HFS

In partnership with

ALTIMETRIK

MARKET IMPACT REPORT

Humans at the Helm of AI

Why enterprise AI deployment has outrun human control

Authors:

Dana Daher, Executive Research Leader
Phil Fersht, CEO and Chief Analyst
Hridika Biswas, Associate Practice Leader
Mayank Madhur, Practice Leader

Contents

Executive summary

AI is already making decisions in your organization. You just haven't decided who owns them yet.

Most enterprises have responded to this reality the same way: put a human in the loop and call it governance. It is not governance, but a feeling of governance. And the gap between the two is where accountability goes to die.

Being in the loop means reviewing outputs. Being at the helm means owning what the machine decides, defining when humans override it, and being able to answer, before something goes wrong and not after, who is accountable. That distinction sounds simple. Closing that gap is one of the hardest challenges enterprises face today, and the cost of delay is compounding.

This is a last-mile problem. Enterprises have deployed AI, but have not designed the human authority, capability, and accountability needed to govern it.

HFS Research partnered with Altimetrik to survey 505 senior executives across Global 2000 organizations to understand how AI decisions are made, who owns the outcomes, how confident workforces are, and how accountability travels across partners and platforms.

What we found is a consistent pattern of breakdown across five dimensions:

1

The helm is empty

Only 14% of organizations have a clear AI strategy with defined goals and outcomes. The CEO owns AI accountability in 6% of organizations day to day, but shows up in 20% of post-incident conversations. Ownership lives in tech. Consequences land in the boardroom.

2

The loop is hollow

Fifty-three percent (53%) name human-in-the-loop as their primary governance mechanism. Only 18% can actually interrogate the reasoning behind what they are approving. The thing enterprises trust most to keep them safe is the thing they have invested in least.

3

Built to comply, not to govern

Fifty-two percent (52%) cite fear of replacement as their biggest barrier to AI engagement. Seventy-two percent (72%) fear being judged if experiments fail. Nearly 80% receive fewer than 10 hours of AI training a year. Enterprises are asking people to govern AI while making it both risky and under-supported.

4

Transition without a plan

More than half expect AI to reduce roles in the next two to three years. Most plan to let it happen through attrition. Only 7% of employees feel in control of what comes next.

5

Accountability without borders

Eighty-three percent (83%) depend on partners to move quickly. Eighty percent (80%) say accountability is unclear when a partner makes the wrong call. The governance failure does not stay inside the organization. It travels.

Humans at the helm is not a cultural slogan. It is an operating shift. It starts the moment leadership stops asking how fast we can scale AI and starts asking whether authority has been redesigned before autonomy is extended.

SECTION 01

The AI velocity gap is compounding

AI is everywhere. Enterprise autonomy is not.

AI is moving into enterprises faster than the institutions meant to guide it. Models are improving rapidly, organizations are embedding AI into everyday workflows, and leaders are under pressure to deploy these capabilities quickly. Yet the human systems surrounding this technology, leadership structures, governance frameworks, workforce capability, and decision authority, are evolving far more slowly.

This imbalance creates what we call the AI velocity gap. It is the widening distance between how quickly enterprises deploy intelligent systems and how slowly they redesign the human systems required to govern them.

AI is active across the enterprise, but institutionalization is rare

Nearly half of organizations (46%) are applying AI to defined tasks or individual workflows. But activity is not maturity. Institutionalization, the point at which AI becomes a standardized, continuously optimized capability rather than a collection of team-level experiments, remains uncommon. Only 13% of enterprises have reached it.

Exhibit 1: AI maturity is unevenly distributed

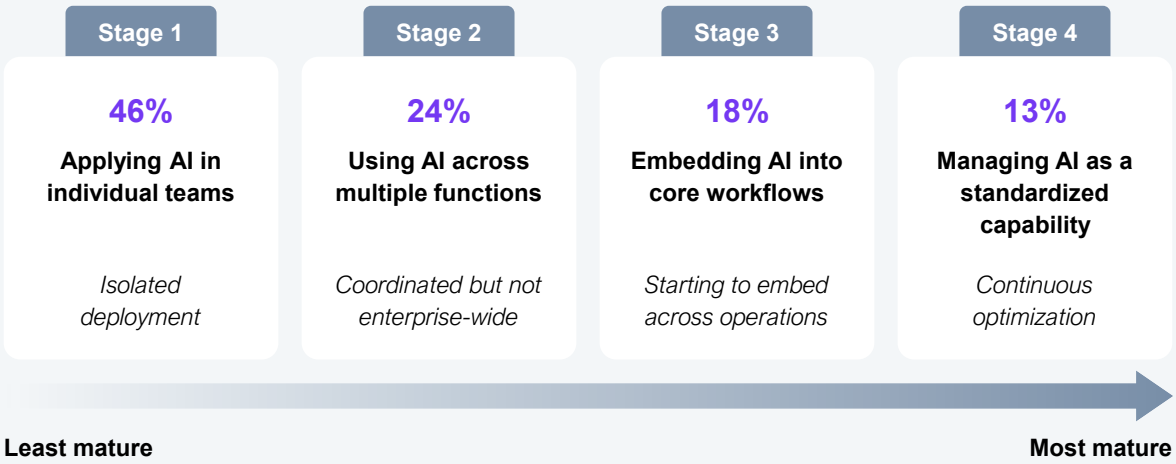

Sample: 505 executives across Global 2000 enterprises
Source: HFS Research, 2026

At the same time, most AI still operates in a supervised execution zone. Enterprises are automating predefined tasks or generating recommendations for human review but rarely extending AI into workflows where it holds meaningful authority. Only 7% operate with high autonomy in critical workflows, the small frontier where compounding advantage begins and where the design of human authority matters most. (See Exhibit 2)

Exhibit 2: Most AI is supervised; very little is truly autonomous

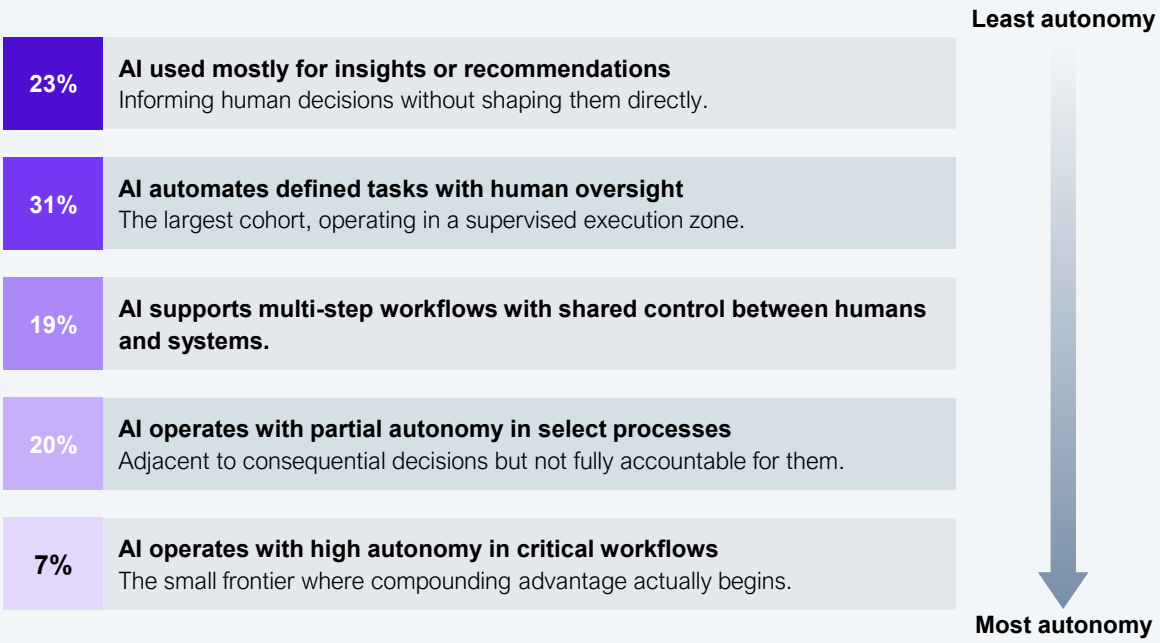

Sample: 505 executives across Global 2000 enterprises
Source: HFS Research, 2026

The velocity gap is the distance between maturity and autonomy, and it is where divergence begins

The organizations that have closed that gap are more than twice as likely to report faster, more accurate decision making, and are dramatically more likely to deliver measurable customer and revenue impact. The performance difference is not marginal. It is compounding in one direction. (See Exhibit 3)

The velocity gap persists not because enterprises lack ambition or good models. It persists because the human systems required to govern AI have not kept pace. Leadership has not declared where AI is going. Authority over what it decides has not been documented. Workforces lack the confidence and capability to challenge it, and accountability dissolves the moment it crosses an organizational boundary.

