepoints
6. AI research workforce
7. Model leadership

Netherlands
3. AI supply chain chokepoints
4. Government AI readiness

South Korea
4. Model leadership
5. Government AI readiness
6. AI supply chain chokepoints

United Kingdom
2. Recorded private AI capital
3. Government AI readiness
4. Verified onshore AI compute

United States
1. In five of six dimensions
2. AI research workforce

Note: Profiles show selected dimension ranks only and are not an overall ranking.
Ranks show position, not distance; small differences in rank should be interpreted cautiously.

Source: Oliver Wyman Forum Sovereign AI Index. Data snapshot Q3 2026

Different economies hold leverage at different points in the AI system

Model leadership

Significant AI models developed by organizations associated with each economy

1. United States	6. Canada
2. China	7. Japan
3. France	8. Russia
4. South Korea	9. Singapore
5. India	10. Italy

88%

of large-scale AI models captured by the index are attributable to the United States and China

AI research workforce

Published AI research authors affiliated with institutions in each economy

1. China	6. Japan
2. United States	7. France
3. India	8. Italy
4. United Kingdom	9. Canada
5. Germany	10. South Korea

69%

of the indexed AI research workforce sits in these ten economies

Verified onshore AI compute

Operating AI hardware the index could verify as physically located in each economy, regardless of ownership

1. United States	6. Portugal
2. Malaysia	7. South Korea
3. China	8. Japan
4. United Kingdom	9. Germany
5. Indonesia	10. Finland

91%

of the verified AI compute captured by the index is hosted in the United States

AI supply chain chokepoints

Ownership or domestic hosting of technologies and capacity judged hard to replace in the AI supply chain

1. United States	6. South Korea
2. Taiwan	7. China
3. Netherlands	8. Switzerland
4. Germany	9. France
5. Japan	10. United Kingdom

17

economies own or host an in-scope chokepoint position; seven own a true chokepoint

Recorded private AI capital

Private investment tagged to AI and assigned to each economy in the 2025 source data

1. United States	6. Germany
2. United Kingdom	7. France
3. China	8. Canada
4. Qatar	9. Singapore
5. Italy	10. Israel

78%

of the recorded 2025 AI-related private capital went to the United States

Government AI readiness

Institutional capacity to adopt, govern, and support AI, as measured by the Oxford Insights Government AI Readiness Index

1. United States	6. Germany
2. France	7. Singapore
3. United Kingdom	8. Australia
4. Netherlands	9. Norway
5. South Korea	10. Denmark

Note: China has been intentionally excluded from ranking due to insufficient coverage in the source material

Note: Government AI funding is omitted from this publication view because the current evidence does not cover several material markets, including China and the United Arab Emirates, and therefore does not support a credible comparative ranking. No replacement composite ranking has been calculated. Verified onshore AI compute reflects documented operating hardware and may undercount capacity in markets with limited public disclosure. AI research workforce counts published authors and does not represent the full engineering or practitioner workforce. Recorded private AI capital reflects one year of deal data and may be influenced by a small number of large transactions. Callout shares describe evidence captured by the index, not complete global totals. The index does not directly score energy, availability, data control, language or cultural context, orchestration, applications, enterprise adoption, or operating security. Small differences in rank should be interpreted cautiously.

Source: Oliver Wyman Forum Sovereign AI Index. Data snapshot Q3 2026. The index covers the world's 100 largest economies

Sovereignty means *having options*

The Sovereign AI Index settles the question of whether anyone can go it alone. The United States hosts 91% of the verified AI computing power in the index and leads five of six dimensions, yet it cannot make an advanced chip without Dutch lithography, Japanese materials, Korean memory, Chinese critical minerals, and Taiwanese fabrication. The leader depends on others as surely as others depend on the leader. And the economies with the most leverage relative to their size did not buy it; each became hard to replace in one category.

That is the lesson the five reality checks share. Every decision on sovereignty comes down to one test: Does it add a credible alternative to keep critical work running? The test cuts both ways. The assets CEOs expect to deliver lasting advantage — proprietary data and processes, not privileged access to a model — are also what make an organization hard to replace, and therefore worth partnering with.

No country or company is a sovereign island, and none needs to be. The question is how to set the terms: what to own, what to secure through partners, what to keep replaceable, and what cannot be allowed to fail.

Every decision on sovereignty comes down to *one test:* Does it add *a credible alternative* to keep critical work running?

Methodology

This report draws primarily on three proprietary research streams, complemented by the applied experience of Oliver Wyman experts. By respondent base, the largest source is the Oliver Wyman Forum Global Voices survey; the next-largest primary survey is the Oliver Wyman Forum/NYSE CEO Agenda 2026; and the report also draws on the Oliver Wyman 2026 Parallel Surveys of CEOs and Tech Leaders on Agentic AI and IT. The research further draws on the experience of Quotient, AI by Oliver Wyman, including insights from client deployments and work on agentic AI, value realization, and operating-model change. These sources are complemented by proprietary analysis, executive conversations, and secondary research. Each exhibit identifies its source, survey question, and respondent base so readers can trace every figure to the data behind it.

