Kyle & Co

THE HUMAN-CENTRIC AI COUNCIL

AI Transformation and Change Management

A guide for piloting and accelerating AI adoption

Table of Contents

INTRODUCTION

Leading a people and work transformation

AI adoption is accelerating—and so is the pressure to prove impact. Leaders want speed. Employees want clarity and stability. Meanwhile, the technology is evolving faster than most operating rhythms can absorb.

For HR teams responding to this pressure, AI transformation can look like a typical rollout: Identify a use case, find a tool, implement and hope it sticks.

But the teams that succeed treat AI as a people-and-work transformation, not just a technology upgrade. They build trust, capability, and the habits that change daily work.

That’s why the Human-Centric AI Council (HCAIC) exists--and why we created this guide. We don't need a shallow how-to developed by vendors with a solution to sell. We need a practical, human-centered playbook for moving from pilots to scale.

Alicia Miller answered that call, and spearheaded this resource which is grounded in real, first-hand change leadership experience. It tackles what makes AI adoption different—probabilistic outputs, trust gaps, constant evolution, and role shifts—and offers frameworks you can use immediately.

It's thoughtful, approachable, and extremely useful. Thank you to Alicia for putting in the work, to our HCAIC members who contributed, and to our underwriting partners who make this work possible.

— Kyle Lagunas
Founder, Kyle & Co
Executive Director, Human-Centric AI Council

INTRODUCTION

Leading a people and work transformation

AI adoption is accelerating, yet a lot of organizations are still early in their journey. HR executives are living the tension: Leaders want speed and measurable impact. Employees want clarity and stability. And the technology itself is evolving faster than traditional operating rhythms can absorb.
You may have seen the headlines: “most AI projects fail.” Many efforts do stall in the pilot phase, struggle to move beyond novelty, or break when scaled. But not all. The winners treat AI as a people-and-work transformation as much as a technology upgrade. As AI becomes embedded in products, platforms, and everyday workflows, one simple fact remains: work is changing.

Successful AI requires many ingredients—data and infrastructure, governance and risk management, vendor strategy, and smart opportunity scoping. But in every case, the bottleneck is the same: we are asking people to change how they do their work. We are asking managers to lead in a new way. We are asking organizations to rethink how work gets done.

And that’s why this guide exists. It is not a technical manual. It is not a vendor comparison. And it is not a vision piece about the future of work. It is a practical, human-centered guide for enterprise and people leaders to lead the people side of AI transformation.

The goal isn’t to make your workforce “use AI.” The goal is to help your organization evolve into new ways of working where AI is embedded, outcomes improve, and people remain at the center of how value is created.

HOW AI TRANSFORMATION IS DIFFERENT

1. PROBABILISTIC VS. DETERMINISTIC MODELS

Traditional software development involved clear logic with consistent, repeatable results. Today's AI models produce outcomes that can vary by user and also change overtime. This is because AI is trained to predict likely outcomes by calculating probabilities leveraging inherently incomplete and imperfect data.

2. CONTINUOUS FEEDBACK LOOPS AND LEARNING SYSTEMS

AI models today can be built with ongoing learning mechanisms, resulting in systems improving (or degrading) over time. These require ongoing monitoring and evaluation.

3. OPAQUE "INTELLIGENCE" AND + TRUST ISSUES

AI often feels like a "black box" that can be unpredictable and difficult to explain. These systems will need to build trust for adoption, particularly amidst high-visibility examples of "hallucinations" and failures. Highly regulated and risk-averse cultures will be expected to establish higher accuracy thresholds before deployment. To build trust, piloting and evaluation will be critical.

4. CHOICE AND EXPERIMENTATION

The number of models and ways to interact with models are expanding rapidly. Users aren't just presented a single solution. Rather, they may be presented with a multitude of options, requiring them to determine the best one.

---

5. CONTINUAL LEARNING AT THE FRONTIER

Models are continuously being refined and also paired with new features to improve performance. With each change, users may need to interact differently to not only improve outputs, but sometimes generate the same outputs. These constant changes require users to explore, experiment, and adapt continuously, rather than simply learn how to use a fixed set of features.

---

6. BLURRING BOUNDARIES OF ROLE DEFINITIONS

The technological evolution has repeatedly redefined what can be done by machines. AI isn't just executing tasks. It is now co-creating and generating. With the activation of AI demonstrating aspects of human agency (creativity, empathy, and interactive communication), we now must answer what is the role of human vs. machine while revisiting job and skill architectures.