The organizations that have closed the gap did not do it by deploying better models. They did it by answering a question most enterprises have never formally asked: What does the human at the helm of this system actually have the authority, visibility, and accountability to do? That is the question this report is built around. The sections that follow examine why most enterprises cannot answer it, and what it takes to change that.

Exhibit 3: The AI velocity gap is compounding

AI constrained by structural friction	**Positioning**	AI operating as a compounding decision engine
Only 39% report faster and more accurate decision making, less than half the rate of mature peers	**Decision advantage**	82% report faster and more accurate decision making, a 43-point advantage
Just 21% achieve double-digit improvement in meeting rising customer expectations	**Customer impact**	88% achieve double-digit CX improvements at scale
0% report double-digit new revenue or product impact from AI initiatives	**Revenue and growth**	41% report double-digit impact in new products or revenue expansion
47% require 6–12 months to move from idea to implementation	**Execution speed**	26% move from idea to deployment in weeks, more than twice the rate of early-stage peers
25% cite lack of clear AI ownership, 8x the rate of more mature organizations	**Ownership and governance**	Only 3% report unclear AI ownership
39% report insufficient AI skills or confidence	**Capabilities**	Just 14% cite skills or confidence as a barrier
—	**Trust**	Distrust in AI outputs drops to 5%

Sample: 505 executives across Global 2000 enterprises
Source: HFS Research, 2026

The bridge forward

Humans at the helm is the answer to the velocity gap—not as a constraint on AI, but as the design that makes AI authority governable, defensible, and compounding in the right direction.

SECTION 02

The helm is empty

Strategy is absent, ownership is misplaced, and accountability appears only when something goes wrong.

Enterprises are not failing to adopt AI. They are failing to lead it. The tools are deployed, and the budgets are committed. What is missing is a declared destination, documented decision rights, and accountability structures that follow outcomes rather than org charts. Until those exist, the helm is empty, not because leadership has abdicated but because no one has been asked to sit in it.

AI deployment is accelerating faster than leadership commitment

Most enterprises are not executing a strategy. They are managing a portfolio of experiments that has been left to find its own direction. Most are still developing a strategy (39%) or operating in disconnected pockets (32%). Without a declared destination, AI keeps moving while leadership debt compounds quietly behind it. (See Exhibit 4)

Exhibit 4: Strategy is thin—most organizations are still developing direction or operating in fragments

Which statement best reflects the current state of your organization's AI strategy?

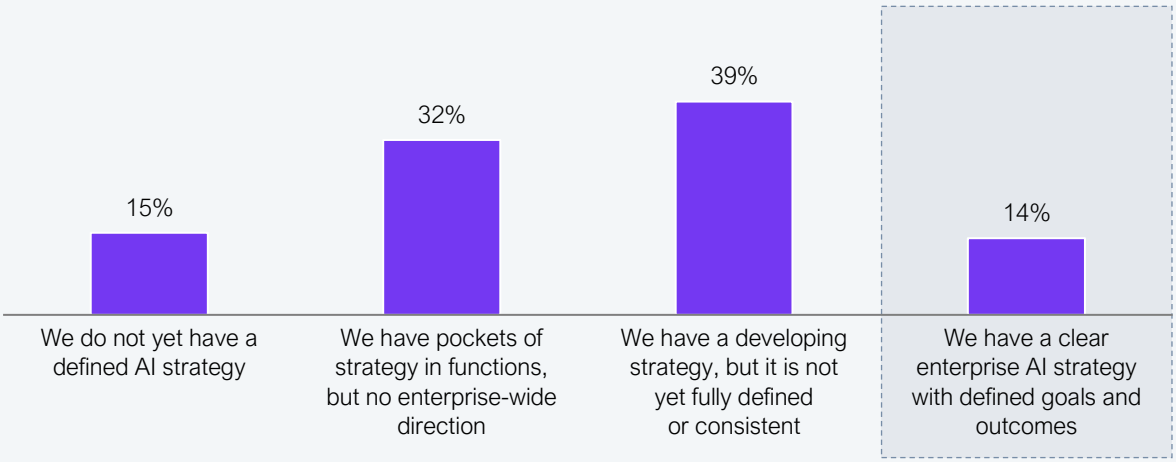

Sample: 505 executives across Global 2000 enterprises
Source: HFS Research, 2026

Cost reduction is not a strategy; it is what fills the space where strategy should be

In the absence of strategic direction, AI gets justified through the safest available narrative. Cost reduction is the top driver for 52% of organizations. Revenue impact still lags far behind at 15%. Cost reduction requires no vision, no ownership model, and no declared direction. It survives every board presentation precisely because it commits to nothing. (See Exhibit 5)

Consider a global manufacturer that deployed AI across procurement and back-office operations with a single mandate: reduce costs. It did. Eighteen months later, leadership wanted to use the same infrastructure to build new revenue streams. They could not. The systems had been optimized purely around elimination. Nobody had decided what the enterprise wanted AI to help it become, and by the time anyone asked, the architecture had already answered.

Exhibit 5: The pressure driving AI adoption: cost reduction leads, reinvention trails

What pressure is most directly driving your organization to accelerate AI adoption today?

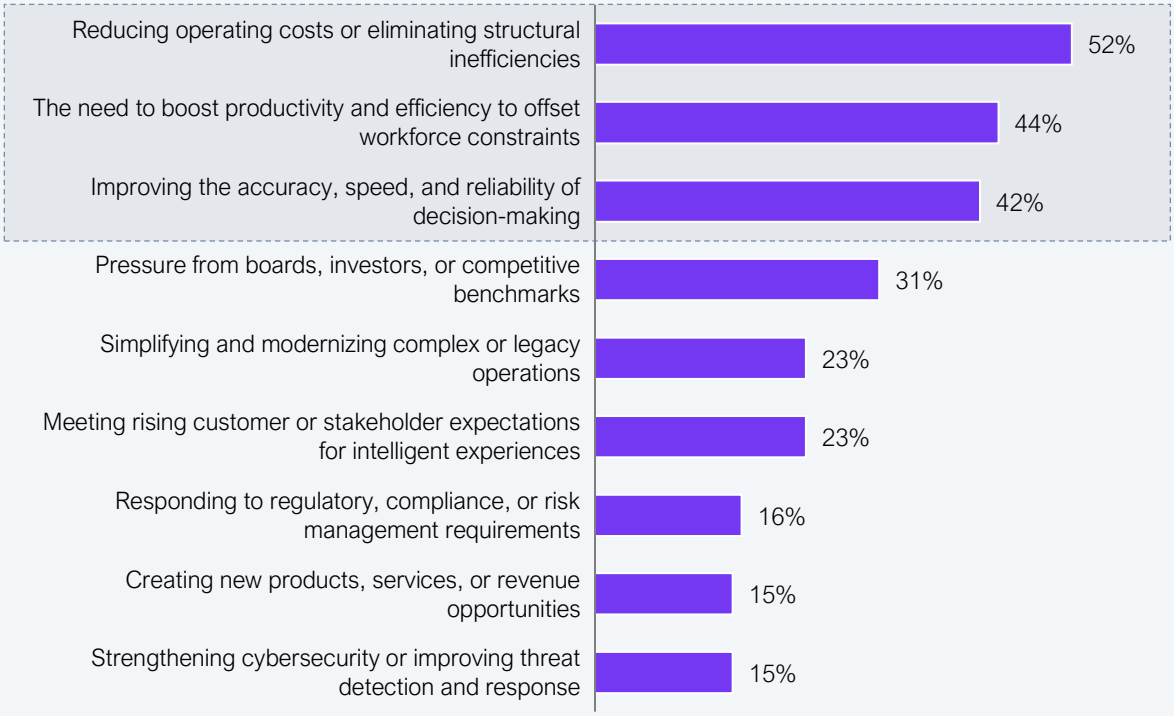

Sample: 505 executives across Global 2000 enterprises
Source: HFS Research, 2026

Ownership diffuses down, and accountability flows up

AI accountability rests with the CIO, CTO, or technology function in 37% of organizations day to day. The CEO owns it in 6%. But when AI initiatives fail, CEO or executive team involvement in the accountability conversation more than triples, rising to 20%. (See Exhibit 6)

When the people responsible for AI performance are not the same people responsible for business performance, the lessons from AI results accumulate in the wrong place. Technology teams learn what broke. Business leaders learn that something went wrong. The accountability loop is not closed. It is triggered by failure and then reset.

For the CIO or CTO, the exposure is specific. You are accountable for deployment, for cost, and for the post-incident conversation, but not for the strategic decisions that would have prevented the failure. That is not a technology problem. It is an authority design problem, and it requires a business leadership response.

