Future-State HR Operating Model Enabled by Agentic AI

A Strategic Framework for Organizational Growth

Future-State HR Operating Model Enabled by Agentic AI

Agentic AI is posing a more fundamental question to the enterprise than whether HR should use new technologies. The question now becomes how HR work should be designed given that execution will be distributed across human, agent, shared services, centers of excellence, and business partner capabilities. It's not a technology question; it's an operating question.

Too much of what's being discussed in the current marketplace still treats AI in HR as a productivity play: draft faster, search better, handle cases more effectively, self-service more responsively. These are important applications, but they're well beneath the real question. Agentic AI is important because it fundamentally shifts how work allocation, decision-making, service delivery, and accountability will be structured.

However, once the AI systems are capable of interpreting the context of the situation, reasoning through multi-step work, taking actions, and executing with minimal supervision, the problem shifts. The problem is no longer whether AI can enable HR functions or not. The problem shifts to how HR functions can operate with some aspects of execution being performed by digital actors within the context of live workflows.

This requires a different level of leadership. This requires a level of leadership that can articulate the implications of AI disruption into the operating model of the organization: the value of the operating model and the fact that most organizations are currently undifferentiated in this space.

The future-state HR model is not an automation program

The future state HR operating model facilitated by agentic AI should be thought of in terms of a series of discrete automation projects. It shouldn't be thought of in terms of a chatbot strategy, a digitization strategy, or a series of pilots linked to legacy operations. It should be thought of in terms of a completely redesigned approach.

This redesigned approach should begin with a fundamental, though obvious, insight: The current HR operating model was built in a world where human activity was nearly all execution, where computers stored data, and where automation was limited and rule-based. Agentic AI fundamentally shifts this dynamic, where digital capabilities now exist that can contribute to work, rather than simply supporting work.

This means that the operating model now needs to address a series of fundamentally different questions: What work should be performed by agents, and what work should be performed only by humans? Where should agents be allowed to initiate work, but not be allowed to finish work without human intervention? Where should authority be required to return from agents to human? How should exceptions be addressed? How should policy be translated into operational capabilities? How should accountability be maintained?

Those aren't secondary design questions. They're the model.

Curtiss J. Shupe

Future-State HR Operating Model Enabled by Agentic AI

Organizations that view agentic AI as an add-on to their existing HR structure will create islands of efficiency and islands of risk. The organizations that view agentic AI as a redesign of service delivery, control, and accountability will create a function that is more scalable and more strategically useful. That's the difference between local optimization and operating model transformation.

The structure may remain familiar, but the roles don't

The vast majority of enterprise HR functions will still be delivered through some form of shared services, centers of excellence, and business partners. Agentic AI doesn't replace this model. It simply changes what each part of this model is accountable for and how they interact.

The future-state model can be thought of as a human agent service network. In this model, agents manage structured execution, workflow progression, triage, monitoring, drafting, and orchestration. In contrast, humans manage judgment, empathy, ambiguity, and decisions, while a control layer between them determines policy logic, boundaries, escalation, data permissions, and supervisory needs.

This control layer is what differentiates a scalable operating model from a series of experiments. Without it, digital capabilities are simply inserted into workflows without authority boundaries, decisions start to become inconsistent, and risk management reverts to humans who were never given a proper design in the first place. The future-state model is less about whether agents can perform tasks and more about whether an organization is disciplined enough to determine how they can be permitted to perform them.

Self-service becomes governed execution, not portal deflection

One of the most obvious changes is the change at the front door of HR. Self-service systems are often built around forms, knowledge articles, ticket queues, and a generally fragmented service model. They're often designed to avoid the actual service rather than deliver it.

In the agentic model of service, the front layer of the service model is more active and can answer questions, interpret intent, gather context, walk the customer through the process, collect the required inputs, initiate the approved actions, and route the cases based on risk, complexity, and policy. This is a significant change in the service model because it fundamentally changes the line between self-service and service execution.

The idea isn't to simply create a more conversational interface to the same fragmented service model. The idea is to rethink the front layer of the service model so it can execute low-risk actions reliably and reduce friction on the employee and the manager. However, this can only happen if the service model is designed with a clear boundary of authority.

Curtiss J. Shupe

The front layer of the service model requires a set of permissions: what can it execute on its own, what requires confirmation, what requires review, and what always requires escalation? Without these boundaries of permission, the service model becomes fast but loses coherence.

The right question isn't whether a conversational agent can do more. The right question is what it should be allowed to do as part of a governed HR service model.

Shared services become execution-and-exception engines

Shared services will be different as well. Shared services are currently designed to optimize transactional efficiency, standardization, and transaction volume handling. In an agentic environment, these goals are the same, but the approach will change from human throughput to workflow orchestration.

Agents can handle tasks such as intake, classification, routing, document preparation, status tracking, process guidance, and moving cases in a structured manner. Humans can then be deployed in areas where they can create the most value, i.e., exceptions, ambiguity, employee issues, judgment, and control.

