WORKFORCE TRANSFORMATION:

The Next Enterprise Growth Function

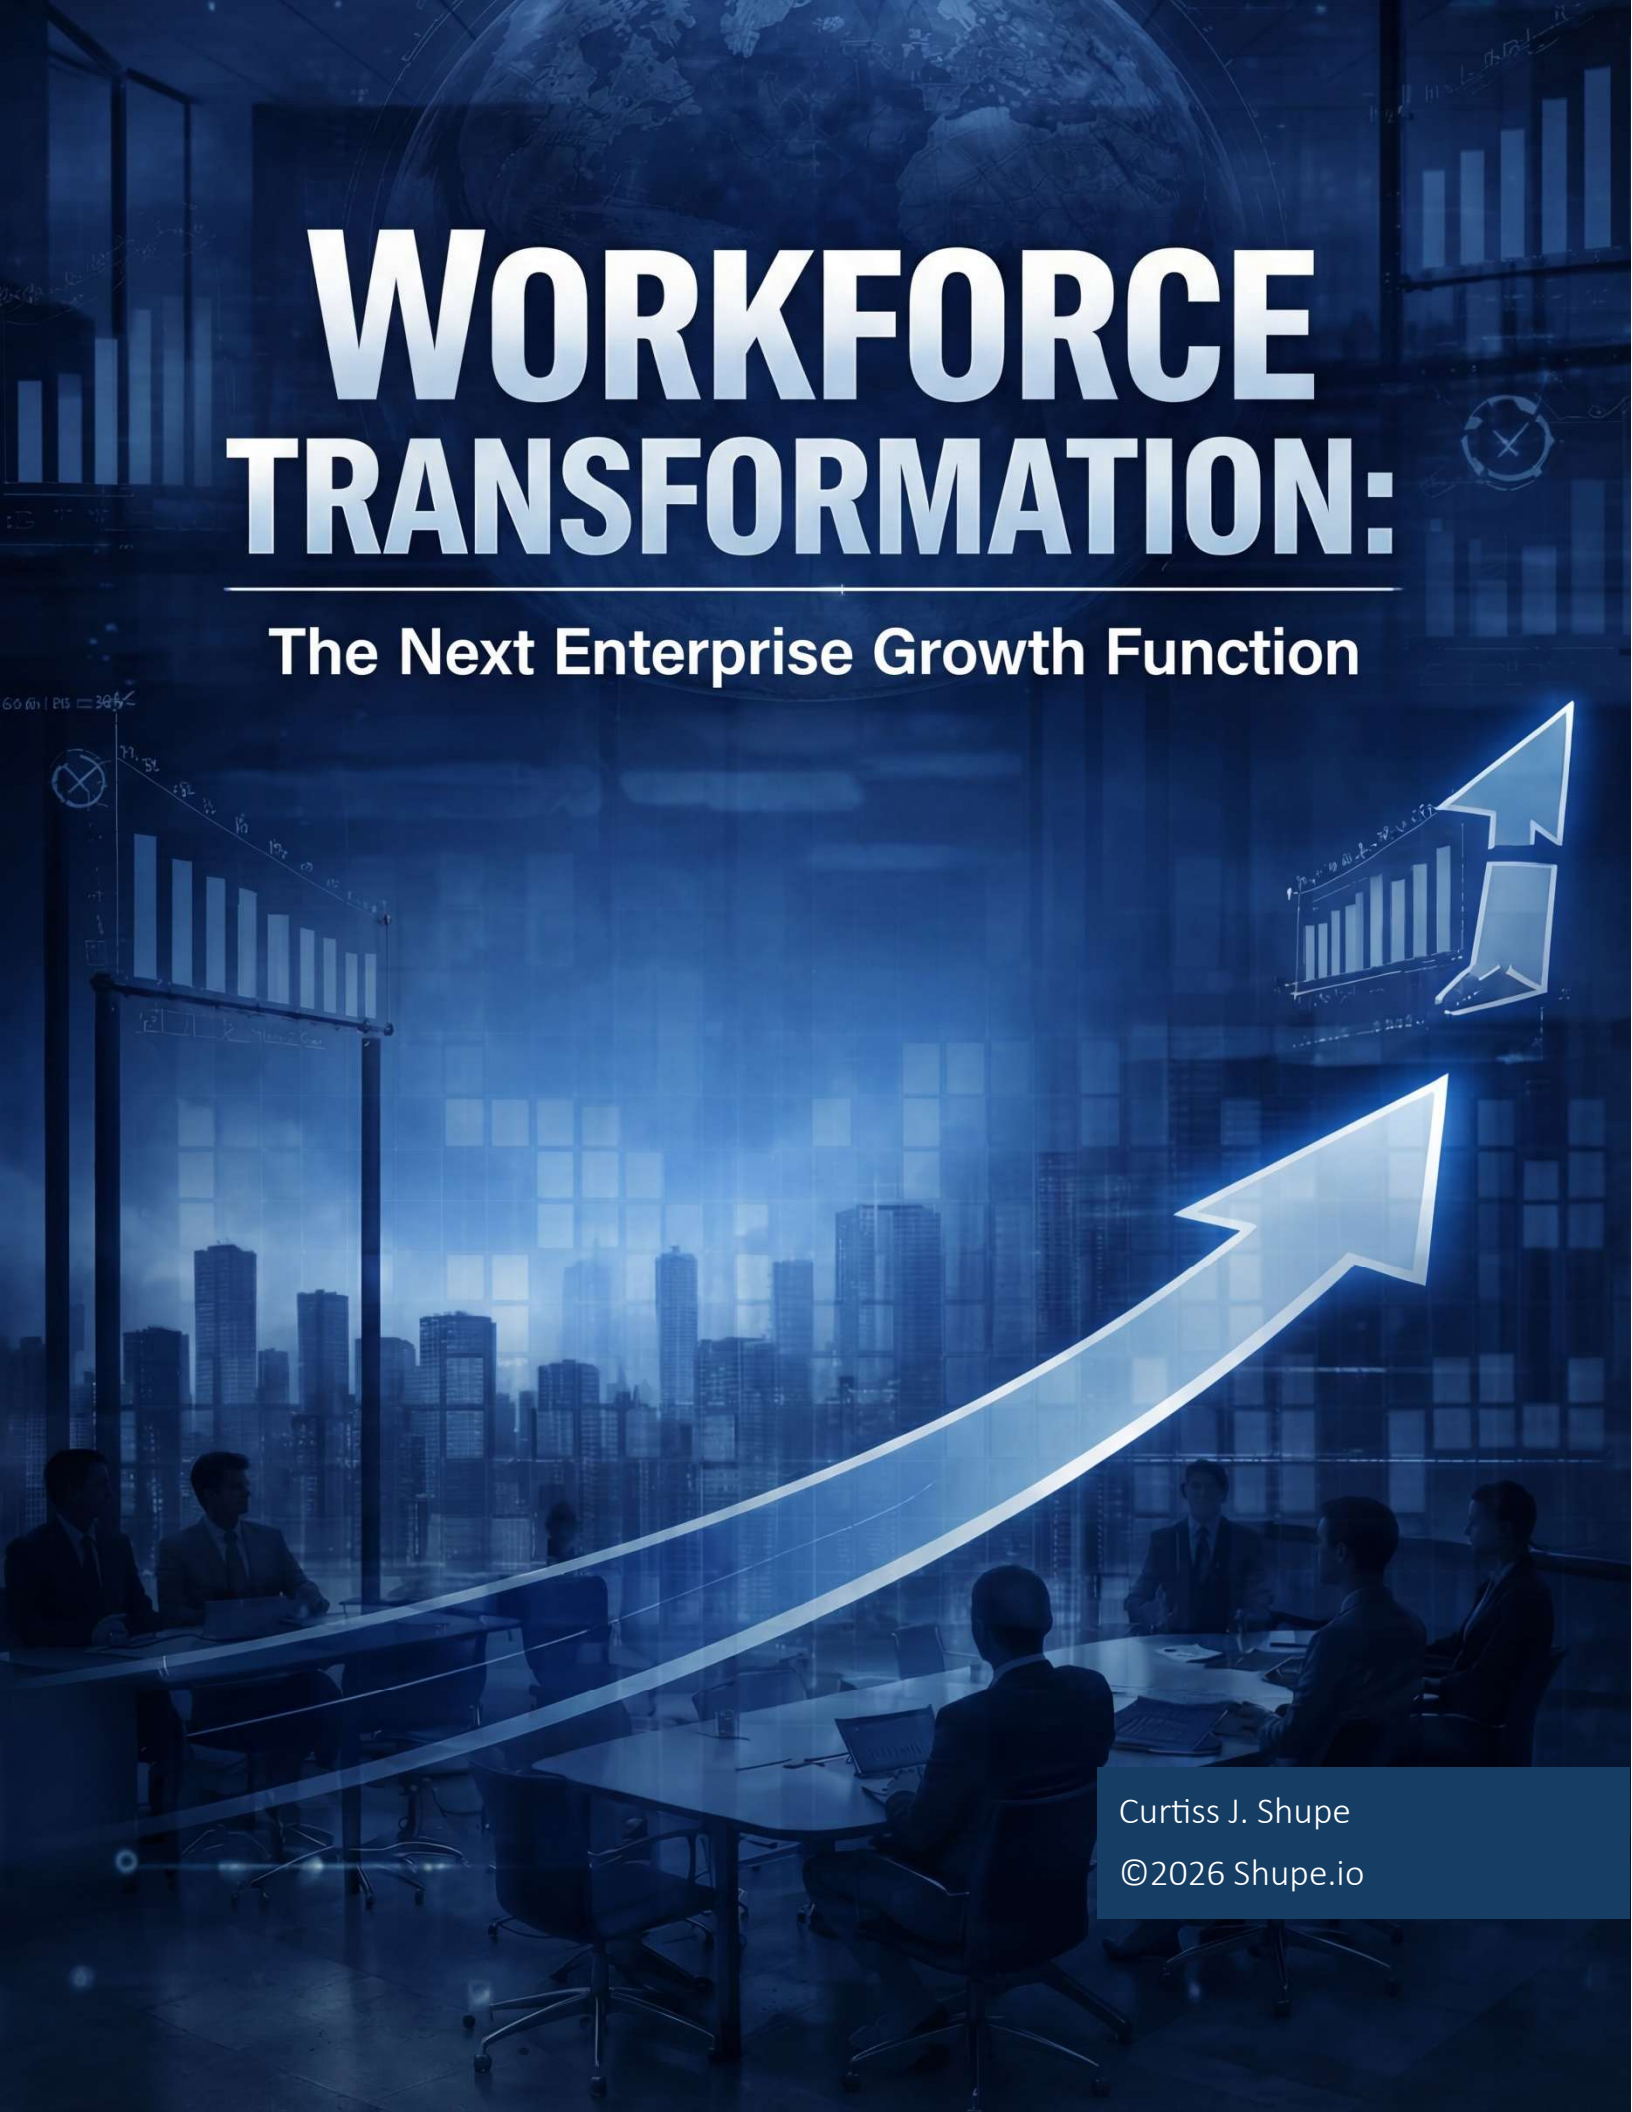

Curtiss J. Shupe

Workforce Transformation: The Next Enterprise Growth Function

Workforce Transformation is moving from AI adoption to work redesign.

That is the central shift. The question of whether AI is important to the market, for the most part, is no longer a question. The more interesting question is how large organizations are adapting their work, their people, their management systems, and their business models fast enough to keep pace with the advancement of AI. The penetration of AI adoption is widespread. In a report from Stanford, as of 2024, 78% of organizations are utilizing AI, compared to 55% the prior year. According to a report from McKinsey, 71% of organizations are utilizing generative AI on a regular basis within at least one function of their business. The mechanisms for extracting value from AI adoption, however, are not well developed. According to a report from McKinsey, fewer than one-third of organizations are following most of the key adoption and scaling practices for value capture. Fewer than one in five organizations are measuring well-defined KPIs for generative AI solutions.

My perspective is that we have now reached the first phase of industrialization in this field. The limiting factor is not the availability of models but rather the ability of the enterprise to leverage these models. The next 12-24 months will not be driven by those who have access to more AI tools but rather those who have the ability to leverage the tools they have to redesign their enterprise, manage risk, and team humans and AI. This is because it's an obvious conclusion drawn from observing how this field has accelerated in terms of technology and how enterprise maturity has lagged. According to Stanford University, this field has seen significant progress in terms of the ability to afford and access this technology. There has been a more than 280-fold decrease in inference costs for GPT 3.5 level performance between November 2022 and October 2024. There has also been a decrease in the gap between open-weight models and closed models in certain benchmarks, where it's been reduced from 8% to 1.7% in just 1 year. However, more than 80% of organizations have yet to see any EBIT impact in terms of generative AI in an enterprise context, according to McKinsey.

The current state of the field

The current state is best described as broad adoption, uneven value, and insufficient redesign.