---

7. PSYCHOLOGICAL AND IDENTITY THREATS

Technological innovations can shift the demand for different skills. People may feel their professional value propositions are diminishing with increasing aspects of their job being done by technology.

IS YOUR CULTURE *AI-ENABLING?*

Adopting artificial intelligence—especially AI assistants—is fundamentally different from previous waves of workplace technology. It doesn't just digitize capabilities, it creates new possibilities.

AI enables customization, insight, and more intelligent production, but requires users to learn new ways of working with technological assistants and adapting to technology that isn't fixed, but evolving. Transformation becomes continuous, shaped not just by technology deployment but by human curiosity, adaptation, and trust.

THE SUCCESS OF AI IN THE WORKPLACE DEPENDS NOT JUST ON TECHNICAL READINESS BUT ALSO ON *PEOPLE AND CULTURAL READINESS*.

To thrive, organizations must foster:

An Experimentation Climate – Leveraging AI requires choosing an AI model and also testing different approaches to work with it. The optimal design will be determined through comparative trials. Organizations must embrace test-and-learn approaches and experimental designs.

Continuous Learning – Employees must expect ongoing evolution of tools and motivation to learn and re-learn how to optimize with AI, not one-time mastery.

Growth Mindset – AI can feel overwhelming and challenging – employees need to believe they can adapt and become successful with new processes. Leaders and teams must embrace iteration and unsuccessful outcomes as a part of the journey to success.

Collaboration and Knowledge Sharing – The quality of AI outputs can vary depending on how employees interact with them. In an emerging field that continues to evolve, the cultivation of networks and sharing of best practices will be critical to unlocking value at scale. Yet this requires employees to share, not hoard, knowledge.

Psychological Safety and Retention – The onslaught of communications around staff reductions and AI replacing humans is creating additional resistance to adoption and knowledge sharing. Companies that can help employees feel more secure with their employment will unleash greater adoption, information sharing, and value delivery.

To build and reinforce this culture:

Habituate Learning. Schedule regular “AI hours” or brown-bag workshops where employees explore tools, or share successes and challenges using AI tools.

Reward Curiosity. Recognize individuals or teams who try novel prompts, build new workflows, or contribute to an internal AI “playbook.”

Celebrate Success. Ensure visibility to successful pilots and employee efforts aligned with the desired culture.

Normalize Failure. Celebrate what experiments teach—even when they don’t pan out—as part of your quarterly reviews and town halls.

Promote Collaboration. Deploy technology that makes collaboration easy. Build the motivation by encouraging and rewarding knowledge sharing, and actively addressing employee concerns. Build systematic processes that invite people to share.

Leadership Modeling. Have managers and senior leaders share and communicate AI usage and exploration in their own work.

AI INTELLIGENCE IS *RESHAPING* HOW WORK GETS DONE

AI is transforming the tasks, decisions, and collaborations that make up everyday workflows.

The emerging opportunities are driving businesses to build an AI strategy, investing in new infrastructure, and embracing a new era of technology — one where people are invited to work with intelligent tools in dynamic, adaptive
ways.

THIS SHIFT REQUIRES ***MORE*** THAN ACCESS TO TECHNOLOGY.

It demands a deliberate effort to build employee capabilities, ensure they have the opportunity to learn and experiment, and foster the motivation to embrace new ways of working.

The following guide outlines how organizations can activate these three levers of behavioral change to successfully embed AI into the fabric of enterprise operations.

BUILDING EMPLOYEE CAPABILITIES FOR AI USAGE

As AI is being productized into broadly accessible tools, it is creating new paradigms both for how we interact with technology and also how we build capabilities. Increasing AI aptitudes will require nuanced approaches that leverage traditional learning and development approaches, but adapted to the new AI landscape. Your learning & development roadmap:

01

STRUCTURED TRAINING:

Address AI training as a portfolio, not a single training. Develop the roadmap to skill building with sequenced steps and courses. Consider when to curate and recommend and when to develop in-house.

02

ROLE-BASED LEARNING PATHS:

Offer tiered and role-based training that match needs for different roles (e.g., AI developer, AI user, manager, senior leader) or departments (HR, Procurement, Sales, etc.).