Exhibit 6: Who owns AI strategy and outcomes: accountability concentrated in tech, away from consequences

In your view, who ultimately owns accountability for AI strategy and outcomes in your organization?

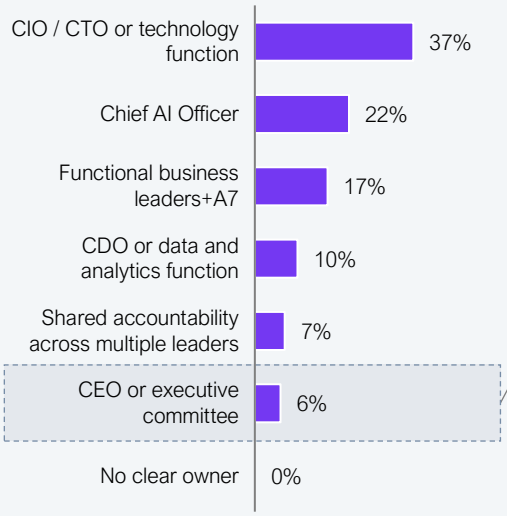

When AI initiatives succeed or fail, who leads the discussion about what happened?

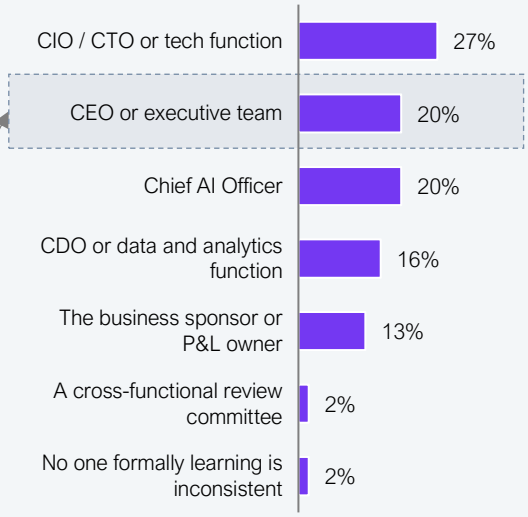

Sample: 505 executives across Global 2000 enterprises
Source: HFS Research, 2026

The bridge forward

Putting humans at the helm begins here. Not with tools, training, or governance frameworks. It begins with the deliberate acts of naming a destination, documenting who decides what, and building an accountability structure that follows outcomes rather than org charts.

SECTION 03

The loop is hollow

When conflict happens, no one knows who wins. When an override is needed, no one can see what to override.

Most enterprises believe they have solved the AI governance problem. Ask how they maintain control over AI systems, and the answer is consistent: humans are in the loop. Someone reviews the output before it becomes a decision. Someone can step in if something looks wrong. The loop exists. Therefore, governance exists.

The data reveals that the loop is largely hollow.

When AI and humans disagree, only 26% say human judgment clearly prevails

Start with the most basic test of human authority: what happens when AI and human judgment conflict? If humans are genuinely at the helm, the answer should be straightforward. It is not. Only 26% say human judgment clearly prevails in those moments. Thirty percent (30%) resolve the conflict through joint reviews. Twenty-nine percent (29%) say the outcome varies by function or process. Fourteen percent (14%) say AI recommendations often carry more weight. For three in four enterprises, the question of who wins when AI and humans disagree has no documented answer. It depends on who is in the room. (See Exhibit 7)

That is not a governance system. It is a negotiation with no rules, repeated across thousands of decisions, with no consistent principle determining the outcome.

Exhibit 7: When AI and humans disagree, decisions are negotiated, not owned—only 26% say human judgment clearly prevails

In your organization, when AI recommendations differ from human judgment, how are decisions typically made?

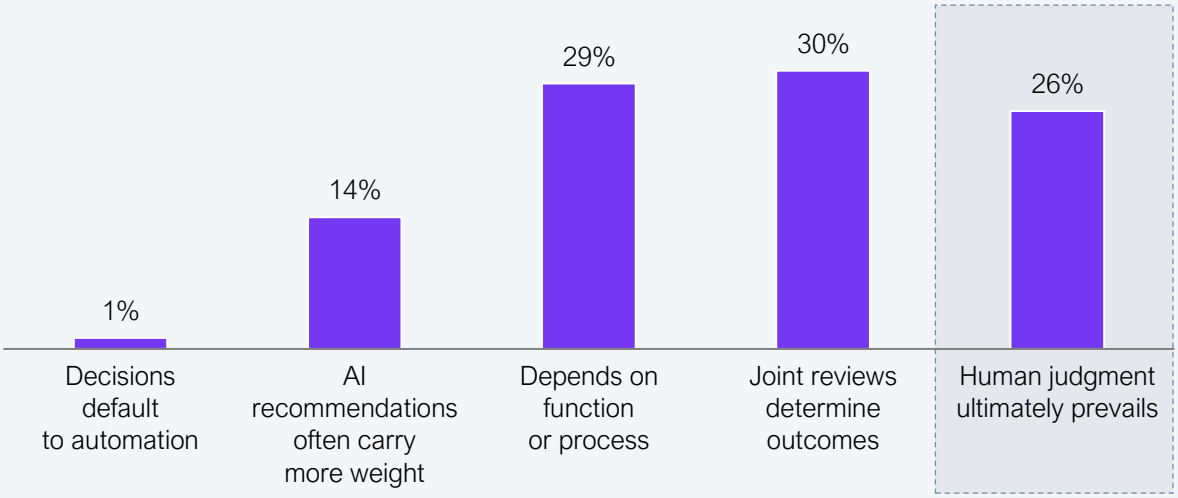

Sample: 505 executives across Global 2000 enterprises
Source: HFS Research, 2026

Only 18% can interrogate the decisions they are supposed to be overriding

Now test the second assumption: that the humans reviewing AI outputs can see what they are reviewing. Only 18% of organizations have clear visibility into both what AI recommends and the reasoning behind it. Thirty-two percent (32%) understand the outputs but not the underlying logic. Seven percent (7%) say their teams rely on AI decisions they do not fully understand. The remaining majority sit somewhere in between, approving things they can describe but not explain. (See Exhibit 8)

An approval you cannot justify is not a governance act. It is a signature on something you cannot read. And if something goes wrong, the human who approved it becomes the accountability surface for a decision they were never actually equipped to make.

Exhibit 8: Only 18% can clearly see how AI decisions get made—the rest are overseeing outputs they cannot fully interrogate

How well do leaders and teams understand how AI decisions are made in your organization?

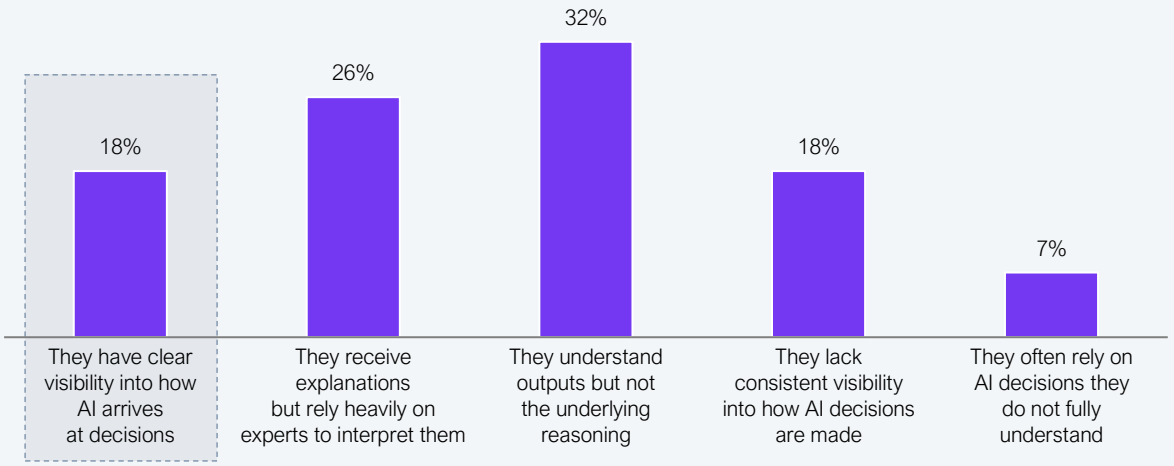

Sample: 505 executives across Global 2000 enterprises
Source: HFS Research, 2026

The primary trust mechanism is the one enterprises have done the least to make work

Fifty-three percent (53%) of enterprises name human-in-the-loop processes as their primary mechanism for building trust in AI, ranking ahead of explainability tools, bias audits, and training programs. But the humans in that loop cannot define when their judgment prevails, and most cannot see into the reasoning behind what they are approving. The governance mechanism enterprises depend on most is the one they have invested in least. (See Exhibit 9)

Enterprises have not failed to put humans in the loop. They have failed to make the loop mean anything.

