

McMillan AI

The Human Factor that Makes or Breaks AI Transformation

Behind the hype, from the C-Suite: What it actually takes to get your organization ready for AI

Prepared by Stephanie Biernbaum & Jeff McMillan

June 2026

The headlines in early 2026 tell two contradictory stories. In one, AI is eliminating jobs at historic speed: entire functions removed, white-collar work disrupted, companies announcing layoffs in the same breath as AI investment. In the other, AI is a productivity revolution every organization needs to get ahead of before it's too late. Both stories are loud. Neither is particularly useful to the people sitting in the executive suite, trying to figure out what to actually do.

What gets lost between the two narratives is the human work. The tools are proliferating faster than any organization can absorb them. Most executive teams have a working thesis on AI strategy. But strategy and tools are only two legs of a three-legged stool. What is almost universally missing is the third: the organizational infrastructure to turn AI investment into workforce capability — the communication, the capability-building, the manager enablement, and the organizational design that determines whether AI transformation actually reaches the people expected to carry it out.

That gap is what this report is about. Across hundreds of conversations and interviews we've held with CEOs, CHROs, CTOs, and senior leaders, a consistent picture has emerged. The questions below surface in nearly every boardroom and executive team conversation.

The Questions Executive Teams Are Grappling With Right Now

Where are we, really, relative to where we need to be?

This is the question executives ask most quietly, because it's the hardest to answer with any rigor. Every leader has a sense of their organization's AI maturity, but very few have a defensible way to benchmark it. Which functions are genuinely ahead versus performing AI investment theater? How much runway does the organization have before the gap becomes a competitive liability?

The absence of a credible benchmark has a specific downstream effect: decisions about where to invest, what to prioritize, and how urgently to move are being made without a reference point. Organizations that have found a way to answer this question with rigor and specificity — not at the level of “we're on a journey” — have a material advantage in how they allocate resources and set expectations.

How do you connect AI strategy to business strategy — and make that connection visible to the workforce?

Organizations are distributing AI tools at scale. Licenses, tokens, and Copilot seats are flowing to employees who have been told, in some form, that AI is important. What is far less common is a coherent answer to the question every employee is actually asking: not what AI can do in the abstract, but whether their job will still exist, and in what form, when the dust settles.

Most organizations are not yet having that conversation coherently — and the reasons have evolved. First, it was not being ready to commit: models and tools were changing so fast that anything said publicly risked being wrong within months. Then it was fear of destabilization: committing to an AI strategy meant acknowledging roles would change, and leadership worried about triggering anxiety before the organization had answers to give. Now it's something more complicated: backlash, re-entrenchment, and a cultural skepticism that has made even well-intentioned communications land poorly.

And underneath all of it is the simple fact that this is genuinely hard. Saying enough to give the workforce a vision while remaining transparent about what isn't yet known, and agile enough to adjust as the landscape shifts, requires precision and courage that most communication playbooks weren't built for. The workforce is eager for more answers than leadership can yet provide. That gap doesn't close by saying nothing.

It is also, too often, a gap that strategy and technology functions are not structured to close. AI investment decisions are being made by technology and finance, while the human implications — how work will change, what will be valued, what employees need to learn — are handed to HR and L&D teams who were not in the room when the strategy was set. The result is a workforce that has heard “AI is a priority” but has not been given a meaningful answer about what that means for how they work. Closing this gap requires deliberate organizational architecture: a mechanism for the executive team's AI priorities to be translated into role-specific context, reinforced in the flow of daily work rather than delivered once in a town hall.

“Firms need to come up with an answer for when an executive gets stopped in the hallway and asked, is this going to take my job? I have had disagreements with other executives about what we think our answer should be.”

How do you actually prepare the workforce — at scale, without pulling people away from their day-to-day work?

Organizations know that doing nothing isn't viable. Equipping people means more than awareness: knowing how to use tools on the right use cases, how to redesign workflows to embed AI within them, and how to design and work alongside AI agents while keeping human judgment where it belongs.

The harder problem is reach. Getting the workforce as a whole up to speed, speaking a common language, without creating a second job for already-stretched employees, requires delivery that meets people where they are — and relevance: generic AI literacy is far easier to procure than content that connects to what a given employee actually does.

Real behavior change doesn't happen through training alone. It requires equipping the manager layer to reinforce new behaviors in the flow of work, building new norms, role-modeling them

from above, and the discipline of spaced practice to build skills into habits over time. This is where most current programs fail.

“Learning without any kind of follow-up is pointless. You want them to immediately apply it to something relevant to them — because if you build canned use cases that are out of context, you’ll lose attention in a second.”

Is the technology infrastructure actually ready?

Workforce capability is only one part of the readiness picture. For AI transformation to pay off, the organizational systems underneath it also have to work: data that is clean and trustworthy, a semantic layer that makes it useful without engineering intervention, governance structures that determine who can use what and under what conditions, and security architecture that allows AI to operate at scale without creating new liability.