03

ON-DEMAND LEARNING:

Provide on-demand training via tutorials, recorded demos, and other trainings. And don't just make them available -- ensure people know how to find and navigate them efficiently and effectively.

04

HANDS-ON EXPERIENCE:

Orchestrate hands-on workshops, AI sprints, and hackathons to activate learning. Use sprints to test different approaches, measure output quality, and document findings.

05

KNOWLEDGE-SHARING:

Leverage Centers of Excellence, Communities of Practice, and lunch-and-learns to build networks and prompt knowledge-sharing. Establish employee communication channels (Slack group, MS Teams channel, etc.) to build peer engagement and real-time solutioning.

HOW SHOULD YOU ADAPT YOUR LEARNING STRATEGIES FOR AI?

DON’T WAIT FOR THE PLAYBOOK. BUILD IT.

Best practices for specific contexts are still being discovered. Invite employees to contribute new insights, share pitfalls, and curate a living “AI playbook”.

EXPERIMENTATION FIRST

Traditional software training involves training “how-to” steps that produce the same outputs for each user. Modern AI requires a training shift toward introducing generalized capabilities and setting the stage for experimentation. Explorers and creators will be your most valuable participants.

FRAMEWORKS OVER SCRIPTS

Frameworks provide a tailored blueprint for problem solving in changing contexts. When users understand the problem statement, what a solution looks like, and how one would solve this without AI; they will have better foundations for building a successful AI-enabled solution.

TRAIN THE TRAINER

Enable employees to share their customized solutions and success stories with peers to scale learning and impact. Unlike traditional train-the-trainer, secondary trainers aren’t teaching what they were taught, but what they learned and created.

CONTINUOUS UPDATES

Technology features and offerings are changing fast, and so are perspectives on “what works best.” Plan for ongoing technology refreshers to address new features and model changes. Likewise, existing training and resources can quickly become outdated. Ensuring reference materials are not only relevant, but up-to-date, will require intention and effort

FRAMEWORKS OVER SCRIPTS

Frameworks provide a tailored blueprint for problem solving in changing contexts. When users understand the problem statement, what a solution looks like, and how one would solve this without AI; they will have better foundations for collaborating with AI on a digitized solution.

MAKE BIG THINGS FEEL SMALL

Carefully crafted systems are critical to help employees efficiently navigate the “AI expanse.” Consider implementing AI copilots that can query resource repositories, recommend courses, and provide tailored instructions for current tools and best practices.

PAIR AI TOOL AND PLATFORM TRAINING WITH:

POLICY OVERVIEWS

Clearly communicate company policies and processes for AI use, including governance and escalation pathways.

AI LITERACY

Provide employees at all levels a fundamental understanding of AI, including what it is, what it can do well, what the limitations are, how to recognize and evaluate risks, and other considerations for ethical and responsible use.

RESPONSIBLE AI

Ensure developers and users understand core principles of responsible AI and how they apply to their organizational context. Developing, deploying, and using AI responsibly can foster trust, mitigate risks, promote ethical use, and also improve adoption and performance.