Exhibit 9: Human-in-the-loop is the primary trust mechanism, but it is the least enabled

What are the top three mechanisms your organization relies on to build trust in AI systems?

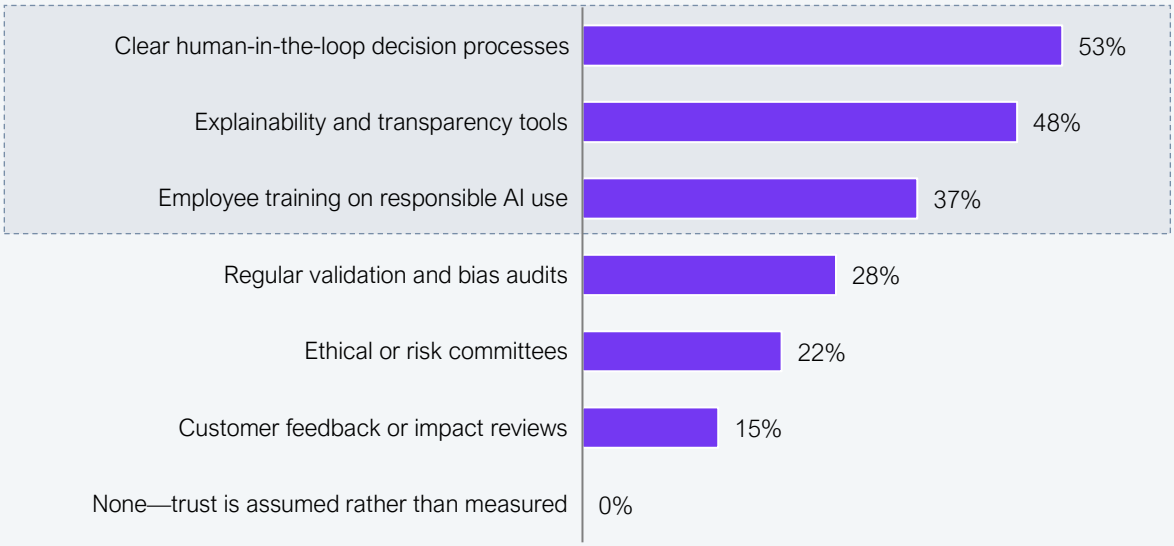

The bridge forward

The question is not whether humans should be in the loop. It is whether the humans in the loop have been given the authority, visibility, and architecture to do anything meaningful with that position.

SECTION 04

Built to comply, not to govern

Enterprises expect humans to govern AI they are not equipped to use. Fear, not capability, is the primary brake.

Governance requires people who can interrogate AI, challenge it, and override it when it is wrong. Most enterprises have built the conditions that make all three of those things professionally risky. Not by design. By default.

Fear, not access, is the biggest barrier to AI confidence

Ask employees what is stopping them from engaging critically with AI, and the answer is not tools, time, or access. It is fear. Fifty-two percent (52%) cite perceived stigma or fear of being replaced as their biggest barrier. Thirty-five percent (35%) fear making mistakes or attracting negative visibility. The barriers that rank last are the practical ones that enterprises are most likely to solve.

Exhibit 10: 52% say fear, not access or time, is the biggest barrier to AI confidence

What barriers most limit employees from gaining confidence in using AI?

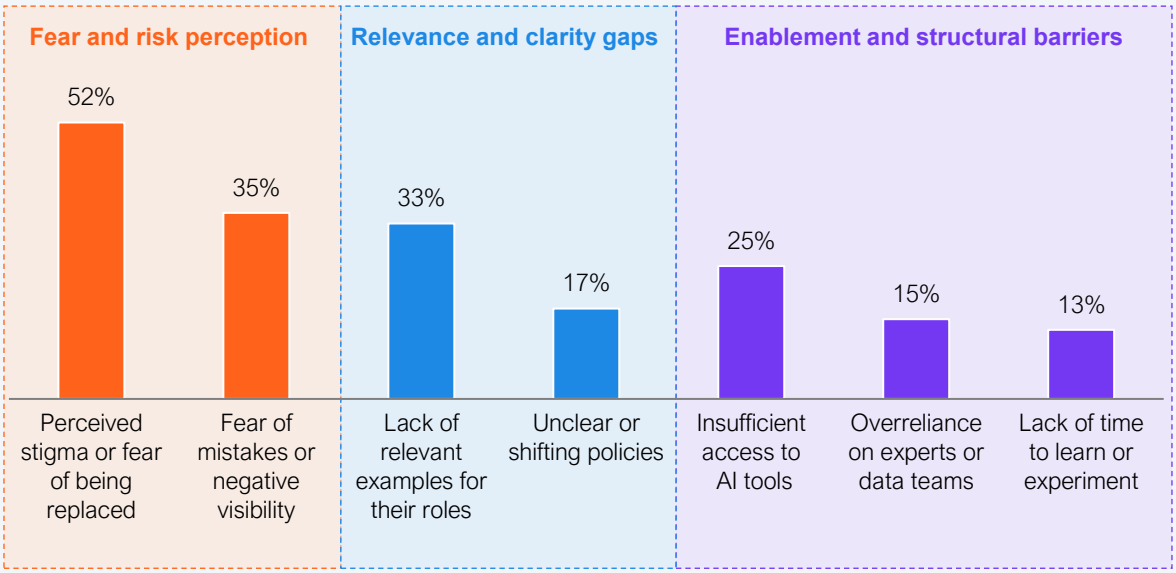

Sample: 505 executives across Global 2000 enterprises
Source: HFS Research, 2026

72% of employees fear being judged if experiments fail—confidence cannot develop where failure is penalized

A slim majority of employees feel safe trying new things with AI (52%). But 72% fear being judged if those experiments fail. Only 40% say they understand the boundaries of safe experimentation.

Employees who are afraid of being judged for getting it wrong will not experiment. Employees who do not experiment will not develop judgment. And enterprises that penalize the visible cost of failure are not building a workforce capable of governing AI. They are building one capable of complying with it.

Exhibit 11: 72% fear being judged if experiments fail — safety exists on paper, not in practice

Sample: 505 executives across Global 2000 enterprises
Source: HFS Research, 2026

Only 17% of employees feel confident and empowered; the rest are somewhere between uncertain and afraid

The consequences of unaddressed fear are visible in how the workforce now experiences AI. These are not isolated issues; they form a reinforcing pattern. Fear makes experimentation risky. Limited experimentation prevents judgment from developing. Without judgment, confidence does not form. And when confidence stays low, employees defer to AI rather than challenge it.

Only 17% feel confident and empowered. Seventeen percent (17%) are hesitant, worried about making mistakes or being judged. Twelve percent (12%) are fearful and actively avoid using AI. The majority sit in between, using AI without fully trusting their ability to question or override it.

This is not a mindset problem. It is the predictable outcome of an environment where experimentation is constrained and capability is underdeveloped.

Exhibit 12: How employees feel about using AI at work—from fearful to confident

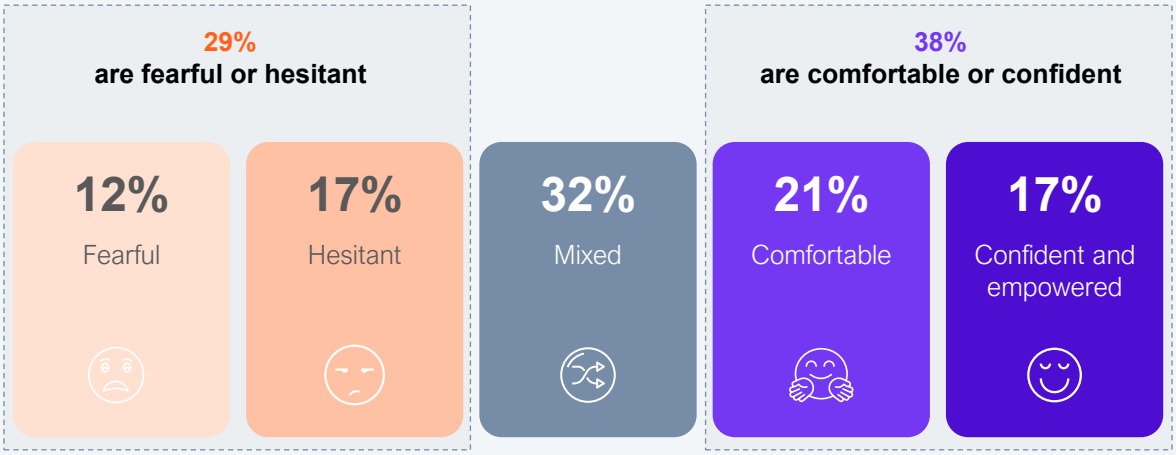

Sample: 505 executives across Global 2000 enterprises
Source: HFS Research, 2026

Most employees receive fewer than 10 hours of AI training a year

Nearly 80% employees receive fewer than 10 hours of formal AI-related training per year, with only 10% report 20 or more hours of structured learning (See Exhibit 13). You cannot build the capacity to challenge AI, override AI, or exercise meaningful judgment alongside AI by giving people less than one working day of training a year and hoping the rest follows. Confidence is a capability, not an attitude that emerges on its own.

