THE
kyndryl
institute

Trend Topic: **Leading in the Age of AI**

By

Sangeet Paul Choudary

Author, 'Reshuffle, Platform Revolution'
Founder, Platformation Labs
Senior Fellow, University of California Berkeley

Sangeet Paul Choudary is a global authority on AI strategy, platform economics, and industry transformation.

He is the award-winning author of 'Reshuffle: Who Wins When AI Restacks the Knowledge Economy', which received the 2025 Thinkers50 Strategy Award for the most impactful idea in strategy.

Sangeet has advised executive teams across more than 40 Fortune 500 companies and served on advisory councils for leading global institutions.

He is currently a Senior Fellow at University of California, Berkeley and Non-Resident Scholar at Dartmouth.

Explore more from The Kyndryl Institute
kyndryl.com/institute

Stop managing your AI 'workforce', start allocating AI capabilities

In the early 1980s, Detroit believed it was at the cusp of the next phase of industrial progress.

General Motors had poured more than forty billion dollars – an almost fantastical sum for the time – into robots that promised to rescue American manufacturing from its long slide.

Automation would make things faster, cheaper, better – replace expensive unionized labor with tireless machines, preserve the familiar workflow, and capture the productivity gains. What unfolded, though, was a reminder that brute technology adoption into legacy systems rarely solves the problems plaguing those systems. GM positioned robots exactly where human workers had stood. Job classifications remained intact, and the pacing and sequencing of the assembly line was left untouched. The robots were treated, in effect, as compliant substitutes for labor rather than as catalysts for redesign.

MIT's Made in America report, chronicling the consequences, explains that the conceptual framing of robots as units of labor made automakers overlook how much they altered the underlying dynamics of production. Robots introduce new coordination demands and shift the economics of flow and quality. They invite reconsideration of layout, timing, and supervision. Yet these implications stayed dormant because the organizations

adopting the robots used the existing workflow as the reference frame. The underlying architecture was preserved, and so were the constraints that limited its performance.

Toyota, in contrast, approached the same tools from a different vantage point. Instead of asking which tasks robots might take over, its engineers examined how the presence of robotic capabilities changed the logic of the production system itself. They reconfigured plant layouts, redesigned work cells, and tied quality feedback tightly into every movement of the line. Human workers shifted from task execution to managing production lines, enabling rapid detection and correction of variation. The factory became a coordinated human-machine environment rather than a human workflow with machines attached.

Both companies had access to similar robotic technologies – yet, only one reconsidered the architecture around them. The long-term divergence that followed – visible in productivity and competitive advantage – was rooted in that initial framing choice.

The rise of agentic execution – AI agents that accomplish complex workflows in pursuit of goals instead of being limited to task performance – positions us at a similar moment today. Executives describe agents as digital co-workers, AI interns, or virtual employees, as if the most natural way to adopt a new computational capability is to place it inside an existing job structure. The metaphor seems harmless, even helpful, because it allows a complex technology to fit within familiar mental models. But it produces the same conceptual narrowing that constrained Detroit four decades ago.

Once an agent is cast as a co-worker, the unit of adoption becomes the role, not the workflow or the system in which that workflow is embedded. This pattern creates the impression that the firm is modernizing its operations, even though the underlying architecture remains unchanged.

The co-worker framing also changes the expectations placed on leadership. A CEO who absorbs the metaphor too literally begins to imagine a blended workforce composed of people and agents. Questions emerge about supervision, performance management, and headcount equivalence. Should agents be onboarded? Should they be managed by HR? Should there be guidelines for managing digital workload? These questions carry the logic of workforce management into a domain where workforce logic is a poor – often dangerous – fit.

The metaphor seems harmless, even helpful, because it allows a complex technology to fit within familiar mental models. But it produces the same conceptual narrowing that constrained Detroit four decades ago.

What is missing in this framing is the deeper shift in the structure of work. Agentic systems do not work with the constraints of human labor and the workflows designed around them should not reflect those constraints either. They create opportunities to redesign workflows and reallocate decision rights across the organization, but these implications remain lost if the organization treats agents as occupants of roles rather than sources of new capability. The metaphor of the co-worker directs attention toward staffing questions and away from questions of architecture. This results in local gains showing up in the form of individual employees and teams accomplishing more but fails to unlock productivity that would come from reimagining the system of work around these new capabilities.