This is when the concept of operating model maturity becomes apparent. A less mature operating model will simply insert an AI component into the current workflow, hoping to gain labor efficiencies. A more mature operating model will change the service paths, thresholds, and supervisory structures so that the benefits are derived from better orchestration rather than uncontrolled delegation.

In the end, the benefits will be derived from better orchestration rather than uncontrolled delegation. That's the difference between automation and architecture.

Centers of excellence become design authorities

The importance of centers of excellence, therefore, increases rather than decreases. In a traditional HR model, centers of excellence often own policy, program, framework, and subject-matter expertise. However, in a future-state model, enabled by agentic AI, this has to be more operational and more architectural.

Centers of excellence increasingly become design authorities, defining policy logic, decision rules, taxonomy standards, risk tiers, exception criteria, skills architectures, evaluation approaches, and conditions under which humans and agents are allowed to act. This is no longer a question of publishing guidance, running programs, and managing centers of excellence; this is a question of encoding how the system should operate.

Curtiss J. Shupe

The importance of this shift lies in the fact that successful execution at scale is directly related to successful design. If policy remains ambiguous, taxonomy remains inconsistent, escalation remains weak, and process definitions remain inconsistent across domains, agents will simply accelerate those problems. The organization will operate faster, but it will operate faster in the wrong direction.

Thus, a mature future-state model increases the strategic importance of centers of excellence. This isn't a question of service ownership; this is a question of those groups defining the conditions under which scalable, governed execution can occur.

Business partners move up the value chain

The HRBP role also needs to be reset. In far too many organizations, business partners still act like transactional brokers, policy interpreters, and process navigators. They still invest time and energy in connecting different parts of the organization because the underlying operating model is still broken.

This isn't an operating model fit for an agentic environment. If the operating model at lower levels is more reliable and more effective at structured execution, then business partners should be moving more rapidly up the value chain. Their activity should be more strategic, more diagnostic, and more connected to decisions that have more impact on the organization: workforce mix, role creation, leadership implications, shifts in capabilities, operating model trade-offs, and human implications of work redesign enabled by AI.

This is where disproportionate value can be added. The business partner of the future isn't a manual connector who tries to hold together an inefficient HR operating model. The business partner of the future is a human advisor who helps leaders make good decisions about how work, talent, accountabilities, and organization need to change.

This won't happen by accident. It'll only happen if the rest of the operating model is redesigned effectively enough to remove transactional drag from this activity. If it isn't, then the business partner still remains stuck in a brokerage activity rather than moving up the value chain to more strategic value creation.

Work design becomes the central HR discipline

The most profound implication of agentic AI, therefore, is that work design takes center stage in HR thinking. The relevant unit of analysis shifts from job description alone to task, handoff, decision, control, and accountability.

This is where AI fundamentally shifts the economics and design of work. It's also where most traditional HR thinking has been least developed. A viable future state HR model, therefore, begins with work design, rather than only thinking about jobs at a high level.

Curtiss J. Shupe

Which tasks are highly structured and rules-driven? Which tasks are repetitive, requiring periodic review? Which tasks involve sensitive employee implications, requiring contextual interpretation, negotiation, or empathy? Which tasks should be initiated by AI, requiring human sign-off? Which tasks should remain human only because of their legal, ethical, relational, or strategic implications? Which tasks are simply not ready for agent involvement because of data, governance, or workflow maturity issues?

These distinctions matter because they ground AI's abstractions in concrete operating model choices. They drive the function to think about how work should actually be redesigned. They also drive accountability because, with work design, one can now actually assign proper authority, supervision, and performance metrics.

Without work design, most AI-related work in HR will be superficial. It'll be productivity-enhancing, perhaps, but won't fundamentally redesign how the function operates. Hence, work design, therefore, isn't an adjacent capability in this model; it's the core discipline.

Governance is not adjacent to the model; it is the model

This is where many organizations still get it wrong. Governance is often considered an adjacent workstream to address issues like legal checks, ethics reviews, technology policy, data standards, or risk signoff. While all of these are important functions, they aren't enough on their own.

Within an HR operating model facilitated by agentic AI, governance must be an integral part of how the operating model executes. The key governance issues here aren't abstract risk-based considerations. They're more like: Who or what can act? Upon what data? To what degree of autonomy? At what thresholds? Who reviews the result? What data is captured? How do exceptions get surfaced? At what point is human authority re-engaged? How do errors, changes, biases, privacy issues, or control issues get detected over time? These aren't abstract risk considerations. These are considerations about whether or not the model can be trusted.

Any employment-based application of AI is, by its nature, a high-stakes activity because it's close to issues of opportunity, mobility, judgment, access, and working conditions. That means the standard for operating discipline needs to be higher. It needs a governance architecture for how it runs, including permissions, approvals, audit trails, policy traces, and ongoing monitoring of its performance and its harms.

That's why governance can't be considered an adjacent process to how an HR operating model runs. Governance is the operating model. It's how the distributed execution of an HR operating model remains legitimate, reviewable, and safe.

