# Governance First: Why SAP's Autonomous Enterprise Vision Is the Right Move for Utilities

*A Utility.Community Perspective on SAP Sapphire 2026 and the Path to Production AI*

## The Table Stakes That Went Missing

For those of us who have spent careers in regulated utility environments, the concepts of security, governance, compliance, and performance are not aspirational — they are prerequisites. These were the table stakes. The assumption was always that nothing moved to production until those boxes were checked.

Then came the exuberance of AI.

Across our industry, teams have built impressive proof-of-concepts. Predictive maintenance models, DER dispatch optimization tools, outage analytics dashboards, intelligent metering pipelines. The demos were compelling. The business cases were sound. And then, one by one, many of these projects stalled — not because the technology failed, but because when the questions of security, performance, and governance were finally raised in earnest, the answers were not ready.

> ***We had the innovation. We had forgotten the infrastructure that makes innovation safe.***

This is not a failure of ambition. It is a natural consequence of moving fast in an industry where we are not accustomed to moving fast. And it is precisely why what SAP announced at Sapphire 2026 matters — not as a limitation, but as a foundation.

## What SAP Actually Announced

At Sapphire 2026, SAP unveiled what it calls the Autonomous Enterprise — a unified platform built around three pillars: the SAP Business AI Platform (context, build, and governance layers), the SAP Autonomous Suite (50+ domain-specific Joule Assistants orchestrating 200+ specialized agents), and a new user experience designed for human-AI collaboration.

For utilities specifically, SAP and Anthropic announced a partnership to build custom agents and agentic workflows targeting the utilities sector — covering asset management, field operations, and the integration of operational and transactional data that has long been our holy grail.

Preceding Sapphire, SAP also updated its API Policy (v4/2026), restricting the use of SAP APIs for unsanctioned agentic AI integrations — specifically those involving autonomous systems that execute sequences of API calls, and large-scale data extraction outside endorsed architectures. This policy update has generated significant community debate.

We believe that debate, while valid, is missing the larger point.

## Why the API Policy Is the Right Move — Even If the Communication Was Not

Let us be honest about something: the AI integrations many of us have prototyped involve third-party tools calling SAP directly — often via undocumented interfaces, often at scale, often without formal governance frameworks. Some of these have been enormously valuable as demonstrations. Very few have been ready for production.

The reason they were not ready for production is exactly what SAP's new policy and Autonomous Enterprise framework are designed to address.

### Security

Agentic AI systems that can read — and increasingly write — to production systems represent a categorically different security surface than traditional integrations. A human makes one decision at a time. An AI agent can execute thousands of actions in the time it takes a security team to be notified. SAP's Agent Gateway and endorsed architecture model create a controlled, auditable channel for agent activity. This is not lock-in. This is the security perimeter that every utility CISO has been asking for before approving any agentic AI in production.

### Governance and Compliance

We operate in one of the most regulated industries in the world. NERC CIP, state PUC requirements, FERC oversight, and an evolving landscape of data privacy and cybersecurity mandates all touch our systems. SAP's governance layer — built on LeanIX, Signavio, and Cloud ALM — provides auditable agent telemetry, verified agent enforcement, and workforce impact mapping. These are not nice-to-haves for utilities. They are prerequisites for any regulator conversation about autonomous AI in production.

### Performance

Our production systems — IS-U billing, outage management, Plant Maintenance, PSCD — were designed for human-paced transactional workloads. Agentic AI creates a fundamentally different load profile: high-frequency, parallel, potentially recursive API calls that can destabilize systems not designed for them. SAP's concern here is legitimate and practical. The endorsed architecture model is as much about system stability as it is about commercial interests.

### Audit

When an AI agent acts on production data — closes a work order, adjusts a rate schedule, dispatches a field crew — who is accountable? Under current third-party integration models, the answer is often unclear. SAP's framework, with its agent activity logs, policy enforcement layers, and company memory audit trails, creates the accountability chain that regulators will eventually demand and that our internal audit functions need today.

## The Vendor Lock-In Question Deserves an Honest Answer

The concern about vendor lock-in is real and should not be dismissed. But it deserves a more complete analysis than it has received in the community conversation so far.

> ***Every AI vendor is building a proprietary moat. The question is not whether lock-in exists — it is whether you are choosing your lock-in deliberately.***

Consider what is happening across the AI landscape right now. Microsoft Copilot captures your employees' interaction patterns, document workflows, and decision logic in its memory layer. Google's Gemini Enterprise builds organizational knowledge graphs from your workspace activity. Anthropic's Claude, used at scale, accumulates company-specific context through memory systems that are not portable across vendors. While open standards like MCP (Model Context Protocol) and A2A (Agent-to-Agent) protocols enable interoperability at the integration layer, the company memory — the captured tribal knowledge, the learned preferences, the institutional process context that makes an AI genuinely useful — is proprietary to each vendor.