We can better appreciate Toyota's insight in the 1980s when we closely examine how its factories changed shape as robots appeared on the floor. Engineers at Toyota did not imagine machines stepping into human roles the way Detroit had. They began instead with a different, almost architectural, question: What becomes possible when a new capability enters the system? A robot could weld with greater consistency than a human, but the implication was never limited to welding. It altered the pacing of the line, the spacing between stations, the timing of inspections, and the logic of how defects were monitored. Quality cycles could be tightened because the system could rely on predictable execution. With workers freed from monitoring each movement, decision rights moved from the shop floor to those who could monitor the behavior of the system. The physical layout of the plant changed in response to these shifts as well.

Once new capabilities, unbounded by legacy limits, enter a system, the system either adapts to exploit them or resists and diminishes them. Toyota adapted, while Detroit resisted. The Japanese saw a new system, while Detroit's factories only saw a new workforce. The lesson from Toyota is that capabilities, not headcount, are the true units of analysis. When the capability changes, the architecture must adjust.

This narrow framing already gives rise to predictable fallacies in today's discourse around AI adoption. When agents are framed as digital labor within established structures, executives tasked with driving adoption of these new capabilities fail to pause and reimagine the system.They continue working on the assumption that jobs remain stable while tasks get redistributed. But once an agent reliably executes a cluster of tasks, the boundary of the traditional job dissolves, and workflows reorganize around the new capabilities introduced. The notion of augmenting a role presumes the role remains stable, which is highly unlikely once agents are effectively implemented.

When we stop seeing agents as digital co-workers and treat them as strategic, evolving capabilities that expand what the system can do – and therefore what the organization can be – the real leadership challenge comes into view. This is because every time agentic capabilities improve and gain greater reliability, three important changes happen.

Once new capabilities, unbounded by legacy limits, enter a system, the system either adapts to exploit them or resists and diminishes them.