No longer is AI a niche phenomenon in the enterprise; it is now a mainstream phenomenon. But it is still a shallow phenomenon. It has spread to more business functions. A report from McKinsey claims that now organizations, on average, use AI in three business functions, most frequently in the fields of IT, marketing and sales, service operations, product and service development, software engineering, and knowledge management. Well done! But it is still a minority of enterprise functions. So, we are not in a period of mature transformation. We are in a period of expanding deployment with selective pockets of real redesign.

The speed at which the technical frontier is growing is far outpacing the ability of enterprise redesign systems to react to it. In addition, according to Stanford's 2025 AI Index, not only is artificial intelligence becoming more efficient, affordable, and accessible, but the technical frontier is becoming crowded and competitive. This means that large organizations cannot afford to have long planning cycles, static job architectures, or post-hoc change management approaches because they would have already moved on by the time they react to the changes.

Curtiss J. Shupe

The value story is real, but still concentrated. McKinsey has found that "organizations are increasingly seeing revenue growth and cost savings within business units through generative AI, although the financial impact across the entire organization remains limited." That is a critical finding. It tells us the technology is already valuable, and the value is not necessarily translating into business improvement. The basic challenge of the field is not to demonstrate the potential for AI to help. It is to build the organizational foundation for help to become value.

The value story is real, but it is still concentrated. McKinsey has found that "organizations are increasingly seeing revenue growth and cost savings within business units through generative AI, while financial impact across the entire organization remains limited." That is a very important finding. That finding tells us that the value story is real, but it is not necessarily translating to business improvement. What is the central problem in the field of AI, then? It is not to prove the value of AI. It is to build the conditions in which AI can help.

It is more likely to be a transformation rather than a destruction of the workforce, at least in the near future. According to the International Labor Organization update in 2025, one in four jobs is potentially exposed to generative AI, but it is more a case of transformation rather than destruction. It is also true, as highlighted by this update, that there will be no complete automation of jobs because there will always be a chance for humans to get involved. This is the correct perspective for large enterprises, not job destruction but rather job transformation in terms of distribution and bundling.

Nonetheless, in terms of skill, it still lags behind exposure. This is because, according to OECD, training supply may not be sufficient to meet the rising demand for basic AI literacy. The 2025 survey carried out by BCG also drives this point home in an operative sense: only 36% of employees are considered to be adequately trained in AI, and the percentage of employees using AI regularly has leveled off at 51%, with a notably higher percentage of AI usage among leaders and managers. This is not a trivial point; it is stating that the next bottleneck in the field is not just access to technology but also distribution and enablement of the same.

However, there is a rise in governance pressure along with the adoption. NIST’s Generative AI Profile clearly positions generative AI as a technology with new risks, for which special actions for managing the risks are necessary. According to Stanford, despite a sharp rise in AI-related incidents, standardized evaluations of responsible AI are rare among big model developers. It is moving to a new era where governance maturity is a prerequisite for scaling, rather than a barrier.

The trajectory over the next 12–24 months

Over the next 12–24 months, I expect five shifts to define the field.

The first change is to move from horizontal copilots to workflow native human-agent systems. The general-purpose assistants will undoubtedly be around in the future, but the area where the impact of the enterprise will be felt is in the workflow where the AI is being applied to solve a well-defined set of tasks. This is where the economics will be felt. BCG has written that business value is achieved not just by deployment but also by deep workflow redesign. Deloitte’s 2026 enterprise survey found

Curtiss J. Shupe

that organizations succeeding with AI are intentionally using AI in their business workflows and pairing humans with machine intelligence.

The second shift is from job-centric to task-centric and capability-centric workforce planning. With AI systems drafting, analyzing, summarizing, and increasingly initiating the work, the enterprise of the future needs to think in terms of tasks and decisions, and risks, not just jobs. The results of the ILO study, namely, the high percentage of exposure and the differentiation in implementation, are in line with the second shift. The results of the McKinsey study, namely, the differentiation in the expected impacts in terms of headcounts, and the fact that service operations and supply chain are likely to experience decreasing headcounts, while IT and product development are likely to experience increasing headcounts, also confirm the second shift. That is, the enterprise of the next two years will not be faced with one AI workforce issue; it will be faced with many.

The third shift is from experimentation to platformized execution. As model performance is improving rapidly and costs are coming down, the basic question in strategy is no longer which model to choose but how we design an enterprise architecture to route through models, manage access, assess results, log all decisions, monitor all risks, and manage all agents in a similar fashion. The data on cost and competition at Stanford supports the platform approach, and risk perspective also supports governance by design as per NIST.