Exhibit 13: Most employees get fewer than 10 hours of AI training a year—confidence is expected, not built

Roughly how many hours of formal AI-related learning or training does the average employee complete each year?

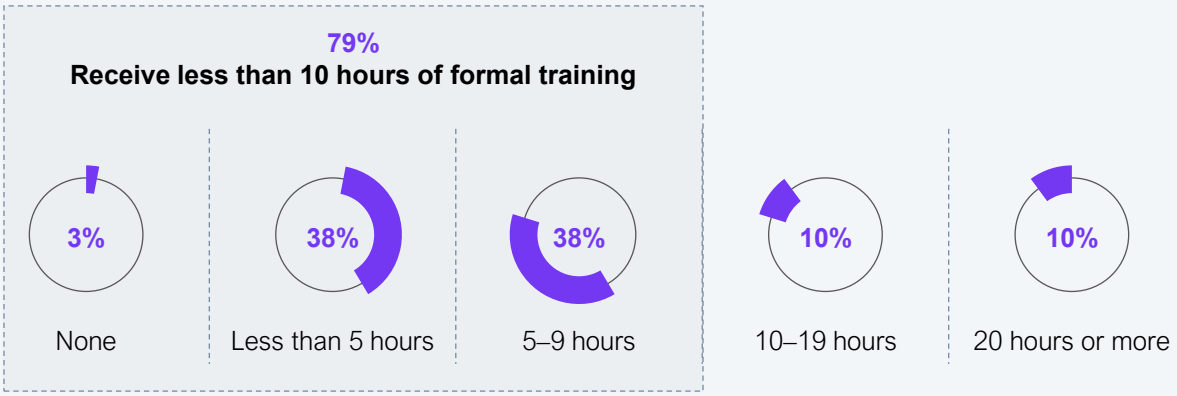

Sample: 505 executives across Global 2000 enterprises
Source: HFS Research, 2026

AI imposter syndrome is widespread

More than four in ten organizations say AI imposter syndrome is either common (30%) or widespread (13%), meaning self-doubt is actively slowing adoption across a significant portion of the workforce (See Exhibit 14).

Imposter syndrome emerges when people are placed in situations where they are expected to perform competently without being given the foundation to do so. When enterprises deploy AI broadly, expect employees to use it, and invest fewer than 10 hours a year in building the capability to do so well, imposter syndrome is not surprising. It is the predictable result of a capability gap that was never addressed.

An employee who does not feel legitimate engaging with AI critically will not challenge it, will not override it, and will not be the human at the helm that the organization needs them to be.

Exhibit 14: 43% say AI imposter syndrome is common or widespread—self-doubt is a structural outcome, not a personal one

How common is AI imposter syndrome in your organization, meaning employees feel they do not know enough to use AI confidently?

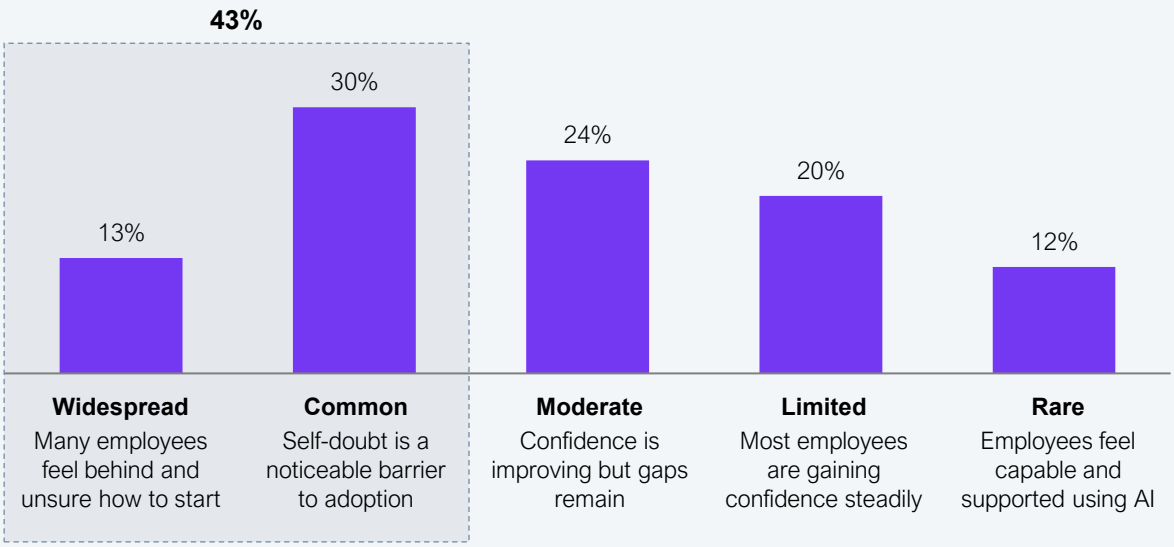

Sample: 505 executives across Global 2000 enterprises
Source: HFS Research, 2026

46% of leaders are concerned employees are already over-relying on AI

Nearly half of leaders say employees are deferring to AI even when human judgment is required. That is the predictable result of everything above. Employees who fear judgment will not experiment. Employees who cannot experiment will not develop judgment. Employees without judgment will defer. The chain is not complicated. And it leads in one direction—toward a workforce that has learned to follow AI rather than govern it.

You cannot build the capacity to govern AI while penalizing the behavior that would develop it

Closing the capability gap requires more than training programs. It requires redesigning the social contract around AI use—making experimentation safe, valuing challenge, and treating the development of human judgment as a deliberate organizational investment rather than something left to accumulate on its own. A workforce built for compliance will follow AI, not govern it.

Exhibit 15: 46% of leaders are concerned employees are already over-relying on AI

How concerned are you that employees may rely too heavily on AI when making decisions?

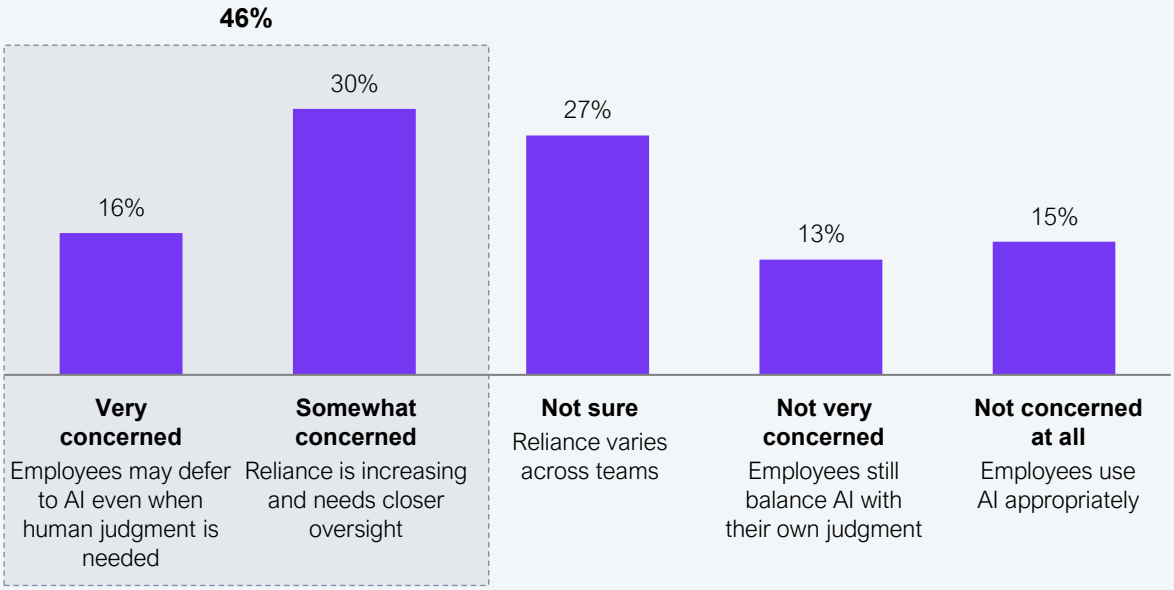

Sample: 505 executives across Global 2000 enterprises
Source: HFS Research, 2026

The bridge forward

Confident humans are not the opposite of capable AI. They are the condition that makes AI governable.

