

EVERYONE IS FOCUSED ON AUTONOMOUS EXECUTION

The Real Risk is Whether It Actually Works

SAP Sapphire 2026 made one thing clear: enterprise systems are evolving beyond traditional systems of record into platforms designed to execute business processes autonomously. Across keynotes, analyst discussions, and partner conversations, the direction was consistent, AI is being embedded directly into workflows, interfaces like Joule are enabling intent-driven interaction, and systems are increasingly expected to not just support decisions but carry them out.

When that execution doesn't behave as expected, the impact doesn't stay technical. It shows up in operations, compliance, and financial outcomes.

In most environments, that risk is already there, it just hasn't shown up yet.

While this shift represents a meaningful advancement in enterprise technology, it also introduces a critical assumption that is often overlooked: that the underlying environment can support consistent, reliable execution under real-world conditions. In practice, this assumption does not hold true in most SAP landscapes.

The transition to autonomous execution fundamentally changes how enterprise systems are expected to behave. Historically, systems served as tools operated by users, with human intervention acting as a buffer for inconsistency, error, or misalignment. In the new model, users define intent, and systems are responsible for executing outcomes. This places significantly greater pressure on the accuracy, consistency, and predictability of the environment itself.

At the same time, AI is no longer being treated as an experimental capability. Organizations are actively working to operationalize AI across core business processes, tying it directly to measurable outcomes in finance, supply chain, and other critical functions. However, embedding AI into execution introduces a new dynamic: when issues occur, they do not remain isolated. Instead, they scale, and they do so without the same level of visibility or intervention. What might previously have been a contained exception or manual workaround can become a repeated, automated outcome with broader operational impact.

The introduction of Joule further reinforces this shift. By abstracting system navigation and allowing users to interact through intent rather than transactions, it significantly reduces friction. However, it also changes the system's role, from supporting decisions to making them. As a result, execution accuracy becomes non-negotiable. When systems act autonomously, any misalignment between design, configuration, and real-world operations becomes more consequential, more immediate, and more difficult to detect.

AI is only as reliable as the environment it operates in.

AI-driven execution depends entirely on the quality and consistency of the underlying environment. In many organizations, that environment is still characterized by fragmented data, evolving processes, and controls that drift over time. While the vision of autonomous execution is compelling, the reality is that most environments still rely on manual intervention when execution doesn't behave as expected.

The commonly referenced “enterprise AI gap” is often framed as a technical challenge, legacy systems, lack of cloud adoption, or integration complexity. In practice, the gap is more operational. Most SAP environments were not designed to execute predictably at scale. As organizations move toward embedding AI into core processes, this limitation becomes more visible and more impactful.

Sapphire also reinforced the importance of cloud adoption and clean core strategies as prerequisites for enabling AI at scale. While these are necessary steps, they are not sufficient on their own. Modernization does not eliminate risk; it changes where and how that risk manifests. In cloud-based and highly integrated environments, ownership can become less clear, dependencies less visible, and control validation more complex. As a result, systems may be technically managed, but not always fully understood.

SAP’s focus on industry-specific AI reflects an important recognition that context matters. Generic AI models often struggle in complex enterprise environments, and domain-specific logic is essential for achieving meaningful outcomes. However, even with more advanced models, the fundamental challenge remains the same: ensuring that systems behave correctly under real operating conditions. This is ultimately where success or failure is determined.

For enterprise leaders, this shift changes the nature of the problem.

The question is no longer simply whether the organization is adopting AI or modernizing its systems. Instead, it becomes whether the environment can support consistent execution when it matters most. If it cannot, issues do not remain isolated, they propagate. Control gaps do not stay contained, they scale. What appears to be increased efficiency can quickly translate into operational risk.

The most significant risk in this transition is not whether organizations adopt new technologies, but what becomes embedded into the environment during transformation. Design assumptions that are not fully validated, execution shortcuts taken under time pressure, and controls that are not fully aligned with evolving operations can all become part of the system itself. These issues often remain invisible during implementation and early operation.

They surface later, when the system is relied upon at scale.

SAP Sapphire 2026 clearly signals the direction of enterprise technology: toward autonomous execution and AI-driven operations. However, the success of that vision depends not just on the sophistication of the technology, but on the reliability of the environment in which it operates.

Autonomy does not eliminate risk. It amplifies it when the underlying system is not prepared.

Ultimately, the challenge is not reaching autonomous execution. It is ensuring that, when organizations do, their systems behave predictably, their controls hold under real conditions, and their outcomes align with business expectations.

Otherwise, what appears to be innovation at the surface becomes risk embedded into the business, often only visible once it’s already impacting outcomes. And by the time it does, the impact is already operational.