

People & Organizational Performance Practice

How to close the agentic adoption gap

At their core, AI transformations are a reinvention of how work gets done. This requires change leadership, not just change management.

This article is a collaborative effort by Arne Gast, Ben Fletcher, Karim Thomas, Natacha Catalino, and Thibaut Larrat, with Rahul Varma, representing views from McKinsey's People & Organizational Performance Practice.



Successful AI transformations follow the 1:3:5 pattern: For every dollar invested in agentic technology, organizations spend three on process redesign and five on capability building and adoption. However, most companies invert this formula entirely. Technology captures the lion's share of executive attention and investment while capability building and behavior change are treated as implementation details. But these are not supporting activities—they are the real work.

In [McKinsey's latest State of Organizations research](#), thousands of global business leaders identified change management and siloed ways of working as major obstacles to scaling AI, outranking concerns about adequate technology infrastructure.

Traditional change management isn't enough to spur new ways of working that lead to value from agentic adoption. Periodic communications, one-time training, static guidance documents, and top-down rollout plans can create awareness, but they rarely build the trust, habits, and operating discipline required to scale agentic ways of working. At the same time, limited AI adoption constrains value creation, while ineffective adoption can generate significant token, compute, governance, and oversight costs without commensurate improvements in performance.

To close the adoption gap, organizations can create a more dynamic approach to change that redefines how work and judgment are distributed between humans and machines. In the agentic era, change leadership should be continuous, behavior-led, and embedded in the work itself.

The leaders that succeed will treat agentic AI as a reinvention of work by deliberately moving employees from awareness to belief, from belief to commitment, from commitment to developing capabilities, and from there to enforcing new behaviors through a new operating system. In this article, we discuss how organizations can focus on these five areas to make sure their people are central to creating value from AI.

Why agentic transformation feels different

The human dynamics that accompany agentic AI are more complex than those of previous technology transformations. Learning a new workflow tool or system of record may be frustrating. But an AI agent that drafts, analyzes, recommends, negotiates, or resolves problems touches something deeper in people: a sense of their own competence, judgment, identity, and future relevance.

Four fears sit at the heart of agentic resistance, as detailed below. Leaders who treat them as obstacles to push through will likely fail. Leaders who name and design for them will win.

- ***Fear of “being found out.”*** Most people don't know where to start with agentic AI. That alone would be fine, except they are also afraid to admit it. Asking a basic question, using an agent in front of a colleague, or experimenting openly can feel like advertising that you're behind. So people stay quiet. They wait to see what others do before committing themselves. Learning happens in private or not at all. The result is not refusal; rather, it is paralysis dressed up as busyness. Teams look engaged, but nothing actually changes.

- ***Fear of accountability without control.*** Agentic AI takes on real tasks: drafting, recommending, analyzing, and sometimes deciding. But employees know that when something goes wrong, accountability lands on them: “If the agent gives me the wrong answer and I act on it, it’s still my mistake.” Because that fear is rational, people are reluctant to rely on tools they don’t fully understand, particularly in work where accuracy, professional judgment, and reputation are on the line. Until employees clearly understand whether agents are reliable and how human review works, trust will be fragile.
- ***Fear of stepping into the unknown.*** People navigate change best when they can see where they’re heading. Agentic transformation offers no such comfort. The tools change every few weeks, and the final operating model can’t be defined in advance, because the technology and the work are co-evolving. Leaders ask people to move forward but can’t say where they will end up. This is more than a communications gap—it is genuine uncertainty. And it’s more psychologically challenging for people than any previous technological change.
- ***Fear of losing your professional edge.*** This may be the deepest and hardest fear to acknowledge. It is not simply about job security, though that is real. The more unsettling questions are: “If an agent now handles the analysis, drafting, synthesis, and judgment calls I spent years building, what exactly is my contribution? What is my professional identity now?” This fear shapes behavior quietly and powerfully, turning experimentation into a private hedging activity.

Together, these four fears explain why organizations so rarely get genuine adoption. What they get instead is imitation: the workflow that looks transformed on paper but hasn’t changed in practice, the employee who says yes but doesn’t alter any work, the tool that is technically deployed but barely used. Leaders who treat this as resistance they must manage are misreading the signal.