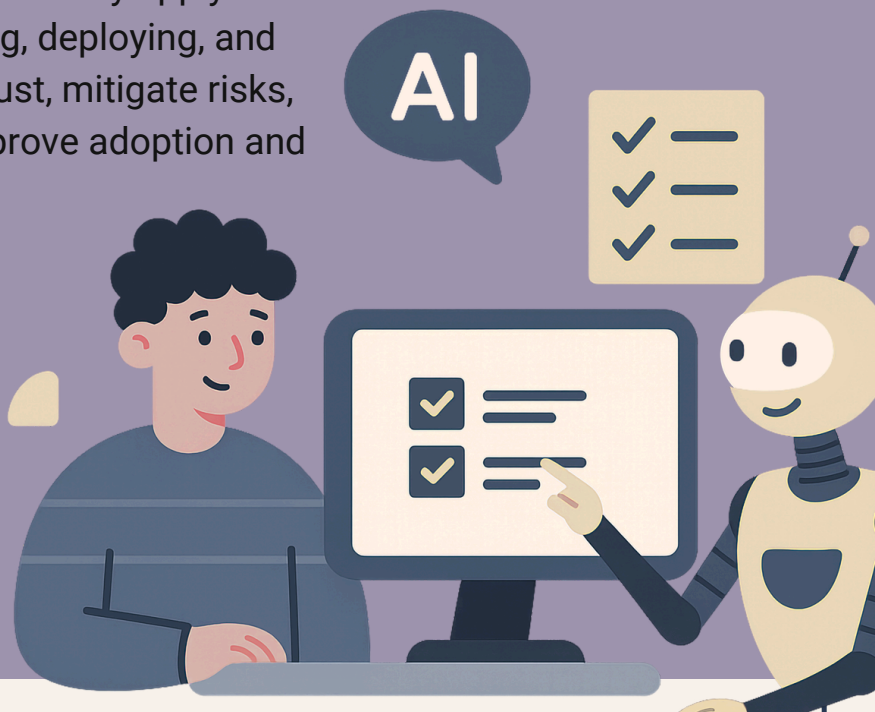

CREATING OPPORTUNITIES FOR AI USAGE

Even when motivated and capable, employees often lack the time or space to explore AI. The most experienced employees—the ones whose judgment matters most—are frequently the busiest. Without dedicated opportunity and access to tools, adoption stagnates.

01

DEDICATED EXPERIMENTATION TEAMS

Assign groups whose primary role is to explore, test, and codify best practices. That team can accelerate ROI and scale adoption by delivering proofs-of-concept, templates, and "practices that work" that business teams can adopt.

02

TIME ALLOCATIONS

Give employees "AI Fridays" or set exploration quotas—time explicitly protected for experimentation.

03

RESOURCE ACCESS

Provide open sandboxes, prompt libraries, and access to enterprise AI tools without bureaucratic barriers.

04

LEADERSHIP ACCOUNTABILITY

Ensure staff have the time and resources to succeed.

ONE KEY CHALLENGE TO BE AWARE OF IS OVERLOAD

While AI offers the promise to provide more efficiencies, it must first be learned – and sometimes built. The process of getting from zero to efficiency will require extra time and work. Building processes that may generate future time savings will still require more time up front to design and implement.

SUSTAINING MOTIVATION FOR AI USAGE

Motivation grows when the purpose is clear, when individuals see how AI connects to their goals, and when fears are addressed directly.

HERE ARE THE KEY TACTICS:

01

CONNECT OPPORTUNITIES TO INDIVIDUAL GROWTH

Help employees answer the question *"What's in it for me?"* with clear articulation of career benefits and opportunities that could be unlocked by upskilling or reducing time on activities that could be automated.

02

ADDRESS FEARS DIRECTLY

One of the most unique barriers to AI adoption is the fear of being automated out of a job.

03

FOCUS EFFORTS

The extent of information and pace of change can quickly feel overwhelming. You can help employees start the journey by communicating where to start and providing clear, tangible directions.

04

SHOW SUCCESS STORIES

Skeptics abound. For many, it takes seeing it work to believe it can.

CHALLENGES YOU MAY ENCOUNTER

MARKET SHIFTS

Companies are making strategic shifts to realign investments toward strategic ventures and AI-enabling teams, often resulting in layoffs in other areas. Automation announcements combined with layoffs have fueled perceptions that AI is causing job losses and enhancing skepticism.

BELIEFS AND PERCEPTIONS

Employees want authentic and transparent communications on organizational priorities and AI strategies, including job protections. And yet perceived intention and perceived outcomes will be more powerful than actual intentions.

FROM PILOT TO SCALE

Proofs of concept and pilots can be essential to de-risking AI adoption, building organizational confidence, and generating the evidence needed to scale across an enterprise.

Pilots provide a controlled, time-bound way to validate that an AI capability works in your environment—using your data, within your security and compliance constraints, and inside real workflows. These are particularly important for high-impact and novel solutions, but even more so for AI, where results are shaped not only by the technology, but also by how people use it.

When done correctly, pilots test more than technical feasibility: they quantify business value, surface adoption barriers early, and clarify what is needed for a successful transformation. Pilots may reveal what roles will shift, which capabilities are missing, where governance must be clearer, and how training and performance expectations need to evolve to sustain adoption.

By carefully choosing when to pilot and evaluating results, organizations can move from experimentation to enterprise transformation responsibly and effectively.