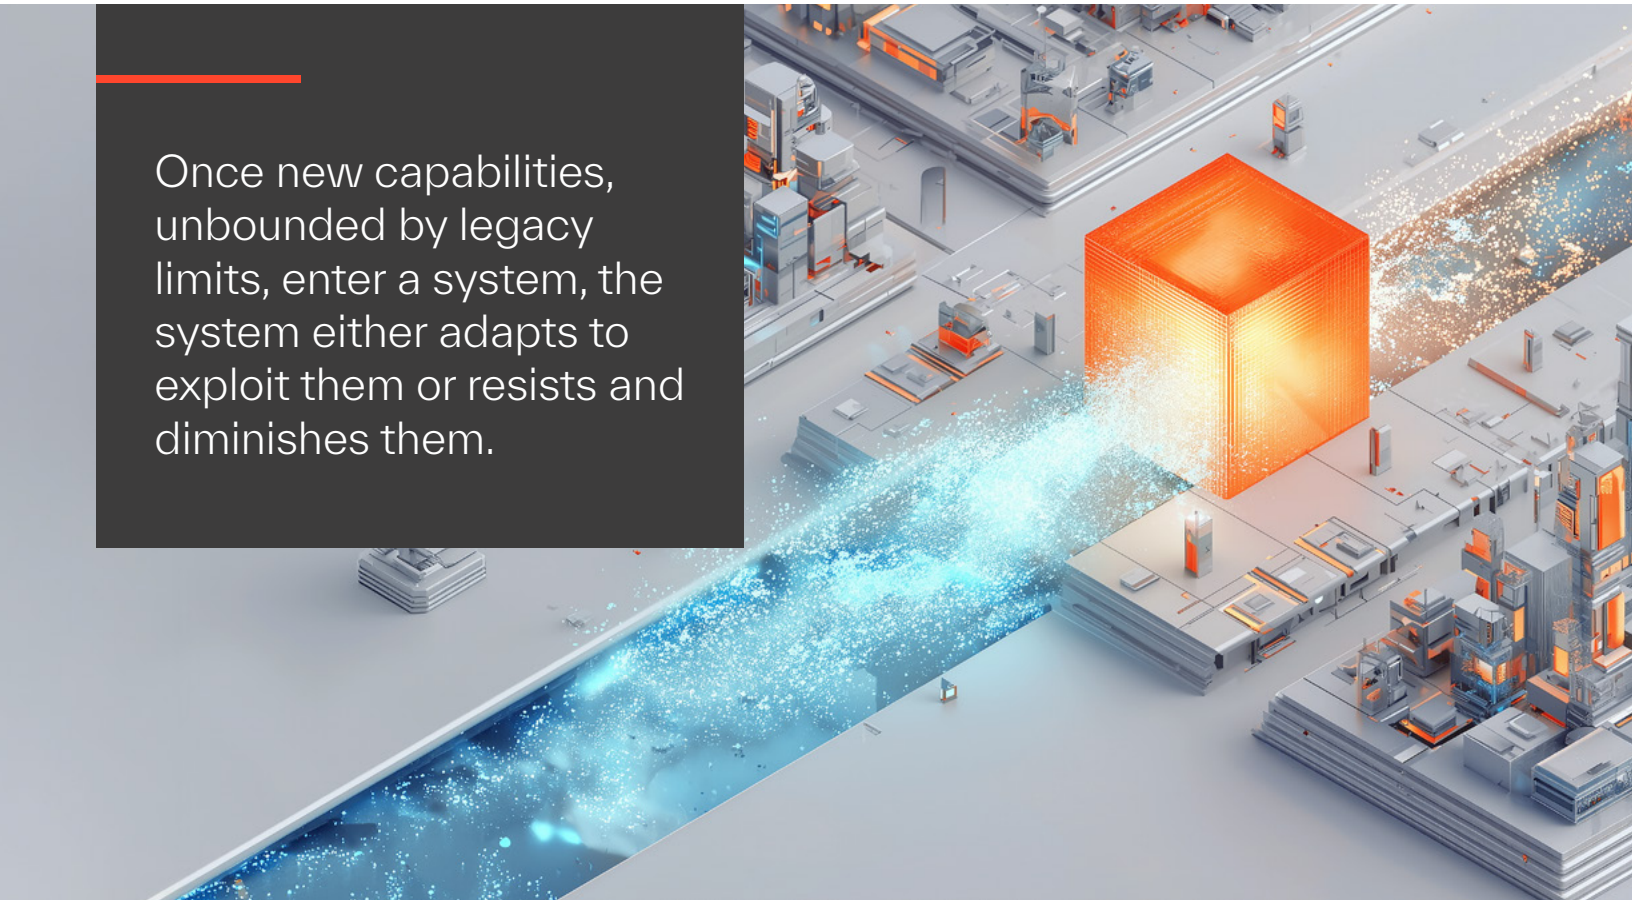

First, improving agentic capabilities constantly 'revalue' human capabilities. When an agent becomes reliable at a task, the economic value of the human capability that formerly performed that task declines. At the same time, the value of complementary human capabilities rises. This dynamic seems similar to the idea of augmentation but is actually the polar opposite. Augmentation assumes that the worker will always be at the center of the workflow and treats AI as a powerful extension of their abilities. But as agentic capabilities improve and force the structure of the work to reorganize around them, without proactive human-agent systems design, workers no longer define the workflow; they occupy the segments where agentic capabilities fail, get bottlenecked, or hit their limits.The effect may look like augmentation from a distance, yet the underlying logic is different: the system expands outward through agentic execution, and humans are continually reallocated to the frontier where reliability breaks and human capabilities like judgment, interpretation, and governance create distinct value.

Consider the case of customer support personnel. As agents become sophisticated enough to handle 80–90% of tickets, the human role no longer augments the agent. Instead, the support representative is pulled into the boundary region where a customer is distressed but not saying why, or where empathy for a particular customer's context requires overriding policy. Human value is no longer about achieving more by leveraging agents but about managing ambiguous situations. Or consider the case of an urban planner. Agents may simulate traffic patterns and model urban density to generate zoning options. But human roles gain more value in mediating between incompatible visions of what the city should be, incorporating moral, cultural, and historical context that would not be adequately factored in by agents.