It is worth noting that these fears coexist, in many people, with genuine excitement about what agentic AI makes possible. The task for leaders is not to eliminate fear—they cannot—but to ensure that it doesn’t freeze behavior. The organizations that move fastest are those that can hold both the anxiety and the possibility, and design for both.

Middle managers reflect the most dramatic version of this challenge. Unlike senior leaders, who set the vision, or frontline employees, who experience the tools directly, middle managers must translate vision into reality while running a dozen other change programs and maintaining credibility in across levels. They carry the burden without having time to experiment themselves—and are often the people most fatigued by change, as well as the ones most anxious about AI-related changes (Exhibit 1).

Exhibit 1

Employees across levels, especially middle managers, report anxiety about AI-related changes in their jobs.

Share of respondents who feel anxious about AI-related changes at work,¹ %



¹ Respondents were asked to indicate their level of agreement or disagreement with the following statement: "I feel anxious about AI-related changes at work." Percentages shown are the share of respondents with a score ≥ 4 on a 5-point agreement scale.
Source: McKinsey AI Individual & Organizational Readiness Assessment Panel Survey, Feb–Apr 2026, analyzing 750 English-speaking employees

McKinsey & Company

From change management to change leadership

Most organizations approach agentic transformation with the tools of traditional change management: communications campaigns, training calendars, KPI dashboards, and project plans. These were designed for a different kind of change, where the destination is known, the answer is largely determined, and leaders' role is execution.

Agentic transformation doesn't fit that model. It belongs to the category of change that [McKinsey calls "C4 reinvention"](#): not moving from project A to B but reimagining an entire system from A to unknown (Exhibit 2). The CEO can't credibly say, "I have the answer." The honest statement is, "I don't know exactly where we're going, but I'm learning alongside you—and the organization that learns faster will win."

That reframing from change to learning is the strategic differentiator. When employees believe the organization is genuinely learning its way forward, resistance softens. Experiments become evidence, and mistakes become curriculum. The question shifts from whether they're doing things right to whether they're learning faster than the competition is.


C4 leadership requires skills, such as modeling uncertainty rather than projecting certainty, that run counter to most executive instincts. Leaders who admit they are still figuring things out accelerate adoption more than those who issue confident strategy narratives. In C3 change, leaders know the answer and manage people through resistance. C4 leaders don't know the answer and lead people through learning.

Organizations that treat agentic AI as a project with a defined end state consistently underinvest in the human system and discover too late that technology adoption and value capture have diverged. Instead, the goal should shift from compliance to capability—not whether employees completed training but whether they are learning to redesign their work faster than last quarter.

Exhibit 2

Agentic transformation is the next frontier of change.

Each change level must fully incorporate and build on prior levels



McKinsey & Company

The playbook for agentic change, from A to E

The organizational playbook for agentic adoption contains five main chapters—awareness, belief, commit, develop, and enforce. Each move builds on the last: people must first become aware before they can believe; they must believe before they are ready to commit; they must commit before they can develop capabilities; and they must develop these capabilities before the new operating system is fully embedded.

Leaders should understand that momentum can stall or even reverse. When trust breaks down—from a poor decision, a policy reversal, or a model failure—employees can backslide. Leaders need to monitor where people and teams are in the agentic adoption journey and intervene at the right moment with the appropriate solutions (see sidebar, “Step into the shoes of your people”). Once they understand the complex human dynamics that accompany adoption, leaders can focus on five areas that create clarity and trust.

Awareness: Make the change tangible

Leaders take the first step in building conviction by providing opportunities for people to dive in and create direct experience with the technology. This tinkering changes the narrative they tell themselves and their colleagues. People who have had concrete experience with AI are more likely to become credible advocates, describing with specificity what they have seen change, what humans continue to own, and how decisions are made.

Step into the shoes of your people

Leaders should understand that agentic transformation is not experienced uniformly. A category manager, a finance controller, a compliance officer, a software engineer, and a customer service manager face different anxieties, workflow constraints, capability gaps, and definitions of what constitutes a good outcome. Generic change programs designed for “the workforce” are usually too abstract to change real behavior.