SECTION 05

Transition by drift

AI is already reshaping work, but most of that reshaping is happening without a plan, without a conversation, and without the people it affects most.

AI is not approaching the workforce. It is already inside it. Roles are shifting, work is changing texture, and the experience of showing up and doing a job is different than it was two years ago. What is striking is not the pace of that change. It is how little of it has been designed. The transition is real. The strategy for it is not.

More than half expect role reduction, but most are letting it happen by drift

Fifty-two percent (52%) expect AI to reduce roles over the next two to three years, with most expecting that change to happen through reassignment or natural attrition (33%) rather than active workforce reduction (19%).

Attrition is not a transition strategy. It means the people whose roles are changing are not being told, not being prepared, and not being given any agency over what comes next. The enterprise is not managing a transition. It is waiting for one to happen and calling it planning.

Exhibit 16: 52% expect role reduction—mostly through attrition

How do you expect AI to affect your workforce size in the next two to three years?

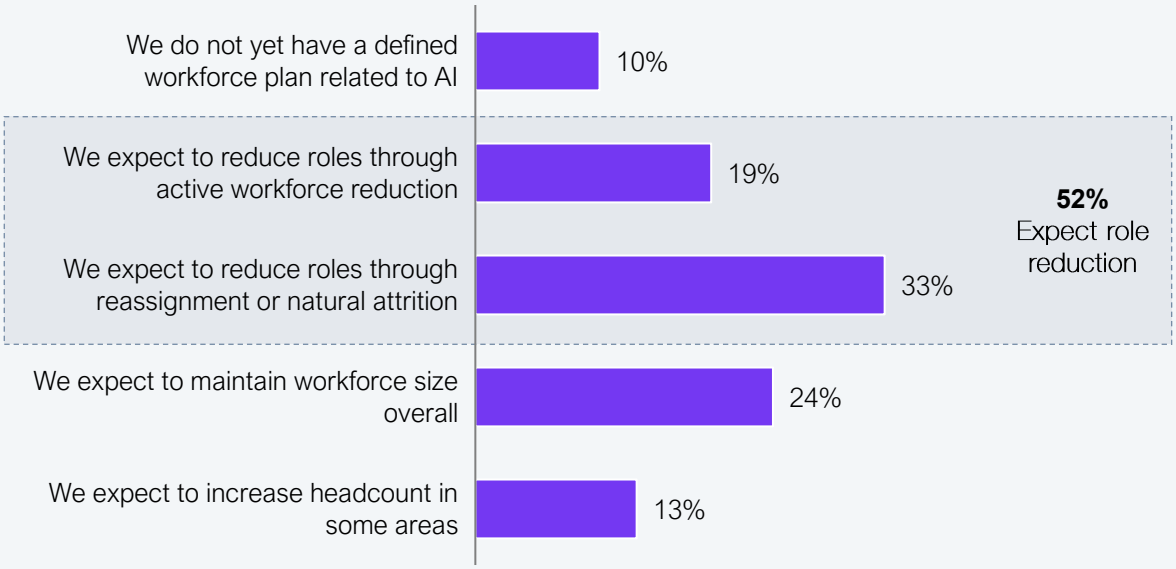

Sample: 505 executives across Global 2000 enterprises
Source: HFS Research, 2026

Only 7% of employees feel in control and shaping outcomes

Most leaders believe employees see themselves executing work directed by AI, monitoring outputs, or feeling increasingly secondary to AI. Only 7% think employees feel in control and shaping outcomes.

AI is being experienced as something done to people, not with them. If the people inside the enterprise feel like passengers in a system they did not design and cannot direct, then there are no humans at the helm.

Transition by drift is a leadership choice, not an inevitability

Putting humans at the helm of the transition requires the same deliberate act as putting humans at the helm of AI governance: someone has to decide it matters, name what the enterprise will look like on the other side, and build a path that brings people with it rather than leaving them to find out what happened after the fact.

Exhibit 17: Only 7% of employees feel in control and shaping outcomes—most feel like passengers

How do you believe employees perceive their role in your organization's AI journey?

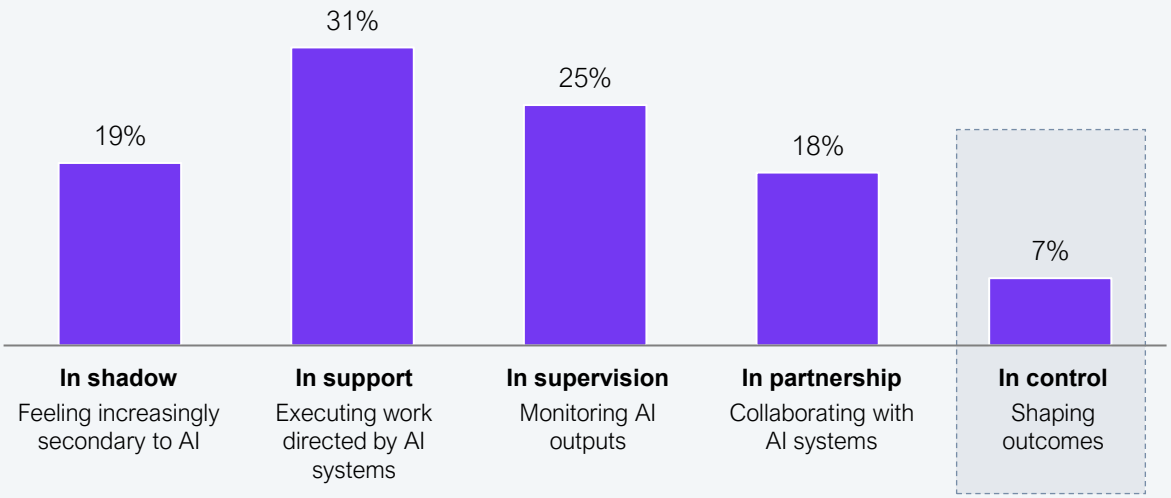

Sample: 505 executives across Global 2000 enterprises
Source: HFS Research, 2026

The bridge forward

The workforce transition is not something AI is doing to enterprises. It is something enterprises are choosing not to design.

SECTION 06

Accountability without borders

Enterprises are outsourcing AI speed to partners faster than they are defining who owns the outcome.

Every governance failure we've discussed has an external dimension that extends beyond the organization to every partner, platform, and system integrator building and running AI on the enterprise's behalf. The helm is not just empty inside the organization. In many cases, it has been handed to someone outside it.

80% say accountability is unclear when an AI partner makes the wrong call—the helm has no address

Enterprises are leaning on AI partners to move faster, with 83% citing internal process friction as the reason. The speed is real. So is the accountability gap: 80% say responsibility is unclear when AI fails.

When the machine is wrong and no one knows who owns the outcome, the human at the helm does not exist. There is only a liability with no address. Enterprises have not outsourced a service. They have outsourced a decision, failing to document who is responsible for it.

Exhibit 18: Speed is outsourced, but accountability is not

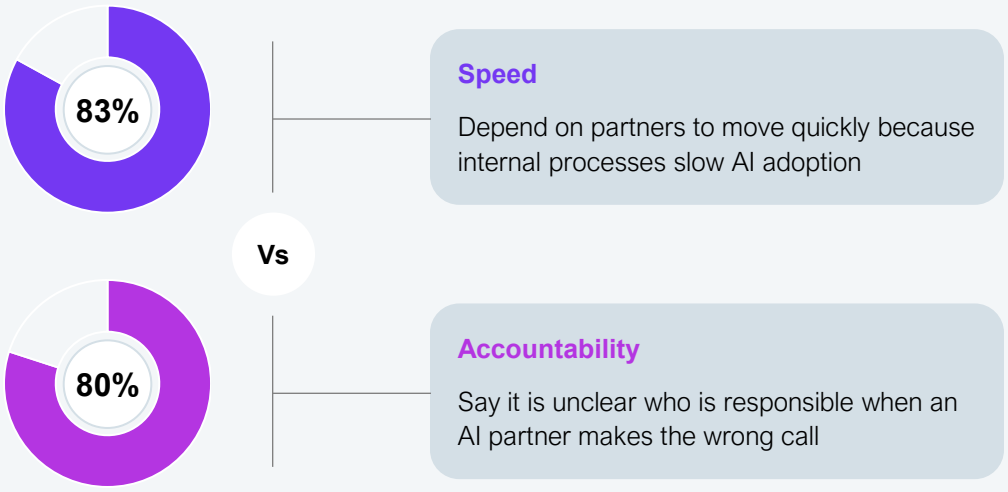

Sample: 505 executives across Global 2000 enterprises
Source: HFS Research, 2026

75% say teams defer to partners because they lack confidence, and the confidence gap has been outsourced too

Seventy-five percent (75%) of organizations say their teams defer to partners because they lack the confidence to challenge AI recommendations. That number is not a partner problem. It is a direct extension of the confidence gap inside the organization. The same fear, underinvestment, and imposter syndrome that prevent employees from challenging internal AI outputs also prevent them from pushing back on partner decisions.

