


People & Organization Practice

# AI is changing work. Now it has to change the organization

The first phase of enterprise AI put powerful new tools in people's hands. Capturing the larger opportunity from agentic AI will require leaders to rethink the shape of the organization itself.

*by Alexis Krivkovich, Brooke Weddle, Holly Price, and Vik Sohoni*



**Agentic AI changes** three fundamental constraints that have shaped the modern organization: human capacity, expertise, and coordination. For more than a century, the familiar building blocks of organizational design, from tasks and jobs to teams, management layers, and hierarchies, have reflected those constraints. AI doesn't remove them entirely, but it alters them enough that leaders can ask a more consequential question: What can this company do that wasn't possible before?

Most organizations are not yet asking, let alone answering, that question. AI adoption is scaling rapidly, and individuals are already seeing meaningful productivity gains, but enterprise value has not kept pace. In McKinsey's [2026 state of AI survey](#), 80 percent of respondents say AI has improved their individual productivity, while only 37 percent of respondents report positive EBIT impact—and just 6 percent qualify as AI high performers. Tellingly, nearly three-quarters of those high performers report substantially redesigning workflows, compared with only one-quarter of other organizations. This is the gap between using AI and redesigning the organization around it.

[Recent McKinsey AI readiness research](#) reinforces that gap. Just 11 percent of organizations have reached the “reinvention” horizon, where they rethink how work gets done and how value is created. Reimagining means going beyond individual AI tools and workflow automation to redesign roles, workflows, and operating models around what AI now makes possible. As AI capabilities advance, the existing organization can become the constraint. Yet many transformation efforts address only parts of that system—workforce, organizational design, or change management—rather than considering how they need to change together.

In this article, we explore four interconnected shifts that can help leaders capture the larger opportunity from AI: organizing work around outcomes, reshaping the talent system to reflect the changing mix of human and AI work, building the mechanisms that allow change to compound, and putting clear authority and accountability behind the transformation.

## **Operating model: From org chart to orchestration chart**

For much of the modern corporation's history, companies have organized people into functions and hierarchies, with work passing from one part of the organization to another. AI enables a different logic: Start with the outcome, then assemble the humans and agents needed to deliver it.

There are signs that this shift has already begun. [New McKinsey research](#) surveying more than 700 executives and senior leaders finds that companies furthest along in AI reinvention are more likely to organize cross-functional teams around end-to-end products, customer journeys, and business processes, putting decisions and accountability closer to where value is created. There is no single winning operating model structure; rather, the common thread is organizing work end to end, reducing unnecessary handoffs, and moving decisions and resources faster when priorities change.

As agents take on meaningful parts of work, organizational design increasingly becomes a question of orchestration: which humans and agents should come together around an outcome, what each should do, and which handoffs and interfaces are still necessary. An “orchestration chart” can make those choices explicit. In some cases, agents may do more than simply join an existing team; they may replace the need for coordination between people or functions altogether, changing the network through which work moves. These shifts also make [risk mitigation](#) an integral part of orchestration—particularly when organizations are moving quickly—as decisions about agent autonomy, access, oversight, and escalation become embedded in the design of work itself.

No company offers a finished blueprint, but early examples provide signals of what these changes might look like. In one organization’s product development function, traditional pods of eight to ten people shifted to hybrid teams of four to six people supported by agents. An orchestration chart of the new team shows which work has moved to agents, which human roles now span more of the product development process, which handoffs have disappeared, and where human judgment and accountability matter most.

That same shift may change coordination mechanisms such as meetings, which are essentially an organization’s way of transferring information, creating alignment, and making decisions. In a human–agent organization, agents can increasingly track which decisions have been made, what work is still outstanding, and what needs attention next. Meetings don’t disappear, but their purpose shifts away from the routine transfer of information toward building judgment, commitment, and trust.

The practical test of a redesigned operating model is whether the organization can make decisions, move information, and coordinate work with fewer layers and handoffs and less manual translation between functions. Across our client work, we have seen “coordinator roles” (internal roles focused largely on moving decisions across the organization, such as project management) grow at roughly 1.5 to 2.0 times the rate of line roles over the past five years, signaling how complex organizations have become. AI creates a significant opportunity to reduce some of that complexity rather than add to it. But the flatter pyramid can enable value only if the lost coordination is replaced by better intelligence, clearer decision rights, and stronger human judgment.

Concrete markers of operating model progress:

- *Seek structural change, not just speed:* Work itself is being reorganized rather than simply accelerating the old model.
- *Remove unnecessary handoffs:* Teams own outcomes across functional lines, while agents take on more execution and coordination.
- *Change the shape:* Team size, composition, and management layers materially change as agents take on more work.
- *Map the new model:* Orchestration charts show how humans and agents combine around outcomes, who does what, and where human judgment and accountability are essential.

