delaware

Why SAP and Anthropic Are Converging on the Most Important Idea in Enterprise AI

we commit. we deliver.

Table of contents

Value is a kept promise. The question AI must answer is not how to accelerate work, but how to govern outcomes.

Executive argument

SAP has moved closer to the value stream than any previous ERP architecture allowed. Anthropic has previewed a form of agent memory that points beyond retrieval toward consolidation. Separately, these are important developments. Together, they suggest the emerging conditions for coherent enterprise AI: systems that can observe value flow, retain context, consolidate experience and govern toward the commitments organisations make to their customers.

This is not yet a finished architecture. It is a direction of travel. A direction that will not be more automation in more places, but intelligence designed around value, flow and outcomes.

The Moment

Within a few days of each other in May 2026, two significant announcements landed in the enterprise technology world. SAP unveiled the Autonomous Enterprise, its most ambitious repositioning in a generation. Anthropic previewed Dreams in the Claude Console, a research-stage capability that allows AI agents to distil their session transcripts into persistent memory stores that can be carried forward into future runs.

Most observers treated them as separate stories. One came from the world of enterprise ERP. The other came from the frontier of AI research. But they are not separate. They point toward the same architectural problem from different directions.

SAP is moving closer to the flow of enterprise value. Anthropic is moving toward the consolidation of agent experience. Together, they suggest something that has been architecturally impossible until now: the ability to govern the value stream with intelligence over time.

Not merely to monitor it. Not simply to optimise fragments of it. But to govern it in service of the promises an organisation makes and must keep.

This paper explains why that matters, what it requires and why the framework to understand it was already set out in From Efficiency to Intelligence before either announcement existed.

What SAP Actually Signaled

Much of the commentary on SAP's Autonomous Enterprise announcement has read it as a story about automation: faster agents, smarter workflows, fewer human touches and more intelligence in the transactional layer. That reading is not wrong, but it is incomplete.

The more important signal is that SAP has placed a foot in the door of the value stream. It has not yet fully entered the space of governed flow, but it has crossed an important architectural threshold: from recording enterprise activity to participating more directly in how value moves through the organisation.

That matters. SAP's historical position has been as a system of record, the place where transactions land, are validated and become enterprise truth. The Autonomous Enterprise points to a different role; a system that can coordinate action across procurement, finance, supply chain, fulfilment, service and delivery.

In other words, SAP is no longer only concerned with where work is recorded. It is moving toward where value is created, constrained, protected and lost.

But placing a foot in the door of the value stream is not the same as governing it and the distinction matters.

SAP's agent architecture can coordinate workflows across the enterprise. It can trigger actions, route decisions and connect processes that previously required human intervention. What it cannot yet fully do is develop the judgement that coherence requires, to hold the whole value stream in mind over time, anticipate where the promise is at risk and intervene with intelligence rather than simply with speed.

This is not a criticism of SAP's architecture. It is a precise description of the journey SAP is on, the gap that remains and why that gap is strategically important.

What Coherence Actually Means

The word governance carries baggage in enterprise settings. It can imply oversight, compliance, audit trails, policy boards and the bureaucratic machinery that surrounds decision-making in large organisations. That is not what this paper means by coherence.

Coherence is closer to the engineering meaning of governance. The original mechanical governor on Watt's steam engine did not administer the engine. It regulated it. It observed the system continuously, detected deviation from the desired state and corrected it in real time.

A governor does not do the work. It regulates the system that does the work.

Coherence is what emerges when AI operates at that level across the enterprise value stream. Not fragmented automation of individual tasks or orchestration of predefined workflows alone, but active, intelligent regulation of the flow of value, preventing leakage, resolving constraint and protecting the organisation's ability to keep its commitments.

This is what the quality movement were always pointing toward. They understood that local optimisation can destroy systemic performance, and that variation is felt by the customer before it appears on the balance sheet. They understood that value does not sit inside isolated processes, but within the commitment that those processes exist to fulfil.

What they could not do was close the loop in real time at enterprise scale. The technology simply didn't exist. Until now, it was possible to describe the value stream, measure parts of it and improve sections of it, but it wasn't possible to govern it continuously with a persistent intelligence layer.

The Four-Layer Framework

From Efficiency to Intelligence proposed a four-layer framework for understanding where AI creates value in organisations. The framework was written before the SAP and Anthropic announcements. Those announcements now make the framework more relevant, not less.

The framework is best understood from the top down.

The top layer is the Coherence Layer. This is where systemic value emerges, not from accelerating activity, but from stabilising behaviour where outcomes are decided. It requires an AI that holds the organisation in mind over time; one that does not simply act, but consolidates experience, reasons across the whole system and develops the judgement that governance requires.

Beneath this sits the three layers where most organisations currently operate and where most vendor platforms, including SAP's, are becoming genuinely strong. The third layer is the Workflow Orchestrator. This is where AI is coordinating activity across processes. The second is Embedded Process AI within specific workflows. The first is the Personal Assistant, supporting individual productivity.

Seen this way, the Personal Assistant is not isolated from the enterprise. It becomes the human-facing edge of a wider intelligence framework. The individual is still supported, but now in context, connected to workflows, informed by orchestration and ultimately aligned to the value stream outcomes the organisation exists to protect.