Most organizations significantly underestimate how much of this work is undone. They have invested heavily in AI applications — the visible layer — without proportionate investment in the infrastructure those applications depend on. Executive teams that have confronted this clearly and made deliberate sequencing decisions are the ones that avoid deploying AI tools that underperform not because the tools are bad, but because the organization wasn’t ready for them.

How do you get enough visibility into AI utilization to make confident decisions?

Token-based pricing has introduced a budget dynamic that few organizations were prepared for. Annual budgets sized in Q4 are being exhausted by Q2. The question is no longer just whether employees are using AI — it’s whether the right employees are using it, on the right use cases, in ways that generate real returns.

Most organizations have no mechanism to answer that. Without visibility into which workforce segments are actively using AI tools, how effectively, and whether they’re focused on the right use cases, token spend becomes an unmanaged cost. The decisions that follow — more tokens or more headcount; broader access or deeper investment in fewer roles — are being made in the dark.

“We’ve put a significant amount of money forward to fund more access. The people who already have access — I don’t know if they’re using it well. I don’t know if they’ve got the right use cases. What are you learning from it, and are you actually seeing results?”

How do you have the conversation the board wants to have?

The board-level conversation has moved. The efficiency question — how much cost can AI take out? — is increasingly table stakes. What boards are now asking is harder: how is AI going to help us grow? What new opportunities does it create? And is the organization building the human capability to capture them?

This puts CHROs in a position they haven't traditionally occupied: accountable not just for workforce readiness, but for connecting talent strategy to AI-driven growth. That means speaking to which roles are most critical to AI-enabled value creation, whether succession pipelines are being built with AI fluency as a genuine criterion, and bringing the board a picture of organizational AI capability that looks more like an asset map than a completion dashboard.

How do you lead the redesign of work itself?

Underneath every workforce readiness question is a harder one most organizations haven't yet named explicitly: what should work actually look like when AI is doing some of it?

This is an organizational design question that belongs to the executive team — in particular, to the CHRO. Someone has to look systematically across the organization and ask: which tasks are candidates for AI automation or augmentation? Which workflows need to be reconceived rather than simply accelerated? Where does human judgment remain essential?

What was once an occasional transformation project is becoming a permanent organizational capability that needs to happen continuously, at the team level. The CHRO is best positioned to lead it, but doing so requires a methodology, a language, and a distributed network of people who can do this work systematically across their teams.

What Senior Leaders Said They Need

These needs surfaced consistently across our conversations — from CHROs accountable for workforce readiness to CEOs asking why AI investment isn't yet showing up in productivity. The specifics varied by industry and organization size. The underlying gaps did not.

A credible benchmark for where they actually stand

Most senior leaders have a felt sense of their organization's AI maturity — shaped by executive team conversations, input from their technology function, and how they read workforce engagement with AI tools. What almost none of them have is a defensible way to benchmark that impression: against comparable organizations, a defined maturity progression, or criteria that distinguish genuine capability from activity.

The absence of a benchmark makes prioritization guesswork. Leaders who cannot locate themselves on a maturity curve cannot make confident decisions about where to invest next or how to set realistic expectations with a board asking increasingly specific questions.

We observed organizations clustering into four broad stages — though few fit cleanly into any single one, and movement between them is rarely linear:

Experimenting — AI tools have been distributed broadly, but without a coordinating strategy. Usage is uneven, driven by individual initiative. Leadership is watching rather than directing.

Activating — The organization has made deliberate choices about where AI should play and has begun building communication and training infrastructure to support adoption. Early proof points exist; scale is the challenge.

Embedding — AI is being integrated into workflows and role design systematically. The manager layer is engaged. Measurement is in place, even if imperfect.

Transforming — AI capability has become a genuine organizational asset: self-sustaining, measured, and continuously updated. The organization builds on a compounding base rather than starting over with each new tool cycle.

Most organizations we spoke with are somewhere between Experimenting and Activating. The leaders furthest along shared one trait: they had found a way to be clear-eyed — with themselves and with their boards — about where they actually were, rather than where they aspired to be.

An upskilling approach that connects to the organization's specific priorities — and to specific roles

The most consistent frustration wasn't about access to AI education — it was about relevance. The market is crowded with programs that cover the fundamentals competently. What's missing is the layer below: content that speaks to what a particular organization is trying to do with AI, and what that means for a financial analyst, an operations manager, or a client service rep specifically. Participants drew a sharp distinction between AI literacy (knowing what the technology can do in the abstract) and AI fluency: the ability to apply it to real work, in context, on problems that matter to the business. The first is widely available. The second is almost nowhere.

Content that reaches people where they are

The problem isn't employee willingness — it's attention and friction. Passive e-learning, scheduled training blocks, and mandatory completion workflows lose a significant portion of their intended audience before the content lands. Format isn't a cosmetic question. It determines whether the program reaches the workforce at all.

"You need a platform where you can say, here's today's thing, click, I'm already listening."