## Talent: How to build a system around work that keeps changing

Organizing work around outcomes rather than inherited roles and functions has consequences for the entire talent system. Most companies still hire, develop, deploy, and advance people through relatively stable jobs and career paths. A human–agent organization needs a talent system that can move people as the skills required by the work change. To make that shift, organizations can focus on rebuilding the pathways through which people learn, advance, and develop expertise.


AI changes the relationship people have with their work. For decades, the job has been the basic unit through which companies have organized work and developed talent. It defined what work someone did, how they created value, how they learned, how they took pride in their output, and how they advanced. A junior analyst learned by doing analysis. A designer improved by making more important design choices. A manager progressed by taking responsibility for larger teams and more complex work. The job has traditionally linked the work people do with how they create value, build expertise, find satisfaction, and advance.

AI alters those connections by completing parts of the work without necessarily replacing the job itself. Agents can take on tasks that once gave a role much of its substance: research, drafting, coding, analysis, monitoring, summarization, routing, and coordination. That creates a productivity dividend that allows time and capacity to be used differently (Exhibit 1).

Exhibit 1

### Where does the AI productivity dividend go?

Customer service representative's role, % of current role capacity, illustrative



While individual productivity gains are welcome, they don't automatically translate into enterprise value. Capturing more value requires an explicit choice about where those gains will go. In a recent conversation with McKinsey, [Stanford economist Erik Brynjolfsson argues](#) that focusing primarily on workforce reduction misses the larger opportunity to reconfigure and rebundle work around what people can now do with AI.

Agentic capabilities also change how people experience the work. If the system writes the first draft, creates the image, summarizes the research, prepares the analysis, or recommends the decision, the human may feel less like the producer and more like the reviewer of someone else's output. In creative fields, that concern is already visible among artists, writers, designers, and actors confronting AI-generated work. Similar questions will appear inside companies as AI moves deeper into professional, managerial, and technical roles. As roles combine human judgment with a changing mix of AI-enabled activities, [companies will also need to manage cognitive load](#): Too little human involvement can lead people to become overly reliant on AI, while too much can leave them overloaded by an expanding mix of responsibilities.

AI also changes career progression. In most organizations, advancement is tied to managing larger teams, coordinating more work, and sitting higher in the hierarchy. If AI absorbs more status tracking, synthesis, routing, and coordination, some of the work that justified management layers will shrink, and [middle management](#) may become less of a default career destination. Companies will need career paths that reward craft mastery, judgment, human-agent orchestration, customer or client trust, and the ability to make high-quality decisions with AI support. A senior individual contributor who can lead agents around a complex business outcome may create as much value as a traditional manager with a large team.

While deep specialization will remain important, there is a trade-off. If agents take on the entry-level work through which people have historically built expertise, today's productivity gain can become tomorrow's judgment gap. [Recent research showing employment pressure among early-career workers](#) in highly AI-exposed occupations makes the question of how apprenticeship is handled increasingly concrete. Companies will need to redesign how people build expertise as the underlying work changes, with managers spending less time coordinating activity and more time creating opportunities for practice, feedback, and increasingly consequential judgment.

Forecasting which jobs will grow or shrink is therefore not enough. [New McKinsey Global Institute research finds](#) that people's ability to move into new work depends heavily on skill adjacency: The most viable moves tend to be into jobs that draw on skills people already have. Some occupations offer many adjacent pathways, while others have fewer options or face credential barriers. For employers, the most salient questions revolve around which skills future workflows will require, who has a realistic path into that work, and where the organization will need to retrain, hire, or create new pathways.

Some organizations are beginning to connect those questions directly to internal mobility. One technology company uses an AI-powered talent marketplace to map employees' existing skills to personalized learning, stretch assignments, and open roles. In a recent quarter, more than half of its hires came from inside the company. That approach connects skills data to actual opportunities rather than treating reskilling as training disconnected from where people can move next.

Companies are also setting principles for how employees will experience the transition rather than going case by case. For example, one global life sciences company has stated that AI may support work but should not replace human creativity, expertise, autonomy, agency, or oversight. Employees remain accountable for AI-assisted work. Such principles don't resolve every workforce decision, but they do establish boundaries before individual decisions about roles and automation arise.