Enterprises are not just buying speed from partners. They are buying the authority they failed to build internally. The partner fills the confidence vacuum the organization created. That is not a sourcing strategy. It is a governance failure that has been contracted out.

Exhibit 19: The confidence gap does not stay inside the organization; it travels

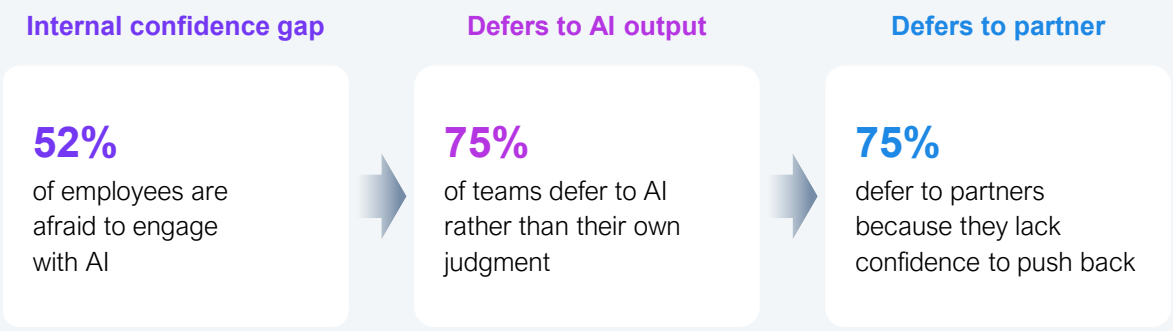

Sample: 505 executives across Global 2000 enterprises
Source: HFS Research, 2026

Enterprises are expanding partner dependency, not reducing it

Over the next 12 months, enterprises plan to invest more heavily in platform providers (36%), expand their partner ecosystem to fill capability gaps (34%), and replace traditional services partners with AI-native firms (29%).

The direction is toward more dependency, not less. That is not inherently wrong. But every new partner relationship that goes live without documented decision rights, override criteria, and consequence ownership is another extension of the authority vacuum beyond the boundary of the organization.

Humans at the helm cannot stop at the front door

Humans at the helm is not just an internal design challenge. It requires the enterprise to hold its partners to the same accountability standards it is trying to build internally—and to recognize that the governance failures described in this report do not remain within the organization.

Exhibit 20: Enterprises plan to expand partner ecosystem—only 3% reducing spend

Over the next 12 months, how is your organization planning to adjust its use of AI partners and service providers?

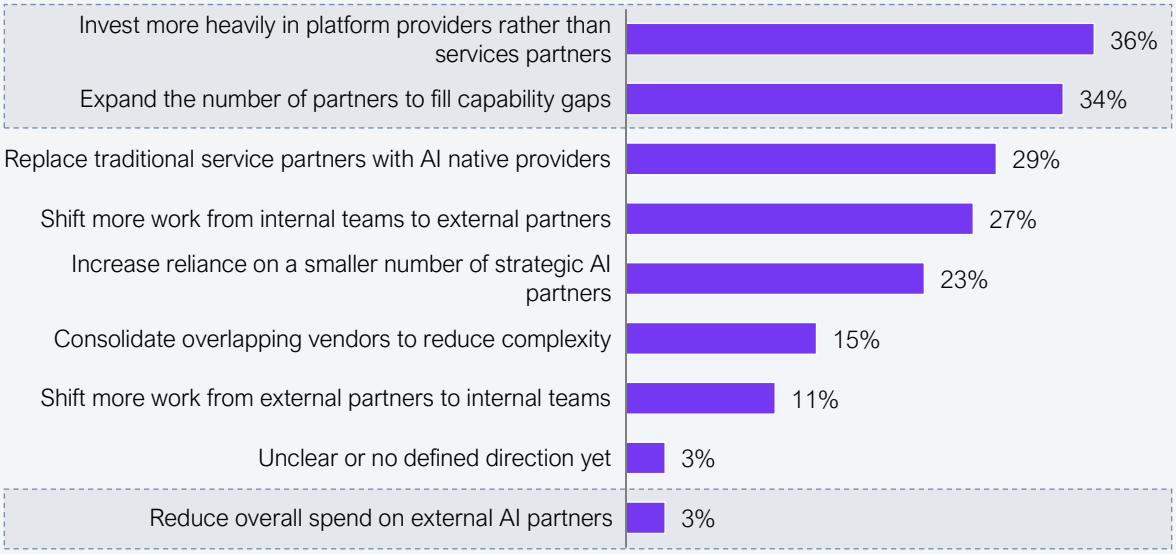

Sample: 505 executives across Global 2000 enterprises
Source: HFS Research, 2026

The bridge forward

You cannot be at the helm if the engine is being steered by someone you have no accountability agreement with.

SECTION 07

Putting humans at the helm

Leading AI enterprises put the human foundations in place before autonomy scales.

The velocity gap does not close on its own. The organizations closing it are not waiting for better models or clearer regulation. They are making deliberate choices about how authority is designed, how governance is resourced, and how people are prepared to engage with systems that increasingly act on their behalf.

What they are discovering is that putting humans at the helm does not slow AI down. It enables AI to move with more confidence, more trust, and more speed because the decisions behind it are owned by someone who can be held accountable for them. (See Exhibit 21)

Exhibit 21: The gap closes from the bottom up—six layers, each the foundation for the one above it

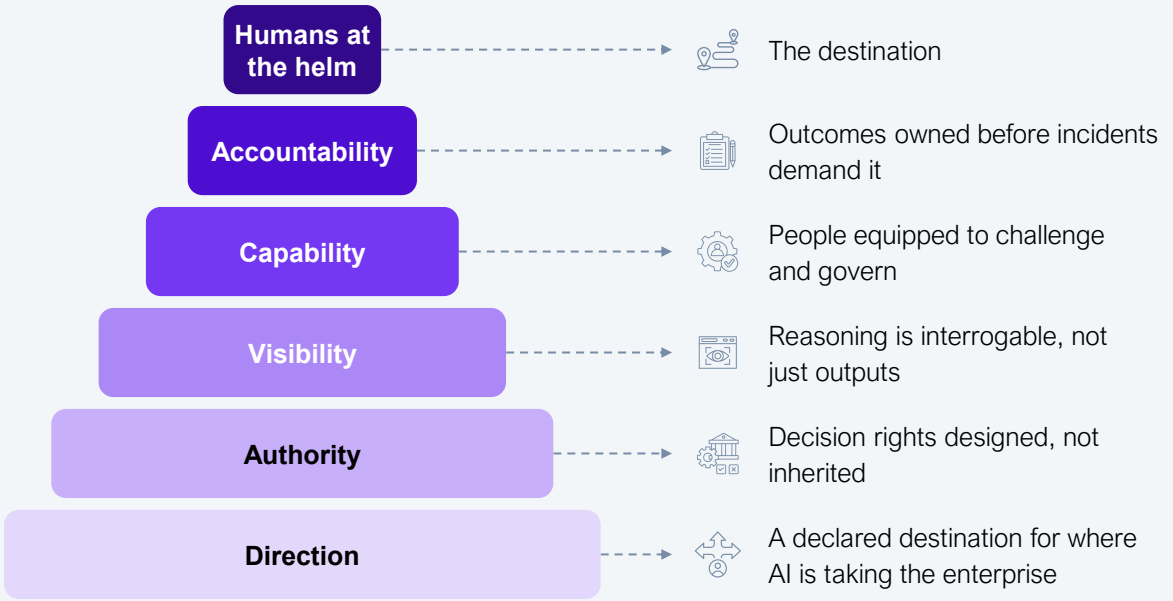

Sample: 505 executives across Global 2000 enterprises
Source: HFS Research, 2026

Direction: Declare where AI is taking the enterprise

Most organizations are letting AI evolve into something rather than deciding what it should become. The organizations that have closed the velocity gap started not with deployment, but with a declared destination—a clear answer to what AI is meant to help the enterprise become.

For employers

Name the destination

Define what AI is meant to change, not just where it can be applied. Assign a business-level owner for whether the enterprise gets there. Architecture will follow that decision whether you make it explicitly or not.

For employers

Design the transition

Tell people what is changing and what they are being asked to become before roles shift, not after. The workforce that feels led through the transition is the one that will show up as a partner in what comes next.

Direction is a leadership imperative. When it is declared clearly, it gives everyone inside the organization the foundation to exercise judgment, challenge AI outputs, and own outcomes with confidence. Without it, none of the layers above can fully hold.