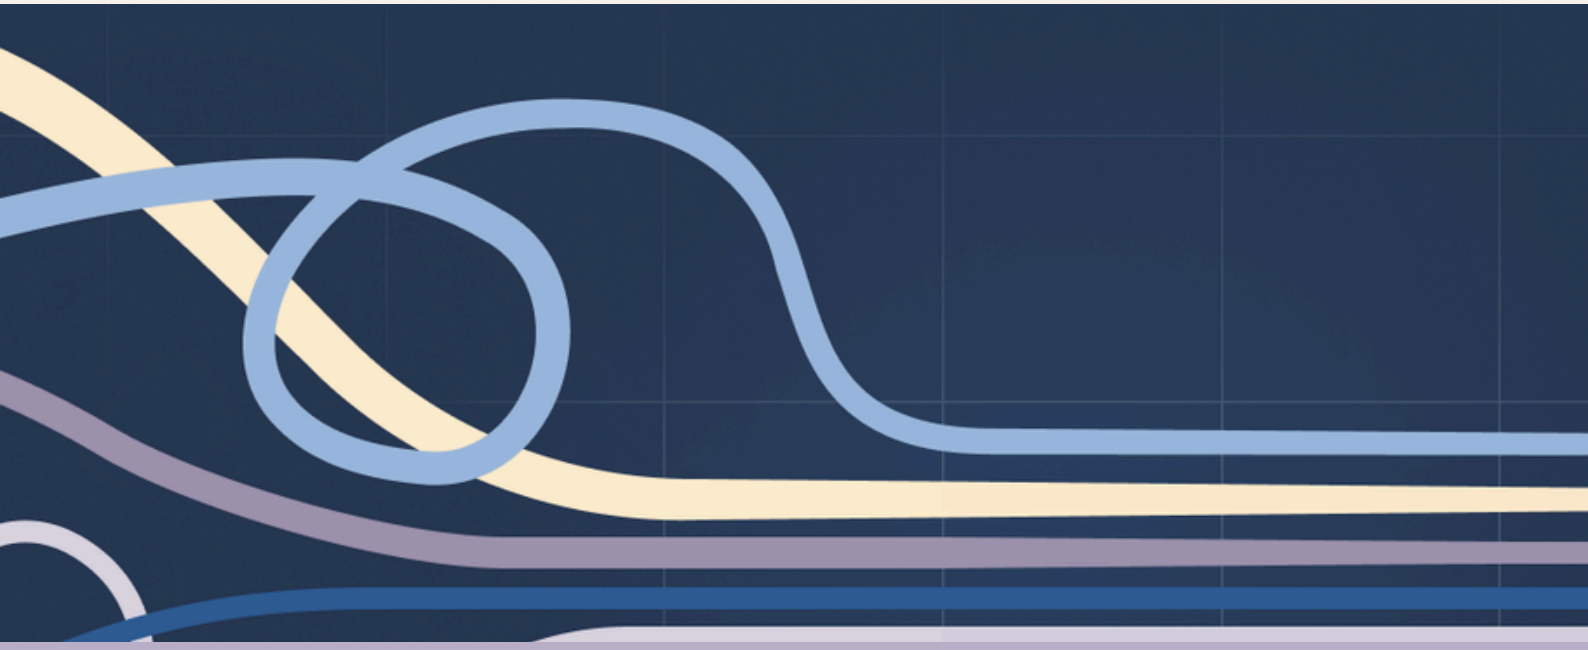

WHAT IS A PROOF-OF-CONCEPT (POC)?

A proof-of-concept is often a small-scale project to develop a prototype or test a new feature to determine technical feasibility. These are often developed in a research or testing environment outside of a production environment. The primary goal is to answer, “can it work technically?”

01

ENABLE INNOVATION IN NEW AREAS

Evolving technology generates new ideas for what is possible. Enabling staff to develop and test these ideas can provide a stage-gate to test whether these can be technically feasible.

02

CONFIRM TECHNICAL FEASIBILITY

Whether an internal development or procured solution, a POC helps validate whether solutions will function as intended within your own data and environment

03

IDENTIFY ADDITIONAL REQUIREMENTS

Running solutions in your environment may identify additional infrastructure, training, or other resource requirements.

04

BUILD THE BUSINESS CASE

Discussions and early stage development can inform minimum viable performance criteria and evaluate how well the solutions meets these. These may include speed, accuracy, and other metrics.

