

McMillan AI

The Human Factor that Makes or Breaks AI Transformation

Behind the hype, from the CHRO's seat: A field guide for getting your workforce ready for AI

Created 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 person sitting in the CHRO's seat, 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 at least a working thesis on strategy. What's missing is the 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. We spoke with CHROs and senior people leaders across industries, to get an honest account of what's actually happening inside real organizations. What we heard shaped both this report and the solution we set out to build: one designed from the ground up to meet the moment these leaders are navigating right now.

The Questions CHROs Are Grappling With Right Now

How do you have an honest, organization-wide conversation about what AI means for your people?

Employees are hearing about AI everywhere. They want to know what it means for them: for their roles, their careers, their value to the organization. Most companies are not yet having that conversation coherently. 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 honest 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.

"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 for what's coming — at scale, without pulling people away from their day to day work?

Organizations know that doing nothing isn't viable. The observation that your job won't be replaced by AI, but by someone who knows how to use it, has become a cliché but is still true. Equipping people means more than awareness: knowing how to use tools like Copilot on the right use cases, how to redesign workflows to embed AI within them, and increasingly, how to design and work alongside AI agents while keeping human judgment where it belongs.

The harder problem is reach. There is no shortage of AI education available. 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 it requires relevance: generic AI literacy is far easier to procure than content that connects to what a given employee actually does and what their organization specifically needs.

What makes this harder than many executives outside of HR appreciate is that 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 and role-modeling them from above, and the discipline of spaced practice to build skills into habits over time. Doing all of this in a coordinated way, without it feeling overwhelming, is an underappreciated art. It is also 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."

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

Token-based pricing has introduced a budget dynamic that few organizations were prepared for. Annual token 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. Licenses have been distributed; usage data has not. Without visibility into which segments of the workforce are actively using AI tools, how effectively, and whether they're focused on the right use cases, token spend becomes an unmanaged cost rather than a

strategic input. The decisions that follow, such as more tokens or more headcount; broader access or deeper investment in fewer roles; are being made largely in the dark.

“We’ve just put a pretty significant amount of money forward to fund more access to Claude. And it’s really expensive. 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 what are you doing with it, and are you actually seeing results?”

How do you have the conversation the board wants to have about AI strategy?

The board-level conversation has moved. The efficiency gains question of *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 channel and revenue 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 a genuine asset map than a completion dashboard. Most CHROs we spoke with are not yet having that conversation with their boards. But they know it’s coming.

How do you lead the redesign of work itself — not just prepare people for change, but determine what the change should be?

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, and it belongs 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? And how do you put the pieces back together into roles and structures that actually make sense?

The work of decomposing roles, disaggregating tasks, and rebuilding workflows around new capabilities has historically lived outside HR, in periodic consulting engagements or industrial engineering functions that most knowledge-work organizations never had. AI is changing that. What was once an occasional transformation project is becoming a permanent organizational capability that needs to happen continuously, at the team level. That changes who needs to own it. The CHRO is best positioned to lead it, but leading it requires a methodology, a language, and a distributed network of people who can do this work systematically across their teams. Most CHROs we spoke with recognize this need. Very few have a clear answer for how to build that capacity at scale.

What CHROs Said They Need

We asked every participant a forward-looking question: not what's broken, but what would actually move the needle. What follows is what we heard consistently across organizations at every stage of maturity.

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.

Delivery 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 consistently 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 described wanting was something more architectural: a credentialed internal cohort equipped with the facilitation tools 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 to remain current. Internal teams cannot maintain content at that pace. What participants described wanting 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 described wanting was the ability to show, by function and level and geography, where adoption is taking hold and where it's stalling. The same visibility that evidences behavior change is what makes token spend manageable: 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? Have I built 20 in the last month and I'm only using one, or using none of them? 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?"

AI at Work: What We Decided to Build — and Why

Research doesn't design a product. But it can make design decisions legible — explain why certain choices were made, what they were trading off against, and what problem they were actually trying to solve. What follows is an account of how what we heard shaped what we built. Not a feature list, but a design philosophy.

The new AI at Work program has three components, intended to run simultaneously: a podcast-based learning series for the full workforce, a manager program, and a champion development program for the internal talent pool positioned to lead AI transformation from within. The three are designed as a system, each reinforcing the others.