Authority: Design decision rights before the system inherits them

Where authority is not defined, it defaults to the model or to whoever happens to be in the room. Documented decision rights are not bureaucracy. They are the condition that makes fast, confident action possible.

For employers

Define the boundary

Specify what AI can decide independently, what requires human review, and what triggers escalation. Assign ownership for outcomes, not just execution. Name that owner before go-live.

For employees

Know and hold your boundary

Understand when you are expected to rely on AI and when you are expected to challenge it. If that line is unclear, surface it. Undefined authority is where risk accumulates and where the velocity gap widens.

Visibility: Make AI decisions explainable enough to govern

Visibility is what converts human review from a checkpoint into a genuine control. The organizations with the strongest human-AI operating models have made reasoning accessible at the point of decision—not as an afterthought, but as the infrastructure of trust.

For employers

Invest in interpretability

Ensure that people reviewing AI decisions can understand not just what the system produced but how it arrived there. A review process that cannot access reasoning provides a record. It does not provide governance.

For employees

Question the opaque

Do not approve what you cannot explain. If you cannot trace the reasoning, treat it as a signal to question, not to defer. That instinct, exercised consistently, is what meaningful oversight looks like in practice.

Capability: Build judgment, not just adoption

Most enterprises are training people to use AI. Very few are training them to challenge it. The organizations outpacing their peers have invested in the human capacity to govern AI, not just operate it—making experimentation safe, rewarding challenge, and treating judgment as a capability to be built.

For employers

Make challenge visible and safe

Reward the behavior that builds governance capacity—experimentation, pushback, expressed uncertainty. A workforce that feels safe challenging AI is the one that will develop the judgment to govern it.

For employees

Engage before you feel ready

Confidence in AI governance develops through use, not before it. Ask why the model produced what it did. Push back when something feels wrong. The judgment governance requires is built one decision at a time.

Accountability: Own outcomes before incidents demand it

Accountability that appears after failure is not accountability. It is escalation. The organizations that have closed the velocity gap have made one consequential shift: they named who owns each AI-driven outcome before the system went live, and they extended that ownership across every partner and platform operating on their behalf.

For employers

Assign ownership before go-live

Name a consequence owner for every AI deployment. Extend that accountability across partners, platforms, and workflows. A governance failure does not stay inside the enterprise; it travels with every system and every partner it depends on.

For employees

Recognize what you carry

Reviewing or approving AI outputs carries responsibility. Ownership does not disappear because a system made the recommendation. That is not a liability to manage around. It is the design of meaningful human oversight.

The Bottom Line: The AI decade will not be defined by who built the most capable models. It will be defined by who built the most capable humans to direct them.

Humans at the helm is not a constraint on AI. It is the design that makes AI authority governable, defensible, and compounding in the right direction. The enterprises that will lead in the AI decade are not the ones that moved fastest. They are the ones that moved with the clearest answer to the question that most are still deferring: what does the human at the helm actually have the authority, visibility, confidence, and accountability to do?

HFS Research authors

Dana Daher
Executive Research Leader

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

Phil Fersht
CEO and Chief Analyst

Phil Fersht is widely recognized as the world's leading analyst focused on reinventing business operations to exploit AI innovations and the globalization of talent. He recently coined the term "Services-as-Software" to describe the future of professional services, where people-based work is blurring with technology.

Hridika Biswas
Associate Practice Leader

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

Mayank Madhur
Practice Leader

Mayank Madhur is a Practice Leader at HFS Research, driving deep research and insights into the healthcare and life sciences verticals. He also brings horizontal depth in IoT, digital engineering, and sustainability, collaborating with industry and technology leaders to deliver cross-functional insights. He holds the Sustainability and Climate Risk (SCR) certification from GARP and is a certified Project Management Professional (PMP®).

Demographics

Role

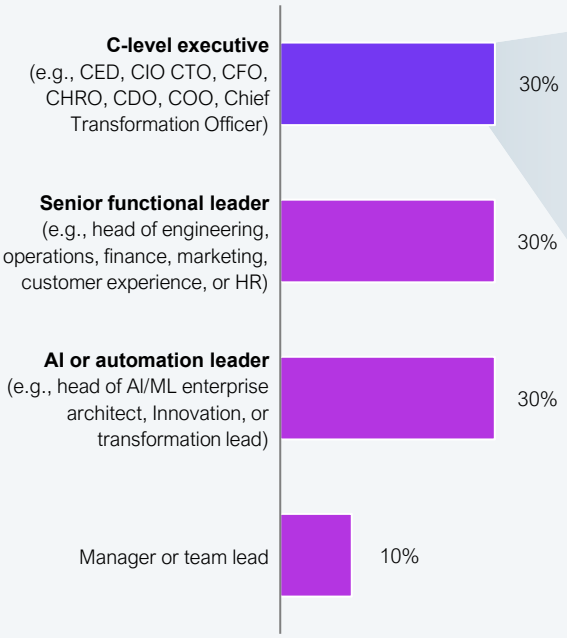

C-Suite Title

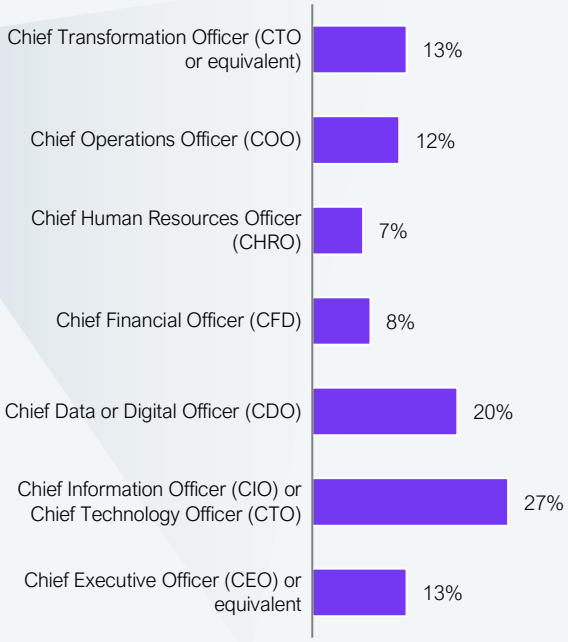

Industry

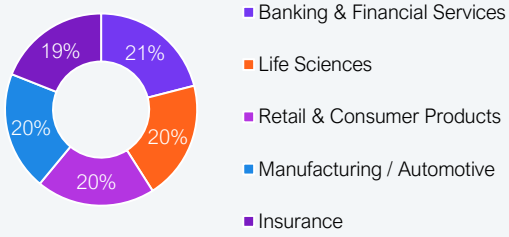

Involvement in AI or Automation Initiatives

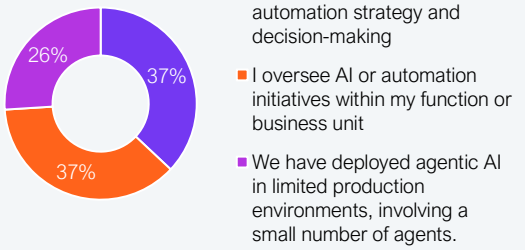

Revenue

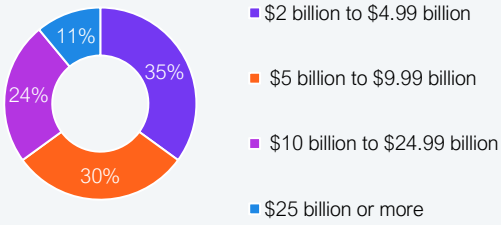

Country

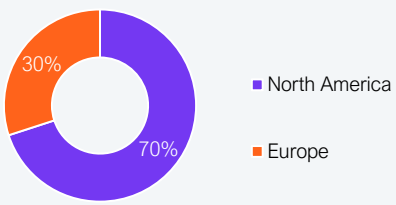

Source: HFS Research, 2026

About Altimetrik

Altimetrik is an AI engineering company, building the systems that power the modern enterprise.

With deep expertise across industries including BFSI, manufacturing, retail & CPG, automotive, healthcare, and life sciences, we help organizations modernize technology, unlock new revenue streams, and build sustainable competitive advantage.

Powered by a global team of 10,000+ practitioners and a foundation of engineering excellence, Altimetrik delivers AI that is not just implemented—but operationalized, governed, and built to continuously evolve.

Recognized by Constellation Research and Everest Group for AI and digital engineering leadership, and named among Glassdoor's Best Led Companies, Altimetrik is redefining how enterprises build for the AI era.

Learn more at altimetrik.com.

HFS

About HFS

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

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

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


www.hfsresearch.com


hfsresearch


www.horsesforsources.com


www.horsesmouthpodcast.com