05

ASSESS VENDOR

Interactions with the vendor to ensure proper integration and trouble-shooting can inform whether you want to invest in this vendor or ecosystem.

WHAT IS A PILOT?

A pilot is a controlled, small-scale implementation of an AI system designed to test usability, effectiveness, and ease of adoption before committing to enterprise-wide deployment. Unlike proofs of concept, which are often exploratory and technical, pilots are operational tests in real-world conditions. They allow organizations to observe how employees, customers, and systems interact with the technology, and whether the anticipated benefits can be realized in practice. The primary goal is to answer, “will it work in operations and deliver value?”

Pilots:

Ensure more robust products and workflows: Test-case scenarios rarely cover the full spectrum of technical
challenges and human perspectives that arise in real-world testing. Pilots enable more robust feedback and stress-testing of the technical product.

Reduce risk: Pilots lower the financial and reputational risks by limiting initial exposure and impact of unintended consequences - such as technical errors or bias - to smaller areas.

Build buy-in: Successful pilots can empower and identify champions to support broader roll-outs.

Assess change readiness and adoption barriers: Pilots allow employees (or customers to test in a limited setting, assess usability and compatibility with existing workflows. Feedback from users reveals adoption barriers (confusion, resistance, training needs) and success enablers (ease of use, trust, perceived usefulness). This human feedback is critical for refining both the technology and the change management plan.

Evaluate business value: Pilots test the technology in real-world workflows, showing whether it drives measurable outcomes such as efficiency gains, cost savings, improved accuracy, or better employee experience. They help leaders answer: “Does this solution deliver enough value to justify enterprise rollout?”

PILOTS ARE BENEFICIAL WHEN

The business case is strong but unproven: There is a clear hypothesis of value (e.g., efficiency gains, improved customer experience) but no evidence under actual conditions.

The risk of failure is high: Deploying untested AI at scale could introduce bias, compliance issues, or workflow disruptions.

The context varies: Different business units, geographies, or customer segments may respond differently, so a pilot provides insight into variability.

The solution involves human adoption: Understanding how employees or customers react is as critical as validating the technology itself.

PILOTS MAY BE LESS APPROPRIATE FOR

Low-risk applications: E.g., when deploying non-critical updates or well-tested software components with minimal impact to core processes.

High-urgency situations: Some scenarios necessitate immediate action and faster deployment than is feasible with pilots. E.g., tools to enable remote operations during the COVID-19 pandemic.

STRUCTURING POCS OR PILOTS TO EVALUATE BUSINESS VALUE

New innovations are often competing for limited resources. Experimental designs offer a way not only to test new technologies and processes, but also to evaluate the extent to which these implementations impact outcomes.

Three potential approaches include:

Concurrent Evaluation: In this approach, an automated or AI-driven approach may be run concurrently or with the same data as a traditional process. These separate, but concurrent work streams, may evaluate the timeliness or accuracy of the two approaches.

There isn't a control group. A given group is tested with a treatment and without a treatment. For instance, you may have a manual approach of human reviews of employee support tickets to triage and determine channels for action.

An automated AI approach may be developed to review and assign a work stream without human review. By enabling human-evaluation and AI-evaluation on the same tickets, one can calculate time and accuracy of both approaches. For accuracy, an expert reviewer will determine what should have been the appropriate determination, to evaluate both the human and AI decisions.

Randomized Control Trials (RCTs):
RCTs involve randomly assigning participants or units (e.g., employees, customers, branches) into treatment and control groups. The treatment group will use the new AI solution. The control group will use the current process. By comparing results, organizations can attribute observed differences to the AI solution with high confidence. These approaches are often used to show causal impact, but due require random assignments to ensure the groups are comparable.

Quasi-Experimental Designs:
When randomization is impractical, organizations can still achieve credible estimates through methods such as propensity score matching. This technique pairs individuals or units in the treatment group with similar ones in the control group based on observable characteristics, helping reduce bias. For example, if a company cannot randomize which employees are given an AI-tool, it could identify employees with similar positions, competencies, experience levels, and performance ratings. By pairing up a treated employee with a similar untreated employee, it can create a control group for comparison.

These methods separate genuine impact from coincidental change, ensuring leaders can make evidence-based decisions about scaling AI solutions.