This is not a criticism of any of these companies. It is the nature of how value is created in AI systems. The more your people interact with an AI platform, the more contextually valuable it becomes to your organization — and the more embedded it becomes. SAP is doing exactly what every other AI vendor is doing. The difference is that SAP is doing it with 50 years of utility industry process knowledge already embedded in the platform.

The question for our community is not "How do we avoid lock-in?" — because that ship has sailed industry-wide. The question is: "With which partner do we want to build our institutional AI knowledge, and do they have the industry depth, governance framework, and regulatory awareness that our business requires?"

That is a question worth asking seriously. And for many utilities already deeply invested in SAP, the answer may well be that SAP's Autonomous Enterprise — with its utility-specific agent roadmap, its Anthropic partnership targeting the utilities vertical, and its governance framework — is the most defensible path to production AI.

## The Data Architecture Shift We Cannot Ignore

For many utilities, we have spent a decade building robust analytical architectures: replicating SAP data into BW, data marts, and data lakes, then running read-only dashboards and decision-support tools on top. This model has served us well. It is also insufficient for the agentic AI era.

The fundamental difference between traditional AI — which recommends — and agentic AI — which acts — is write access to production systems. An AI agent that can update a maintenance work order, adjust a demand response signal, or modify a billing account is not operating on a read-only copy of your data. It is operating on your production system. That requires an entirely different level of security, governance, and performance infrastructure than we have historically built around our analytics layers.

SAP's Autonomous Enterprise framework, with its Agent Gateway and endorsed architecture model, is the first credible industry-specific answer to the question: "How do we give AI agents safe, governed, auditable write access to our production utility systems?" We should engage with that answer seriously — even as we push to improve it.

## What We Still Need to Influence

Embracing SAP's strategic direction does not mean accepting it without question. As a community, we have both the standing and the responsibility to shape how this vision is implemented. Here are the areas where our voice matters most:

- Real-time operational data integration: Our DER management, grid edge analytics, and demand response programs require sub-second data from SCADA, AMI, and OMS systems that live outside SAP's transaction perimeter. We need SAP's agent framework to connect cleanly with real-time operational data, not just SAP transactional data.
- Pricing transparency and contractual protections: The DSAG user group has formally demanded that SAP provide clear contractual definitions, transition timelines, and protection for existing integrations. Our community should align with and amplify these demands, particularly around contract renewal implications.
- Existing AI investment protection: Many of our members have built valuable AI pipelines and data products that currently access SAP via interfaces that may be affected by the new API policy. We need grace periods, migration support, and clear endorsed-architecture pathways that protect these investments.
- Asset management with embedded analytics: Predictive maintenance, failure probability modeling, and maintenance plan optimization that connects SAP PM data with operational sensor data is one of our highest-value use cases. We need SAP's utility agent roadmap to prioritize this.
- 2027 pricing clarity: The Joule agent runtime is free through December 2026. No post-promotion pricing has been disclosed. We need commitments, not promotions, before we build production AI infrastructure on this platform.

## A Call to Our Community

Before dismissing SAP's API changes and Autonomous Enterprise vision as unwelcome constraints, I ask each of you to sit with a harder question: If not SAP's governance framework, then what? How are you planning to operationalize agentic AI — with true write access to your production utility systems — in a way that satisfies your security team, your compliance function, your internal audit, and your regulator?

If you have a better answer, please bring it forward. Our community is stronger when good ideas compete. But "we'll figure out governance later" is not a plan. It is the reason our proof-of-concepts have not become production systems.

> ***AI is not going away. Neither are our regulatory obligations. The goal is not to choose between them — it is to build a framework where both can coexist.***

SAP has put forward a strategic vision for doing exactly that. It is imperfect, it raises legitimate concerns, and it will require significant community engagement to shape into something that fully serves our needs. But the direction — governance, security, and performance as the foundation for production AI, not as afterthoughts — is correct.

Let's read the details. Let's ask the hard questions. Let's engage SAP through ASUG and our direct relationships to influence the roadmap. And let's do it with the recognition that the hardest part of the AI journey for regulated industries was never the innovation — it was always the operationalization.

That challenge is now in front of us. Let’s meet it together.

---



Published May 2026 | Utility.Community | utility.community