Concrete markers of talent system progress:

- *Make the productivity dividend explicit:* Show where AI-released capacity is going, whether to greater output, different work, more sustainable workloads, lower cost, or a combination thereof.
- *Map realistic paths into future work:* Connect changing work to the skills it requires, identify who has adjacent skills, and where reskilling alone will not be enough.
- *Build new paths to expertise:* Where AI absorbs foundational work, replace it with deliberate practice, feedback, and exposure to increasingly consequential decisions.
- *Protect what should remain human:* Be explicit about where human judgment, autonomy, expertise, and oversight remain as roles change.

## Turning AI activity into sustained organizational change


Redesigning workflows and talent systems leads to enterprise value only if people themselves begin working differently. Yet one of the paradoxes of the current AI moment is that widespread adoption can coexist with surprisingly little organizational change. Employees may use AI every day while the same workflows, decision rights, incentives, and management practices remain intact. The challenge is to make new ways of working take hold across the organization.

[McKinsey's recent AI readiness research points to this divide](#): 70 percent of respondents say they are personally ready for AI, while only 27 percent of leaders say their organizations are ready to make the shifts required for an agentic future (Exhibit 2).

Exhibit 2


**Most respondents are ready to work with AI, while organizations overall are much less prepared.**

Personal readiness, % of all respondents



of respondents across job levels say they are personally ready for AI

Organizational readiness, % of leader respondents



of leaders say their organizations are ready to make the shifts needed for an agentic future

Note: Percentages shown are the share of respondents with a mean score  $\geq 4$  on separate composites aggregating across survey questions for personal readiness and organizational readiness; all questions were asked on a 5-point agreement scale.  
Source: McKinsey AI Individual & Organizational Readiness Assessment Panel Survey, Feb–Apr 2026, analyzing 750 English-speaking employees; organizational-readiness percentage is based on a subsample of 608 leaders

McKinsey & Company

That preparedness gap [puts trust at the center of the transition](#). People can embrace AI themselves while questioning whether their organization will manage the consequences well. McKinsey data suggest that this is more than an employee sentiment issue: Trust that the


organization will support people as AI changes their work is associated with greater enterprise value across every stage of AI maturity (Exhibit 3). As roles, career paths, and expectations change, trust will increasingly be shaped by what employees see leaders do, including how openly workforce decisions are explained, whether people have realistic opportunities to adapt, and whether commitments about human oversight and agency survive difficult choices.

Exhibit 3

## Trust in the organization is critical to whether employees are ready for the changes AI brings.

The specific readiness factors that predict enterprise value<sup>1</sup>

● Organizational-readiness factor ● Personal-readiness factor



<sup>1</sup>Statistically significant positive predictors of enterprise value impact for respondents in each horizon based on separate regression models for organizational-readiness and personal-readiness concepts; each horizon builds on the previous one, meaning all the challenges in prior horizon(s) apply; trust in the organization significantly predicts enterprise value impact across all three horizons. Source: McKinsey AI Individual & Organizational Readiness Assessment Panel Survey, Feb–Apr 2026, analyzing 750 English-speaking employees. Results are based on subsample of 582 leaders who could be categorized into one of the three horizons (enablement, n = 267; automation, n = 252; reinvention, n = 63).

Sustained change also requires people to work differently, not just learn new skills. Doing all the work oneself is not necessarily a sign of competence if AI can improve the outcome. Teams can test ideas faster and more cheaply with AI while being clear about where mistakes would carry serious consequences. And contribution becomes less about completing a list of activities and more about owning the outcome. McKinsey's readiness research finds that such organization-wide behavior shifts are associated with enterprise value for companies in the reinvention horizon.

Some organizations are beginning to build those behaviors into how work is learned and reinforced. One large financial institution has evolved its approach from broad AI training toward role-based learning, peer coaching, dedicated development for managers, and repeated opportunities to build and demonstrate solutions in real work. Internal champions, hackathons, weekly demos, and peer learning circles help successful practices spread across teams, and the number of AI solutions reaching production has continued to grow. The model embeds learning and experimentation into how teams work, rather than treating AI adoption as a separate training effort.

Concrete markers for leading sustainable change:

- *Look beyond adoption:* Workflows change as teams learn what AI makes possible, rather than simply adding AI to existing work.
- *Reinforce new behaviors:* Performance management, manager expectations, and incentives change to support how people need to work differently.
- *Build trust through the transition:* Employees trust the organization to support them as AI changes their work, a factor associated with enterprise value across every stage of AI maturity.
- *Get faster at changing:* The organization shortens the time it takes to turn new AI capabilities into new ways of working.

## **Enabling AI reinvention: Building the transformation engine**

Reinvention on this scale can't be carried out through disconnected AI experiments. Companies need a way to repeatedly redesign consequential work, build the technology and controls around it, and change the organization alongside it.