Start with the organization's own voice

The single most consistent signal across the research: employees don't need more generic AI content. They need to hear what AI means for their organization, from their leadership, in terms that connect to their actual work. Every component of the program is built around the organization's own voice. Each episode opens with a C-suite leader speaking to the company's specific AI priorities. Managers have a brief they can reinforce in team conversations. Employees have a reason to believe what they're learning connects to what the organization actually needs from them. This was a foundational design choice, not a feature. It determines everything else about how the program lands.

Build the capacity to redesign work — not just consume content

The workforce readiness question that carries the most long-term consequence is who is going to lead the systematic redesign of work as AI capabilities expand. The Manager Program is designed to build that capacity. Every manager works through a structured process of task decomposition and workflow redesign — not as a theoretical exercise but applied to their own team's actual work. The required build,

an AI tool designed and deployed for their team, is also the proof of concept: the manager has done the work, not just learned about it.

An organization that has put its managers through this program has built a distributed, continuously operating capacity for work redesign at the team level. This investment compounds over time and doesn't require a consulting engagement every time AI capabilities shift.

Reach the workforce where they are

A program that doesn't get consumed doesn't change behavior. The mass workforce program is delivered as a podcast series — mobile, episodic, designed for short windows in the flow of daily life. No log-in screen. No completion workflow. Background play, offline download, on the same device and in the same format as content employees are already choosing to consume. This wasn't cosmetic. It was a prerequisite for reach — and reach is a prerequisite for everything else.

Make the measurement worth something

The measurement architecture is designed backward from one question: what would a CHRO need to walk into a board conversation with something credible? The program produces that picture automatically — through weekly manager submissions, episode-level sentiment tracking, and an aggregated intelligence package giving the C-suite a longitudinal read on workforce sentiment and AI utilization by function, level, and geography. The same data that evidences behavior change makes token spend manageable: connecting investment to demonstrated usage and value, and giving the CHRO a basis for resource allocation decisions that don't have to be made in the dark.

Build a channel, not a course

After the initial series, the program feed stays live. New McMillan AI content keeps the workforce current on developments that matter. Client-produced content — internal success stories, use case spotlights, leadership updates — keeps the channel connected to the organization's own journey. The program doesn't end. It evolves.

Build an internal infrastructure, not a vendor dependency

The Champion Program is designed with a specific failure mode in mind: the cohort that produces personally expert individuals but leaves no lasting organizational capability. Selection is deliberate. The credential requires demonstrated skill growth, a completed AI application build, a facilitation practicum, and peer review contributions — not module completion. Cohort 1 champions graduate with facilitation guides and teaching toolkits to become coaches and facilitators for every subsequent cohort. The organization's capacity for AI capability building becomes self-sustaining. External support required decreases over time rather than remaining constant.

Hold the program accountable for behavior change, not content delivery

Every component has a required build at its core. Managers must design and deploy an AI tool for their team. Champions must build, document, and teach a function-specific AI application. These are credential requirements because the program's accountability is to behavior change, not content consumption. A credential that requires demonstrated application — a real tool, deployed on real work, reviewed by peers — is a more defensible answer to board scrutiny than a completion certificate.

What we chose not to build

We didn't build social features — no peer feed, no leaderboards, no public comment layer. Every feature that makes the platform feel more like a traditional training platform erodes the engagement advantage of the podcast format. Accountability runs through the manager layer and assessment architecture, not peer visibility.

A Note on What Comes Next

The organizations represented in this research are at different stages, facing the same underlying challenge: the investment in AI tools has outpaced the infrastructure to make those tools pay off. That gap is real, it is widening, and it is not going to close on its own.

What we heard consistently was that CHROs know what good looks like. They can describe it precisely: a workforce that understands what AI means for their specific role, managers equipped to lead the conversation rather than avoid it, an internal capability that compounds over time, and measurement that gives leadership something credible to act on. The vision is clear. The path to it has been harder to find.

AI at Work was designed in direct response to what this research surfaced. Not as a course to be completed, but as an infrastructure to be built — one that gives the CHRO a mechanism for the workforce conversation, a distributed capacity for work redesign, and an ongoing channel that keeps the organization current as the landscape continues to shift.

We are now moving into first deployment. If what you've read here reflects your organization's situation, we'd welcome a conversation about what that could look like for you.