A champion infrastructure with the organizational design to sustain it

Most organizations have attempted some version of an AI champion program. Most have found them hard to sustain. The gap wasn't motivation — it was structure. Without deliberate selection criteria, a formal development track, defined governance, and a clear deployment mechanism, champion programs produce individuals with personal expertise rather than organizations with distributed capability. What participants wanted was something more architectural: a credentialed internal cohort equipped to activate colleagues without ongoing external support.

"Where I could use help is how to structure champions and ambassadors in a way that gets maximum impact. It's one thing to say, you're going to be the champion here — but how do we organize that programmatically? How do we make sure it actually has some controls over it?"

A living channel, not a completed course

AI training built today might have as little as 6–9 months before it requires a full update. Internal teams cannot maintain content at that pace. What participants wanted wasn't a better course — it was a channel: something that keeps the workforce informed continuously, surfaces new use cases as they emerge, and never goes dark.

Measurement that means something to a CFO

Completion rates have become a running joke — not because tracking them is wrong, but because they've become a substitute for the harder question of whether anything is actually changing. What participants wanted was the ability to show, by function and level and geography, where adoption is taking hold and where it's stalling — connecting investment to demonstrated usage and value rather than leaving resource allocation decisions to be made in the dark.

"The AI companies can provide some data on how many agents you're building — but are they good agents or bad agents? The sentiment you're getting — what are people feeling, what are people thinking — is just as important as raw data for organizations to say, are we moving in the right direction?"

What the Evidence Points To

Across every conversation in this research, the same underlying picture emerged: organizations have invested heavily in AI tools and AI strategy, and are significantly underinvested in the human infrastructure required to make either pay off. That gap is real, it is widening, and it will not close on its own.

The leaders furthest along are not, in most cases, the ones with the most sophisticated technology. They are the ones who got clear-eyed early about where they actually were, made deliberate choices about the human work required, and built the internal capacity to sustain it without depending on external support indefinitely.

From the full body of this research, five conditions stand out for any organization serious about getting AI transformation right.

Build the communication infrastructure before the skepticism hardens.

The window for a direct, organization-wide conversation about what AI means for your people is not indefinitely open. The organizations that waited — for certainty, for the right moment, for a cleaner message — are now navigating cultural skepticism that well-intentioned communications are struggling to penetrate. Getting ahead of that requires precision and courage, not polish. And it requires recognizing that this is not a one-time message to be delivered and checked off. The AI landscape will keep shifting, roles will keep evolving, and the workforce will keep forming new questions. The organizations building durable AI capability treat internal communication as ongoing infrastructure — a channel that stays live, updates as the picture changes, and never goes dark.

Connect AI strategy to the work people actually do.

Workforce readiness that stops at AI literacy produces employees who know what AI can do in the abstract and struggle to apply it on Monday morning. The organizations making real

progress have translated their AI priorities into role-specific context — not a generic upskilling program, but a clear answer for each function about what AI means for how they work.

Build the manager layer or the program will stall.

Every organization that described successful AI adoption at scale pointed to the same thing: managers who could reinforce new behaviors in the flow of daily work, not just communicate the strategic rationale. Without deliberate manager enablement, capability-building programs produce individuals with new knowledge and organizations without new habits. The highest-leverage version of this goes further than communication. The organizations making the most systematic progress have equipped their managers to be the primary agents of work redesign at the team level — the people who can look at their function, decompose which tasks are candidates for AI augmentation, and rebuild workflows around new capabilities. This is where enterprise-wide AI transformation actually gets done: not in a consulting engagement or a central transformation office, but in a hundred team-level conversations led by managers given the methodology and the mandate. Equipping that layer gives the executive team and transformation function leverage they cannot otherwise manufacture.

Treat measurement as a design input, not a reporting afterthought.

The organizations with the most credible AI capability story — the ones who could walk into a board conversation with something real — had built measurement in from the beginning, not retrofitted it after the fact. Measurement that means something tracks behavior change over time, by function and level, in ways that connect AI investment to demonstrated business value.

Invest in the infrastructure, not just the application layer.

AI tools deployed on top of data, governance, and technology infrastructure that cannot support them will underperform — not because the tools are bad, but because the organization was not ready for them. The organizations that have avoided this failure mode are the ones that confronted the infrastructure question squarely and made deliberate sequencing decisions, even when that meant slowing the rollout. But infrastructure is not only a technology problem. The organizations making the most durable progress have been equally deliberate about building the human capital required to construct and maintain it: identifying the internal talent segments who will own the infrastructure layers, selecting them with intention rather than defaulting to whoever raised their hand, and deploying them on the right transformation tasks. That combination — the right technical foundation and the right people assigned to build and sustain it — is what separates organizations that get a working AI capability from organizations that get a capability that quietly degrades the moment external support steps back.

The through-line across all five conditions is the same: the organizations getting this right are the ones that treated AI transformation as an organizational design problem, not a technology procurement decision. The tools are widely available. The will to do the harder human work is the variable that differs.