As human roles transform around agentic capabilities, the organization becomes a moving boundary between what the agentic work system - the evolving system of work structured around evolving agentic capabilities - can do and what humans must still provide. Leading such an organization requires attention to the evolution of that boundary, because it determines where new opportunities emerge and where hidden risks accumulate.

Second, as agentic capabilities evolve, the value of human capabilities evolves with them as they are constantly revalued. Traditional skill models and job architectures assume that capability value changes slowly. They were built for a world in which experience accumulated over years and organizational roles stayed stable. In a system shaped by agentic capabilities, those assumptions collapse. A skill taxonomy can become outdated faster and a job description can lose relevance when a workflow is redefined. As a result, organizations begin to misallocate their people because their frames for evaluating capability are anchored in a past that no longer exists. Consider a case where airline dispatchers are still judged on plan accuracy even though agents handle routing, leaving humans undervalued in the real work of managing disruptions that can have cascading effects. Or consider the case of scientists measured on experiments run even though hypothesis selection and anomaly interpretation, not lab throughput, increasingly drive discovery. Or consider the case of airline pilots. With autopilot technologies accounting for the majority of flight miles, evaluating a pilot based on flight miles no longer makes sense. Evaluating them based on their ability to manage the constraints of safe takeoff and landing as well as disruptions during the flight would be more appropriate.

The problem confronting organizations today is not really one of managing 'digital labor' as much as one

As agentic capabilities evolve, the value of human capabilities evolves with them as they are constantly revalued.

of better managing workers whose capabilities are constantly being revalued as agentic capabilities improve. What compounds this further is that capabilities are also simultaneously revalued because the nature of competition constantly changes as AI makes previously scarce capabilities more widely available. With that, there's an additional external force that determines which capabilities will hold value and enable the firm to compete.

This creates a capability-sensing problem for leaders. Firms often struggle to understand which human capabilities are rising or falling in value, how internal talent should be redeployed, and where new gaps are emerging. They see the symptoms as some roles feel underutilized and some other roles become bottlenecks but lack the structural mechanism to interpret them. The solution is not to manage employees more tightly but to develop better capability sensing mechanisms and better capability allocation.

This brings us to the third and most important point. As the competitive environment values new capabilities and as agentic capabilities themselves evolve and improve, a capability that was peripheral yesterday may become central to the organization's differentiation tomorrow. Leaders need to actively recognize how agentic capabilities are shifting and realign the organization around those shifts. This resembles capital allocation more than workforce management – investing in the capabilities that create leverage, monitoring the ones that are becoming strategically sensitive, and deciding which must remain proprietary because dependence on external providers would compromise long-term autonomy. These strategic issues remain invisible if we only view agentic capabilities as ways to augment the workforce and not as capital investments that determine long-term competitiveness.

This is not incremental change management but organizational re-architecture, requiring new methodologies and leadership capabilities beyond conventional change practice. The focus shifts to designing how humans and agents interact, governing the logic that determines which decisions remain human-led, and investing in the capability base that produces long-term differentiation. We need to stop taking the digital labor metaphor literally to see the questions that truly matter: Which agentic capabilities must be proprietary? How should workflows be restructured around them? What sensing mechanisms reveal how capability values are shifting? How should humans be repositioned as the boundary evolves? These are not questions that can be answered through workforce management, because they concern the structure and economics of the entire system. They belong at the center of strategy.

In this emerging environment, the CEO's task is closer to shaping a continuously adapting organism than to managing a traditional workforce. The architecture changes, the capabilities shift, the boundaries move, and advantage flows to the firms that can reconfigure themselves. Leaders who take a static, literal view of agents as co-workers will inherit the constraints of their old structures. Leaders who view agents as capabilities will design structures capable of evolving. And just as Toyota discovered decades ago, the real edge comes from seeing that distinction early enough to act on it.