Curtiss J. Shupe

Data, knowledge, and identity become operating infrastructure

The strength of agentic HR will be only as good as the infrastructure supporting it. If there's inconsistency in policies, fragmentation in knowledge, poor data definitions, issues with identity and control, or no evidence being captured, then agentic execution will simply enhance those problems rather than solve them. So, the future state model relies upon many more things than just workflow redesign.

The first component required is a knowledge layer, including policies, procedural logic, role definitions, decision standards, domain guidance, and more, which can be accessed consistently in human and digital workflows. The second component required is data, including clean, standardized, and interoperable data across the entire HR domain and adjacent enterprise applications. The third component required is an identity and control layer, including clear human and agentic permissions, and evidence of what was done, by whom, under what authority, and what was the result.

While typically considered a technical enabler, I'd argue that all three components should be considered part of an operating model, driving whether or not service delivery can be trusted, whether or not accountability can be defended, and whether or not the model will improve over time based upon evidence rather than assumption. Without these, there isn't an agentic HR model; there's just a weak system with unknown failure points.

Digital labor needs explicit management

As agents become participants in live processes, there's a need for a more explicit management model of digital labor as well. Agents can't continue to be invisible utilities locked within tools, whose ownership is unclear. If they're participating in human resource processes, they need scope, ownership, permission, oversight, lifecycle control, performance, and review.

In other words, digital labor must be managed as an operating asset rather than a feature of an operating asset. Organizations need disciplined answers to the following simple questions: what is the agent authorized to do? What processes does the agent participate in? Who is responsible for the agent? What are the implications of agent performance degradation or policy change? What are the fallback options? And how do you retire, replace, or constrain the agent when performance degrades?

These are simple management questions, but many organizations don't have disciplined answers to these questions. This'll become more apparent as agentic systems move from pilot to operating execution. A proper operating model can't leave these questions in an ambiguous state.

Workforce planning becomes mixed-labor planning

Workforce planning too would undergo a significant transformation in this environment. The traditional workforce planning discipline includes dimensions like headcount, requirement for roles, spans and layers, hiring needs, and labor costs. These dimensions still apply, but they'd no longer be relevant for a function operating in a mixed human agent environment.

The future-state workforce planning discipline would have to address a mixed labor model comprising people, contingent talent, automation, and agents. That'd mean a more dynamic workforce planning discipline where business needs would be linked to work decomposition, tasks would be linked to skills and control needs, and would determine effective combinations of human and machine labor. It would also have to determine what new skills would be required for people who would be supervising, working alongside, or intervening in agent-based work.

This is a significantly different workforce planning discipline than what we've traditionally done. It's more like work allocation within an enterprise. It'd require HR to think not just about labor supply, but about how execution is designed.

HR capability requirements rise

However, a future-state HR operating model, enabled through agentic AI, also implies a stronger HR function. This doesn't mean a function in which every HR leader also has to be a technologist, but a function that has much greater fluency in terms of data, architecture, service, instrumentation, governance, and organizational decision-making processes.

The bar for HR function capability has to be higher, and this is partly because the work itself has become more integrated, more technical, and more critical in its implications and expressions.

We also need an HR function that can take a strategy and turn it into a work design, a policy and turn it into an operational logic, an AI capability and turn it into a bounded authority, and a change and turn it into a scalable execution, rather than a communications exercise in its own right. This is a much more significant shift in terms of capability, and one that has much more to do with what sets apart organizations that use AI from those that redesign the HR function around its use. The line here has less to do with tools and everything to do with the maturity and fluency of the HR function that must redesign work around those tools.

Curtiss J. Shupe

The real measure of success

The future state of the HR operating model must not be evaluated solely in terms of efficiency. While speed and cost are clearly important, they aren't the only considerations. A proper scorecard must include service quality, exception rates, escalation, policy, workforce, auditability, manager trust, employee experience, and the quality of decisions in critical moments.

This is the appropriate approach because the proper focus isn't simply to speed up the processing of HR-related work. The proper focus is to develop an HR function that is more scalable, more precise, more governable, and more strategically valuable in an environment in which the nature of work is becoming more distributed among human and artificial actors. This is the true promise of agentic AI in the HR domain.

Not generic productivity gains. Not isolated automation wins. A materially stronger operating model.

What leadership looks like now

The most potent leadership position in this environment isn't generic enthusiasm about AI. The most potent leadership position is operational clarity: the ability to turn the disruptions caused by AI into concrete decisions about the operation of the enterprise, including the operation of the function.

It's the ability to move past the experimental phase and make concrete decisions about the operation of the function when humans and agents interact within the enterprise. This is the challenge facing the HR function. This is the opportunity.

The winners won't be the enterprises that experimented most successfully with the most tools. The winners will be the ones who make concrete decisions about the operation of the function at the level of performance creation: tasks, decisions, handoffs, authorities, controls, evidence, etc. This is where the future state of the HR operating model will be determined.

Curtiss J. Shupe