That starts with clear ownership. Because AI reinvention cuts across business, technology, talent, risk, and other organizational boundaries, no single existing function has a natural mandate to deliver it end to end. Companies need a senior leader with the authority to drive the agenda across those boundaries and make the trade-offs that reinvention requires. The title matters less than the levers attached to the role: the ability to move funding and scarce talent,

shape strategic investment decisions, and resolve competing priorities.

Some companies are beginning to build dedicated execution capabilities for this type of work. One emerging model is what we call an “agentic mission factory.” Its job is to take a consequential business outcome—the mission—and redesign the work through a hybrid human–agent workflow, build the required agents and controls, launch the new way of working, and create reusable assets for the next mission. The accountable leader provides the direction and authority while the mission factory turns that mandate into working solutions. The factory also creates the AI operating system, which compounds institutional intelligence by creating learning loops as individual agents execute their tasks.

In this model, the factory starts with the mission, not the tools. A mission might be to approve and fund a good loan faster, resolve a customer issue in one interaction, reduce claims leakage, accelerate product development, or improve supply chain resilience. The business owner brings the outcome, economics, constraints, and accountability. The factory brings the cross-functional capability to redesign and build the work.

That capability needs more than AI engineers. It brings together business, technology, product, design, data, risk, compliance, cybersecurity, legal, change, and HR from the start. Controls are built into the workflow rather than added at the end, and product ownership is assigned at launch, not after the pilot. Evaluation, monitoring, permissions, autonomy, and life cycle management are part of the operating model.

This model builds on the “[digital factory](#),” which used dedicated cross-functional teams to change the business while the rest of the organization continued to run it. The agentic version extends that model to redesigning workflows around humans and agents. That requires new capabilities for building and governing agents, as well as the shared AI infrastructure, evaluation, permissions, and life cycle management needed to operate them at scale.

The mission factory is also where the organizational changes described above come together. Redesigning a mission means deciding how humans and agents will work together, how people’s roles and development need to change, and how new ways of working will take hold across the organization. Those choices need to be made as the work is redesigned, not after the technology is built. That is what turns a series of AI deployments into organizational reinvention.

Concrete markers of AI transformation progress:

- *Put authority behind the mandate:* The accountable leader can move funding and scarce talent and resolve cross-enterprise trade-offs, not simply recommend priorities.
- *Build together from the start:* Business owners, builders, and control functions work on the mission together, with ownership and accountability established before launch.

- *Build for production, not the pilot:* Evaluation, monitoring, permissions, and life cycle ownership are designed in before agents go live.
- *Make each mission compound:* Completed missions leave behind reusable agents, components, controls, and decisions that shorten the path from redesign to production for the next one.

## The leadership test

No single company can claim to be a finished human–agent organization, but the direction is becoming clearer. AI expands the choices leaders can make about how teams are shaped, how hierarchies work, how jobs develop people, and how quickly the organization can change.

Leaders can test whether their organization is truly ready for AI transformation by asking six questions:

- Where would a different mix of humans, agents, and data create more value, and how should roles, accountability, and controls shift as a result?
- As AI takes on more tasks, which parts of each role remain distinctively human, meaningful, and worth designing the job around?
- Where is AI removing the work through which people have built expertise, and what new apprenticeship ladder will replace it?
- How will people be developed, coached, and promoted to cultivate judgment in human–agent workflows?
- What new structure or capability will redesign workflows, build agents and controls, and change the organization as one coordinated effort?
- What behaviors are we modeling as leaders, and would our teams conclude from watching us work that AI is genuinely changing how we lead?

If leaders can answer these questions, AI is no longer just an experiment. It is changing the organization's ability to break the constraints of the pre-AI age.

---

The companies that capture the most value from AI will not necessarily be those that deploy the most agents or automate the greatest number of tasks. They will be the ones that [use AI to rethink how value is created](#), then make the organizational choices required to deliver it, from how work and talent are organized to how transformation is led and built. The larger opportunity is not simply to make today's organization work better, but to build an organization capable of

Find more content like this on the  
**McKinsey Insights App**



Scan • Download • Personalize



doing what was previously out of reach.

[Alexis Krivkovich](#) is a senior partner in McKinsey's Bay Area office; [Brooke Weddle](#) is a senior partner in the Washington, DC, office; [Holly Price](#) is an associate partner in the Houston office; and [Vik Sohoni](#) is a senior partner in the Chicago office.

The authors wish to thank Laetitia Bauchet for her contributions to this article.

---

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

Copyright © 2026 McKinsey & Company. All rights reserved.