SCALING SOLUTIONS

Scaling is not simply replicating the pilot; it requires strengthening processes and ensuring the AI solution is robust enough for larger and more complex environments. Solutions that will be scaled are generally those determined to be low-risk and/or have met business requirements.
These may include performance criteria, successful demonstration of business value, and alignment with key strategic focus areas.

Enterprise deployments may not always be leveraged by or impact large-scale teams. For instance, the deployment of AI in evaluating employee sentiment in pulse surveys may impact only a single team. In this section, we focus on solutions that are anticipated to have broader employee usage and impact.

To prepare to scale a solution:

- Document lessons learned from the pilot, including technical, operational, and cultural insights.
- Implement critical product or infrastructure enhancements that address feedback during the pilot as well as prepare the solution to accommodate larger user bases.
- Refine and prepare for broader training and change management programs so employees understand how to work with the technology.

In some cases, it may be beneficial to scale incrementally rather than all at once. Decision criteria may include:

Geographic variation: Teams may require additional customization and/or greater hands-on support to ensure adoption and business value in varied contexts.

Limited support resources: Incremental scaling can ensure available support teams will have adequate capacity to engage teams for training and technical support. Focused engagement can take advantage of roll-out momentum and better ensure early and long-term adoption.

Cost controls and strategic deployment: Some AI tools will incur incremental costs per usage. For this reason, it may be beneficial to limit or incrementally deploy to monitor cost usage and focus roll-out on higher-value priorities or specific position types.

CONCLUSION

Momentum is the work

AI adoption doesn't fail because organizations lack ambition or access to technology. It fails because leaders underestimate how fundamentally different this shift is from what came before. AI is not a system you install and stabilize. It is a capability you introduce into a living organization, one that learns, adapts, and reacts alongside the people using it.

Across these chapters, one pattern becomes clear: progress depends less on choosing the "right" tool and more on creating the conditions where experimentation, trust, and judgment can take hold. Probabilistic systems demand tolerance for variation. Continuous learning systems require sustained attention. Opaque intelligence forces leaders to earn trust rather than assume it. And as AI reshapes tasks, roles, and identities, organizations must actively support people through that transition instead of hoping it resolves itself.

This is why pilots matter. This is why culture matters. And this is why human-centric design is not a philosophical stance but an operational necessity.

The organizations that move forward are not the ones racing to automate the most work. They are the ones building confidence, capability, and clarity as they go. They treat adoption as a process, not an event. They create feedback loops between technology, people, and outcomes. And they measure success not by how fast AI is deployed, but by how well it is understood, trusted, and used to deliver real value.

There is no finish line for AI transformation. The work is ongoing by design. Momentum is not a phase you reach. It is the result of deliberate choices made over time about how technology and people move forward together.

The question is no longer whether AI will change HR. It already has. The real question is whether organizations will lead that change intentionally or react to it after the fact.

Human-centric AI is not about slowing down innovation. It is about making progress sustainable.

THANK YOU TO OUR PARTNERS

workday. findem Warden AI isolved GLIDER CodeSignal
Humanly goodtime RIVAL Gem HiBob SmartRecruiters An SAP company

ABOUT KYLE & CO.

Kyle & Co is a modern research and advisory firm helping HR and talent technology leaders make smarter decisions, faster.

By blending practitioner insight, rigorous research, and strategic advisory, we uncover the signals that matter most—then translate them into action. Our work spans custom research, workforce strategy, and market analysis, supporting solution providers and HR leaders alike. Kyle & Co is the team behind the Human-Centric AI Council, the Transformation Realness podcast, and industry-shaping reports on quality of hire, workforce planning, and responsible AI.

We work with clients like Workday, ADP, GoodTime, Findem, and SmartRecruiters.

Learn more at www.kyleandco.com

ABOUT THE AUTHORS

Alicia Miller, member of the HCAIC was the lead contributor in developing this resource.

Alicia Miller is a seasoned leader in innovation and organizational transformation driving tangible business impacts on revenue generation and operational efficiencies. Leading AI literacy and enablement, she has extensive experience working with analytical, research, and product teams to drive data-driven, AI-enabled, and human-centric services and operations. Alicia co-founded an international community on behavioral insights where she drove innovation, experimentation, and evidenced-based practices, growing participation to more than 30 countries. Alicia holds a Masters in Industrial/Organizational Psychology from the University of Maryland-College Park.