Leading organizations invest early in persona-based change journeys. They map what a key role’s week looks like today and what it should look like tomorrow—not at the level of strategy,

but at the level of decisions. Where will an agent enter the workflow? What is the critical moment when the person either acts on the agent’s input or doesn’t? What would build enough confidence to act?

In one European retail transformation, the team analyzed a category manager’s workweek. Today, 50 to 60 percent of the manager’s time is spent on category strategy and planning, 15 to 20 percent on vendor management, and the rest on performance analysis and ad hoc requests. Tomorrow, agents will handle data collection, report preparation, and exception chasing, while the manager focuses on strategy, supplier relationships,

and decisions requiring judgment. The judgment moments—specific decision points where agents provide input and humans make the call—were scripted in advance, rehearsed in pilots, and refined from feedback. That specificity is what makes the change real rather than theoretical.

This persona work surfaces what diagnostics often miss, which is whether a workflow itself has been redesigned to create clear, repeatable moments where agents add value and human judgment stays visibly in the loop.

A European consumer goods company built belief and capability among its top 100 leaders through an intensive three-day program. Leaders spent the first half day demystifying the technology—understanding what agentic AI is, what is new, and what can go wrong with these systems. They also discussed technology coming out of China, the European Union, and the United States.

The next half day featured coaching on tools, with leaders experimenting with coding, text, sound, and video applications while embracing a “beginner’s mindset.” The third day was devoted to working on real business problems, coming up with actions to address them, and debriefing about what they had learned.

Broad statements about becoming AI-first or unlocking productivity can increase anxiety, making employees feel incompetent and leaving them to fill in the blanks. Instead, leaders who are proficient with AI tools can describe with specificity what will change, what humans will continue to own, and how decisions will be reviewed. This change story should address the questions employees are asking: What will happen to my role? Will the organization keep people and shift their work, or will it reduce head count? Transparency about these questions, even when answers are incomplete, builds more credibility than messaging designed to avoid them.

Belief: Leaders adopt first

While awareness creates contemplation, belief moves people into the trial stage. Two groups are crucial to driving this belief: leaders who model the new behavior and trusted peers who demonstrate that the technology is safe and worthwhile.

[A recent McKinsey survey of employees and leaders on AI readiness](#) found that respondents who report low trust in their organizations’ support during an AI transformation are 1.5 times more likely than those reporting high levels of trust to feel anxious about workplace changes related to AI. In a fast-moving transformation, trust is built when leaders credibly show how the organization will learn its way forward.


McKinsey research also shows that high performers' use of AI is more often championed by their leaders. High performers are three times more likely than their peers to strongly agree that senior leaders at their organizations demonstrate ownership of and commitment to their AI initiatives (Exhibit 3).

Leadership role modeling in the agentic era should be business focused and applied to real work—not “happy AI” anecdotes, but workflow examples from leaders' own practice. This means using AI-generated insights in business reviews, asking agents to stress-test assumptions, replacing static reporting with agent-supported decision flows, and publicly noting when human judgment overrode an agent recommendation—and explaining why.

At one enterprise software provider, adoption stalled when agentic AI was treated as an engineering-owned tooling rollout. The company shifted course by putting leaders through immersive bootcamps that prepared them to have practical conversations about workflow change, uncertainty, and evolving expectations.

Exhibit 3

High performers tend to have senior leaders who demonstrate strong ownership and commitment to AI initiatives.



Note: Figures may not sum to 100%, because of rounding.

¹Question asked to what extent respondent agreed that senior leaders in their organization demonstrate true ownership of and commitment to its AI initiatives (eg, championing them across the organization over time, role modeling, providing continued funding and engagement in regular budget reprioritization).

²AI high performers are respondents who said more than 5% of their organization's EBIT and "significant value" are attributable to the organization's use of AI.
Source: McKinsey Global Survey on the state of AI, 1,993 participants at all levels of the organization, June 25–July 29, 2025

Peer networks carry equal or greater weight than leaders do. A global bank's mobilization of more than 2,000 "AI champions" illustrates the power of credible, peer-led adoption at scale. Champions ran floor walks (structured visits to team workspaces to demonstrate tools in context) and hosted office hours and micro-demonstrations (short sessions that show colleagues a specific use case). They made experimentation feel normal, converted abstract use cases into practical ones, and served as a listening system, surfacing where adoption was advancing and where it was stalling.