Oliver Wyman Forum Global Voices survey. The largest source is the Oliver Wyman Forum Global Voices survey. This report uses two waves fielded in 2026: April (N=12,883) and August (N=12,837). Both waves covered the same 15 markets: the United States, Canada, Mexico, Brazil, the United Kingdom, France, Germany, Italy, Spain, South Africa, Saudi Arabia, the United Arab Emirates, India, China, and Hong Kong (reported separately from mainland China). Each market had roughly 900 respondents, except Saudi Arabia and the UAE at roughly 500 each.

Respondents were adults age 18 to 75, surveyed online in local languages, with quotas on age and gender. The survey focused on the employed population; about 96% were employed, so findings for "consumers" or "all respondents" reflect a predominantly working population rather than the general public.

The surveys were sourced from a panel of 67 million people worldwide. To support representative distributions, our respondent pool generally mirrored the demographics of each country, including age, income, education level, political affiliation, and gender. In a small number of waves, certain demographic segments may have been modestly over- or underrepresented due to sampling constraints; where this occurred, results were still intended to be broadly representative at the country level. We also applied standard quality controls to remove low-quality responses.

Each wave combined fixed questions on demographics, employment, and attitudes with modules on AI in the workplace, governance and risk, shopping, and sovereignty. Many questions were asked only of relevant subsets, such as employed respondents, workplace AI users, or white-collar employees by job level. Some analyses also excluded "not sure" responses. Bases therefore varied across exhibits, and the same question could produce different shares depending on the population analyzed.

Throughout the report, we generally used global findings across the 15 markets rather than presenting country-by-country results. Aggregation may smooth over meaningful local differences in attitudes, adoption, regulation, or market context. We reviewed the major findings for directional consistency across markets.

In most cases, while magnitudes differed, the directional story held across the markets analyzed. Even so, this is global research rather than a comprehensive assessment of every local market, and some country-specific differences may not be fully captured.

CEO and technology-leader surveys. CEO findings come fromthe Oliver Wyman Forum x NYSE CEO Agenda 2026, an online survey of 415 CEOs and CEO-equivalents fielded January to March 2026. Respondents were recruited through the NYSE's network, the Oliver Wyman Forum's network, and specialist research partners. By region, 56% of respondents were based in North America, 32% in Europe, 9% in Asia-Pacific, and 4% elsewhere. By revenue, 45% led companies with under $1 billion, 31% with $1 billion to $5 billion, and 24% with over $5 billion. Public companies made up 63% of the sample and private companies 36%, across 10 industry groups.

Year-on-year comparisons draw on the 2025 edition (165 CEOs, March to April 2025). That sample was predominantly North American and used different question wording, so these comparisons are directional only. The report also draws on the Oliver Wyman 2026 Parallel Surveys of CEOs and Tech Leaders on Agentic AI and IT (200 respondents, July to August 2026).

Analytical segments and proprietary analysis. Several survey findings rely on segments we constructed for analysis. These include "AI superworkers" (workplace AI users who say AI at least doubles the work they can complete), companies that employees describe as ahead of or behind competitors on AI adoption, and CEOs grouped as AI ROI leaders or laggards based on the financial impact they report.

Beyond survey data, the Oliver Wyman Forum Sovereign AI Index ranks the 100 largest economies across six dimensions of AI capability, drawing on sources including the Oxford Insights Government AI Readiness Index. Three proprietary analyses of AI shopping inform the findings on agentic commerce: simulated recommendations, a product-page audit, and simulated purchase journeys. These are illustrative samples, and AI outputs can vary with model, prompt, and timing. The 2030 estimate of AI-influenced e-commerce is a directional scenario based on survey adoption data and market growth assumptions. Conceptual frameworks and illustrative exhibits labeled "Oliver Wyman analysis" represent our perspective and are not empirical estimates.

Secondary research and interpretation. Finally, parts of this report are informed by conversations with senior executives and by secondary research used to contextualize and triangulate our findings. Secondary sources include Challenger, Gray & Christmas layoff reporting, Stanford Digital Economy Lab research using ADP payroll data, the UK Competition and Markets Authority's cloud services market investigation, and public reporting on market developments.

Survey responses were subject to standard quality checks. Results were self-reported and descriptive unless otherwise noted, and percentages may not sum to 100 because of rounding or multi-select questions. Readers with questions about bases, markets, or definitions should please reach out.

Clarity in an age of acceleration

In a world of AI hype, endless possibilities, and growing complexity, *The Reality Check* offers a moment of clarity. The minimalist cover reflects this intent, cutting through the noise to focus on what truly matters.

