

# The Enterprise Needs an Execution Control Plane

Curtiss J. Shupe

©2026 Shupe.io

# The Enterprise Needs an Execution Control Plane

## *How workforce operating systems coordinate workflows, agents, governance, telemetry, and intervention at scale.*

Most organizations are still trying to manage human-AI collaboration using a somewhat random assortment of technology. HR management platforms manage roles. Workflow management platforms manage tasks. Enablement platforms manage learning. Analytics platforms provide reporting. Policy management platforms manage governance. The AI platforms reside at the top, or off to the side, as assistants, copilots, agents, and feature implementations. This worked okay when the technology just supported work at the edges. But it doesn't work well when the AI starts to participate in execution.

The direction that is being taken can be seen clearly by looking at the 2025 Work Trend Index from Microsoft. In it, they outline the concept of the "Frontier Firm," which is based on intelligence on demand, human agent teams, and managers who are becoming responsible for digital labor. They also outline that their research is based on survey data from 31,000 people across 31 countries, LinkedIn's labor trends, and Microsoft 365 productivity signals. This is significant because it shifts the discussion away from specific tools. The reality is, the enterprise is not simply using AI; it is using it with human contributors within the workflow.

That is precisely what is happening, and hence the need for the concept of an enterprise workforce operating system, which is what this article is proposing. It is a term that is used as a synthesis term in this series, not as a term for a specific product. However, it is a term that reflects a reality supported by all the evidence in this article and the others in the series. Once the workforce is distributed among humans, agents, automation, decisions, telemetry, and governance, there needs to be a way to bring all these together in one operating system. This is supported by McKinsey's article in 2025 on the agentic organization. They state, among other things, that what is happening is that organizations are moving to a new paradigm in which humans and virtual and physical AI agents cooperate to produce value, and in which operating model, governance, and workforce are not secondary considerations.

In this series, this coordinating function is named an execution control plane. This is also a synthesis term. The external sources don't use it exactly. However, the logic is sound and getting harder to escape. The enterprise needs a coordinating function to translate policy into runtime, route tasks to human and machine contributors, consume telemetry, support intervention, and drive continuous redesign. The National Institute of Standards and Technology (NIST) AI Risk Management Framework is relevant to the underlying need. It's organized around Govern, Map, Measure, and Manage for AI risk management, and GOVERN is a cross-cutting discipline. This isn't static approval logic. It's operating model logic.

# The Enterprise Needs an Execution Control Plane

A good working definition of an enterprise workforce OS has six parts, and they're more specific than they sound. First, there's workflow orchestration. Work has to move from person to agent to automation to decisions through explicit routing logic. If you don't have good workflow orchestration, AI makes some parts of the work go fast, while the overall system is slow, brittle, and opaque. The second part is policy translation. Policy is not really operational unless it's been translated into a form that the runtime system can actually execute. This is one of the largest blind spots in enterprise AI today. Policy is in place, but the workflow can't execute it.

Third is telemetry interpretation. Article 8 established that an enterprise requires visibility into activity, behavior, interaction with humans, quality of decisions, and business outcomes. There is an operating layer that requires ingesting those signals and converting them into actions, which can be escalations, overrides, reallocations, or refinements. Without it, it is reporting, not control.

Fourth is runtime supervision. If agents are operating within the workflow, then an enterprise requires a structure by which humans can observe, allow, or stop actions, and correct actions prior to errors. This is evidenced by both Microsoft and McKinsey, albeit indirectly, by describing an organization that is redefining the work process with human-agent teams and management.

Fifth is guidance delivery. Policies, playbooks, role guidance, and decision rules need to be present at the point of work. Static documentation, disconnected from runtime execution, is a weak form of operating. It forces both people and systems to make up their own rules. Sixth is continuous redesign. ISO/IEC 42001 is particularly relevant to this point, as ISO says, "ISO/IEC 42001 specifies the requirements for a management system to support an organization to achieve its objectives. It also specifies requirements to establish, implement, maintain, and improve an AI management system. The management system concept is particularly relevant to workforce execution. The operating layer is not static. It needs to learn from telemetry, incidents, friction, and results, and then change how it operates."

This is where many organizations will tend to underinvest. They will invest in their models, their copilots, their agents, their workflow tools, their dashboards, and many other items. However, without a coordinating operating layer, they will continue to struggle with many of their traditional enterprise problems: unclear handoffs, inconsistent thresholds, fragmented telemetry, redundant control paths, and accountability. These aren't quoted findings from any one source. They're the operating inference based on the overall direction of the Microsoft, McKinsey, NIST, and ISO findings. In other words, enterprises are adding capability much faster than they're building operating discipline to run that capability in a coherent fashion.

This also explains why workforce enablement changes character in a human-AI environment. In a legacy model, enablement could sit beside execution as training

# The Enterprise Needs an Execution Control Plane

programs and knowledge repositories. In a hybrid human-agent environment, enablement has to be integrated into execution itself: guidance at the point of work, certification where needed, intervention rights where risk is elevated, and redesign where telemetry shows persistent friction. That's the real convergence point between workforce design and operating architecture. It isn't enough to teach people how to use AI tools. The enterprise has to teach, control, and redesign the system through which human and AI work now run together.

That leads to the executive question that actually matters. Not just which AI tools the organization is deploying, but what operating layer coordinates workflows, agents, decisions, telemetry, and governance across the hybrid workforce. Organizations that can answer that question will have more than adoption. They'll have a system that can route work dynamically, enforce autonomy boundaries, supervise agents during runtime, interpret telemetry, and improve continuously. That's what this article means by an enterprise workforce operating system. It's the difference between adding AI to the stack and building an enterprise that can actually operate with it.