The key distinction from traditional champion programs is that agentic champions need to be credible users, not just influential voices. Their authority comes from demonstrated practice—for instance, being the colleague who uses AI to cut promotional-calendar preparation from two days to two hours and can show people how.

Commit: Give people the space to choose

Once employees believe the change is real and worthwhile, they enter the zone of choice: reviewing their options, considering the stamina required, and deciding where they dare to try. This moment of personal commitment is the hinge of the whole model—and the one most often left to chance.

It is striking how often a top team spends multiple days in executive retreats and immersion programs exploring the technology, while frontline employees receive a ten-minute presentation and are told to start using it. That asymmetry signals, unintentionally, that the transformation belongs to leadership and is being imposed on everyone else.

The commitment moment requires protected time to make sense of what the change means personally, to process what is being lost as well as what is being gained, and to decide how extensively to engage. For some employees, this process includes genuine grief, as a role built over years is substantially redesigned by a technology they did not choose. Leaders who acknowledge this directly—who say, in some form, "I understand that this changes work you have spent years mastering, and I take that seriously"—are more likely to convert passive compliance into active commitment.

When adoption feels like a choice, it also tends to be durable. The employee who has committed—who has said, "I am going to try this"—is more likely to persist through difficulty and less likely to lose momentum.

Develop: Build competence in the work itself

With commitment in place, employees move into full action by adapting new behaviors, turning them into daily and weekly practices, and learning to do them better and faster each time. The goal of this stage is to help people redesign tasks, evaluate agent outputs, apply judgment, manage risk, and collaborate with hybrid human-agent teams in repeatable, improvable ways.

One-time training is poorly suited to creating work habits for a technology that changes every few weeks. Capability building must be continuous, practical, and tied to the flow of work. Some organizations are moving from classroom-based training to embedded coaching. For example, a global Fortune 500 technology company placed AI coaches directly into engineering teams, with one coach per five to ten employees. They joined sprint ceremonies and retrospectives, helped troubleshoot workflow issues, refined prompting and review practices, and built confidence in safe operational use. By embedding learning into the actual work, the company converted experimentation into habits faster than training alone could.


The reskilling narrative matters as much as the program design. A connected certification architecture—mapping capability levels from level one (aware) through level five (enforce)—gives employees a visible progression and makes development easier for managers to assess. Employees engage more readily when they can see what higher-value work looks like, and when the organization demonstrates, through role design and career paths, that it intends to invest in their career progression.

As [McKinsey Global Institute research shows](#), while up to 57 percent of current US work hours are automatable, more than 70 percent of human skills sought by employers today will endure. This shows that the future of work is a true skill partnership of people, agents, and robots (Exhibit 4).

Exhibit 4

People, agents, and robots could all play significant roles in the workforce of the future.

Distribution of work hours in the US, by technical automation potential, 2024, %



Note: Technical automation potential shown is the late scenario of expert estimates. The early scenario of technical automation potential in the US is 65% of current work hours. In this research, we use "agents" and "robots" as broad, practical terms to describe all machines that can automate nonphysical and physical work, respectively. Many different technologies perform these functions, some based on AI and others not, with the boundaries between them fluid and changing. Using the terms in this inclusive way lets us analyze how automation reshapes work overall.

¹All work requires cognitive capabilities. Both physical and nonphysical work may also require social and emotional capabilities.

Source: O*NET; US Bureau of Labor Statistics; US Census Bureau Current Population Survey; McKinsey Global Institute analysis

Enforce: Redesign the operating system

The final chapter of the agentic playbook is embedding new behavior and reimagining workflows. This happens only when the formal operating system reinforces new ways of working. If metrics, incentives, decision rights, promotion criteria, and management rituals continue to reward the old behavior, the new behavior won't survive regardless of how well the earlier stages went.