The fourth shift is from periodic change management to continuous adaptation. This means there will be a constant process inside the enterprise to turn technical changes into redesign of work, roles, manager support, and communication of changes to the workforce. This is not about communicating AI changes to the workforce. This is about constantly re-specifying who does what, where judgment is the control mechanism, how exceptions are handled, and what success means inside the system. The inability of BCG to successfully deploy employee adoption and training speaks to the fact that it will not be periodic, but continuous.

The fifth shift is from AI as assistance to AI as managed digital labor. Microsoft indicates that the expectations of enterprise leaders are that agents will be somewhat or substantially integrated into the AI plan in the next 12 to 18 months. BCG indicates that three-quarters of the workforce believe that AI agents are vital to future success, yet only 13% report that they are now broadly integrated into the enterprise and only one-third report that they understand how agents function. I think that is one of the most important market indicators and that the expectations for agentic systems are far outstripping enterprise readiness. The next shift is not going to be about the arrival of agents, it is going to be about the ability of the enterprise to effectively utilize them without sacrificing control, clarity, or the trust of the workforce.

What this means for large enterprises

For large organizations, the space is converging on a new management issue: how to build an operating system for workforce transformation.

This is not an HR subtopic, a technology subtopic, or a change program. This is a strategic capability that is at the intersection of operating model, finance, workforce planning, role, risk, learning, and

business execution. The evidence now supports the same conclusion from multiple directions. McKinsey's research shows that value is correlated with having a defined road map, tracking key performance indicators, and adoption practices. BCG's research shows that workflow redesign, training, and leadership support are the actual drivers of value. Deloitte's research shows that scaling requires the combination of people with machine intelligence within business workflows. These are different ways of describing the same thing: a systemized process for converting AI progress into enterprise redesign decisions.

That system, of course, must address a regular set of questions. What workflows need redesigning first? What tasks need to remain human, become hybrid, or shift to agents under supervision? What new roles will emerge, and which existing roles need rebundling? Where will spans, layers, and organizational ratios shift? What needs building broadly, and what needs building deeply? How will managers manage teams that include both people and AI systems? What controls are required before agents can act, not just advise? These are no longer purely hypothetical questions of organizational design. They are increasingly becoming questions for organizational executives.

The organizations that move fastest will not be those that automate most aggressively. They will be those that make work legible enough to redesign. This requires clarity of process, clarity of decision rights, performance baselines, and risk ownership. Without this foundation, AI will not create leverage, it will accelerate noise, increase exceptions, and exacerbate fragmentation. The trajectory of the field will favor enterprises that excel both in architecture and governance, as well as organizational design and workforce enablement.

The most likely failure modes

The most common failure mode is likely to be the assumption that AI adoption is proof of transformation. It is not. High adoption rates can be accompanied by low enterprise value, poor governance, AI shadow, and lack of structural transformation. McKinsey's evidence on immature adoption practices and low EBIT impact is the most obvious warning.

Another failure mode would be underinvesting in manager readiness and frontline enablement. According to BCG's data, leadership buy-in and training have a big influence on adoption but are currently weak in many organizations. Organizations that invest in training and leadership but don't adequately prepare managers would find it hard to drive AI adoption into team-level performance.

Another failure mode to be considered is the one in which agents are forced into workflows prior to achieving role clarity, oversight, and trust within the organization. The world is rapidly moving towards agentic systems, and while progress is rapid, understanding is still shallow. Enterprises that are using AI agents without clear authority boundaries, auditability, and escalation are likely to generate risks faster than they generate value.

Conclusion

My position is that the field is becoming a discipline of enterprise redesign.

Curtiss J. Shupe

Workforce Transformation: The Next Enterprise Growth Function

The next 12 to 24 months will not be defined by better models, although better models will continue to be built. They will be defined by whether large organizations can build a repeatable capability to redesign work around human-AI teaming. The evidence is already pointing to what the future will look like. It is broad in terms of adoption, growing in terms of capability, real in terms of value, likely to change in terms of jobs, significant in terms of skills, increasing in terms of governance, and moving from concept to execution in terms of agents.

This represents a true blank-sheet opportunity for organizations that are willing to think of workforce transformation as a strategic function rather than a program. The fundamental question now is not whether artificial intelligence has useful work to do. The question is whether the enterprise can be designed to make effective use of that work. The organizations that solve this problem will drive the pace of the next generation of large enterprise performance.