Turn the page, and the world is in motion. People move through shifting environments, reflecting a world in transformation and the human decisions shaping its future. Woven into these scenes, ASCII-inspired patterns of characters and symbols represent AI as an increasingly embedded presence in our everyday world, accelerating the pace of change.

Together, these visual expressions reflect a central idea: **Technology may change the pace, but humanity determines the direction.**

Note: All section images featured in this report were generated using artificial intelligence.

Authors and contributors

CHAPTER 1

The *return on intelligence* (ROI)

Aaron Fine
Ana Kreacic
Ashik Ardeshna
Edwin Anderson
Eric Czervionke
John Romeo
Simon Luong

CHAPTER 2

Rewiring the enterprise for the *bionic workforce*

Ana Kreacic
Helen Leis
Rupal Kantaria
Simon Luong
Terry Stone

CHAPTER 3

The market for *machine attention*

Adrian Oest
Dan Jones
Emma Stall
Joshua Gilbert
Michael Zeltkevic
Muriel Schtickzelle
Rupal Kantaria
Simon Luong

CHAPTER 4

How *guardrails* can enable growth

Daniel Cope
Elena Belov
Frédéric Thomas-Dupuis
Nick Dykstra
Ted Moynihan
Til Schuermann

CHAPTER 5

The terms of *dependence*

Abhishek Sharma
Darren Goldie
Jad Haddad
Jose Lopez
Mario Rizk
Matthieu De Clercq
Xavier Boileau

CONTRIBUTORS

Aashna Sethi, Alice Thanur, Annie Liu, Arran Yentob, Athan Siah, Aymen Saleh, Charlie Zhang, Charlotte Fuller, Chris Murphy, Christina Kyriakides, Claudio Lago de Lanzos, Dan Kleinman, Dustin Irwin, Gilles Roucolle, Heather Stern, Jocelyn Cao, Johaan Wiggins, Laksh Maggoo, Lulu Smith, Marissa Lynch, Monica To, Nick Liptak, Peta Egan, Pierre Bechon, Pieter Durman, Robbie White, Robert Hunter, Sayli Chitre, Tom Buerkle

ART AND DESIGN

Adrien Slimani, Brenda Battaglia, Cindy Sim, Melissa Ureksoy, Ramona Pillai, Simon Crofts, Weronika Talaj, Yireli Pale

About Oliver Wyman Forum

The Oliver Wyman Forum is a community of senior leaders in business, government, and society dedicated to addressing the biggest forces shaping our future. Through forward-looking research and dialogue, we turn diverse perspectives into clearer thinking and bolder action — reshaping industries, navigating global change and complexity, and reimagining leadership.

For more information, visit oliverwymanforum.com

About Oliver Wyman

Oliver Wyman, a Marsh (NYSE: MRSH) business, is a management consulting firm driven by deep industry insight, bold innovation, and a collaborative approach that cuts through complexity to help organizations navigate their most defining transformative moments. Marsh is a global leader in risk, reinsurance and capital, people and investments, and management consulting, advising clients in 130 countries. With annual revenue of $27 billion and more than 95,000 colleagues, Marsh helps build the confidence to thrive through the power of perspective.

For more information, visit www.oliverwyman.com

About Quotient, AI by Oliver Wyman

Quotient, AI by Oliver Wyman, brings together strategy, technology, and implementation to unlock new avenues for growth, enhance enterprise performance, strengthen resilience, and effectively deploy capital across the AI value chain. We work side by side with clients to identify the opportunities worth pursuing, build AI into the way organizations operate, and create the conditions to mitigate risk. Deep industry knowledge helps us focus on what will move the business; hands-on AI expertise helps us build it; and relationships with leading technology partners help us move from idea to measurable results.

Learn more about Quotient, AI by Oliver Wyman

Copyright 2026 Oliver Wyman

All rights reserved. This report may not be reproduced or redistributed, in whole or in part, without the written permission of Oliver Wyman and Oliver Wyman accepts no liability whatsoever for the actions of third parties in this respect. The information and opinions in this report were prepared by Oliver Wyman. This report is not investment advice and should not be relied on for such advice or as a substitute for consultation with professional accountants, tax, legal or financial advisors. Oliver Wyman has made every effort to use reliable, up-to-date and comprehensive information and analysis, but all information is provided without warranty of any kind, express or implied. Oliver Wyman disclaims any responsibility to update the information or conclusions in this report. Oliver Wyman accepts no liability for any loss arising from any action taken or refrained from as a result of information contained in this report or any reports or sources of information referred to herein, or for any consequential, special or similar damages even if advised of the possibility of such damages. The report is not an offer to buy or sell securities or a solicitation of an offer to buy or sell securities. This report may not be sold without the written consent of Oliver Wyman.