This is where most agentic transformations fall short. They measure log-ins, licenses, prompt volume, and agent invocations. What they don't measure is whether the work is better. A stronger approach links agentic adoption to process performance: cycle time, quality, customer experience, decision speed, risk reduction, and value captured per workflow. As usage matures, the measurement system should add economic discipline: value created per workflow, cost per task, human review burden, and quality improvement over time.

Salesforce provides an example of this outcome-oriented reinforcement. With Agentforce, the company tracks adoption against operational outcomes, including the share of support cases agents can resolve autonomously and the reduction in resolution time. This lesson extends well beyond this one company, however. Agentic AI becomes durable when organizations measure whether work is actually better, faster, safer, or more valuable—not simply whether people are using AI tools.

Leaders should note that when trust breaks down, employees can slide back from maintenance into the awareness or belief stages. Change leaders should monitor where people are on the curve (for instance, by measuring how and where users are embracing and rejecting recommendations for using AI tools) and treat loss of trust as a signal to return to the right stage, not to push harder.

The organizations moving fastest treat resistance not as an irrational obstruction but as data. They diagnose where fear, ambiguity, capability gaps, and process friction sit, and they design interventions that address those barriers directly.

The leadership mandate

Change leadership in the agentic era requires special focus from CEOs, CFOs, and chief human resources officers (CHROs), who each play distinct and complementary parts.

CEOs drive the voice of the customer end to end through workflows and across ecosystems. They set the C4 ambition: not executing against a fixed plan but learning how to reimagine work faster than the competition. Their most important signal, in addition to the strategy narrative, includes their own operating cadences, the questions they ask, the agents they use visibly, and the decisions they make that the whole organization can see.

CFOs use budget discipline to find new ways of getting work done in a leaner structure. This is what one leader calls “extreme budget pressure” as a catalyst for learning. CFOs allocate capital to future experiments as a portfolio, understanding that in agentic transformation, as in venture capital, a small number of bets generates the majority of the value. The discipline is not to defund experiments too early, but to double down on those that are teaching the organization something. With cost of tokens becoming P&L line items, the role of CFOs is to demand value creation.

Find more content like this on the
McKinsey Insights App



Scan • Download • Personalize



They supply the budget for tokens, restrict the backfilling of people, and demand that the total productivity of people plus tokens increases.

CHROs have perhaps the most pivotal role: making change personal and designing learning strategies that combine people, agents, and, in some contexts, robots. They rethink what to hire for (resilience, learning agility, judgment, the ability to connect with others) versus what to develop (the ability to reimagine workflows, decision quality, ethics in human–AI collaboration). They design the progression pathways that make the organization's investment in people visible and credible. Bringing in data to predict and follow fundamental workflow shifts and driving adoption for ROI will be [an exciting game changer for many HR departments](#), while they agentify their own workflows at the same time.

For all leaders, three imperatives remain constant: state the direction clearly while admitting uncertainty, reinvent the operating model to support it, and model the learning behavior they are asking of others.

The agentic era will reward organizations that can learn faster than their operating systems were originally designed to move. In addition to the technical transformation of tools, data, and architecture, leaders must manage the human transformation of mindsets, behaviors, roles, and systems. They can take employees on a journey from awareness to belief to commitment to capability and, finally, to a reinforced operating system. When these elements come together, employees don't experience AI as something being done to them, but as path to greater relevance, judgment, and impact. The transformation stops being something that happens to the organization. It starts being something the organization is leading.

[Arne Gast](#) is a senior partner in McKinsey's Amsterdam office, [Ben Fletcher](#) is a senior partner in the Sydney office, [Karim Thomas](#) is a senior partner in the Seattle office, [Natacha Catalino](#) is a senior expert in the Boston office, [Thibaut Larrat](#) is a partner in the Toronto office, and [Rahul Varma](#) is a senior manager of client capabilities in the Gurugram office.

The authors wish to thank Carmen James, Hayden Lindskog, Jasmine Majdpour, Nicolette Rainone, Rishika Bhikha, and Simon Wintels for their contributions to this article.

This article was edited by Barbara Tierney, a senior editor in the New York office.

Designed by McKinsey Global Publishing
Copyright © 2026 McKinsey & Company. All rights reserved.