Most organisations will capture value from layers one and two. Some will progress to layer three and coordinate work more effectively. But the step change sits at layer four, because coherence is not reached by accumulating more local automation. It requires a deliberate architectural decision to design AI from the outcome down.

That means starting with the kept promise, identifying where the value stream is constrained and building intelligence that governs toward the outcome rather than merely optimising within fragments of it.

Why Anthropic's Dreams Matters

Anthropic's Dreams capability is significant because it points toward a different kind of agent memory. Most AI memory is retrieval. A system remembers what happened and can surface it when asked. That is useful, but it is not enough for governance.

The harder problem is consolidation. As agents become persistent, they accumulate transcripts, observations, assumptions, contradictions and outdated context. The question is no longer simply what should be remembered. The question is what should be refined, reconciled and carried forward as understanding.

The sleep analogy should not be pushed too far, but it is useful. In humans, sleep is understood to play a role in consolidating experience into more durable and usable memory. Recent events are replayed, reorganised and integrated into longer-term understanding. A manager, for example, does not become experienced simply by remembering separate escalations, delays and difficult meetings. Over time, those events consolidate into judgement about where confidence breaks down, what signals matter early and how to intervene before the promise is lost. Dreams appears to borrow that architectural idea. Agent sessions are not merely stored for retrieval but distilled into a refined memory store that can shape future reasoning.

That distinction matters. The difference between retrieval and consolidation is the difference between a system that logs and a system that learns.

A human manager does not become experienced simply by remembering past escalations. They become experienced when those events consolidate into judgement, with the ability to recognise patterns earlier, understand what matters and act before the promise is broken.

The same principle applies to enterprise AI. A value-stream agent should not merely remember that a supplier was late, that warehouse capacity was constrained or that an order missed its delivery promise. It should consolidate those events into a more useful understanding, such as when a supplier delay coincides with late-week capacity pressure, next-day commitments become vulnerable unless capacity is protected earlier in the cycle.

Dreams is still at research-stage. It is not yet a complete enterprise governance architecture. But it is directionally important because it points beyond AI that remembers events toward AI that refines experience into future judgement. That is precisely what a Coherence Layer requires.

Ackoff's Ladder and the Direction of Travel

Russell Ackoff's ladder provides a useful lens for this shift. The familiar progression from data, to information, to knowledge, to understanding and finally to wisdom describes more than an information hierarchy. It describes a movement from raw facts toward contextual judgement.

Enterprise AI appears to be following a similar path, not as a single product and not as a single vendor, but collectively. Data became searchable. Information became retrievable. Knowledge became synthesised and conversational. Agentic AI is now beginning to approach something closer to understanding, not human understanding, but an architectural form of understanding, with the ability to retain context, consolidate experience, recognise patterns over time and act with greater relevance in future situations.

Whether AI ever reaches wisdom is a different question. It may develop increasingly sophisticated operational judgement, by recognising patterns, anticipating risk, recommending action and learning from consequence. But wisdom requires more than effective reasoning. It requires responsibility for what should be done, what should not be done and which consequences an organisation is willing to own. That still remains a human and institutional burden.

That is why self-restraint matters. The most intelligent enterprise AI will not be the system that acts most often. It will be the system that knows when to act, when to learn and when to deliberately hold itself back.

The Convergence

SAP has world-class capability across layers one to three. Its agent architecture can coordinate workflows across the enterprise with a sophistication that was not possible only a few years ago. The Autonomous Enterprise represents genuine progress toward a more intelligent organisation.

Anthropic's Dreams points toward the kind of persistent memory and consolidation infrastructure that layer four will require. It suggests agents that do not merely execute, but carry forward refined experience into future decisions.

Neither organisation has yet named the governed value stream as the destination. Neither has delivered the full Coherence Layer out of the box, but the direction of travel is clear for all to see.

SAP is building more of the enterprise operating fabric. Dreams points toward the kind of persistent intelligence that could sit above it. Coherence is what becomes possible when both are deliberately designed around value, flow and outcomes.

The integration does not yet exist as a finished product. SAP workflows, enterprise events, agent memory, MCP-enabled tooling and consolidated reasoning would need to be deliberately architected, but that is not a weakness in the argument, it is the opportunity.

The organisations that design this integration deliberately, with coherence as the objective and the outcome as the governing purpose, may very likely gain a significant advantage that is not merely technical. They will learn earlier where value leaks, where promises become vulnerable, where cost to serve is rising without visibility and where AI should be allowed to act, learn or restrain itself.

The Implication: Excavate, Don't Explore

The conventional approach to AI implementation has been bottom-up. Find a use case. Automate it. Measure the efficiency gain. Find the next one. This approach is rational at the individual initiative level but dangerous at the enterprise level.

It produces disconnected automations that optimise locally while collectively creating bottlenecks, absorbing capacity and working against each other in ways nobody intended. This is how organisations can become faster in fragments while remaining incoherent as a whole.

The coherence framework inverts the problem. You do not search for AI use cases. You excavate them from the value stream.

The value stream is already laid out. The promises are already being made. The constraints are already binding. The leakage is already present. The task is not to explore an infinite catalogue of AI possibilities. It is to find where the promise is most at risk and where intelligent governance would release the greatest systemic value.

The most valuable AI use case is unlikely to be hidden in a Design Thinking workshop or a vendor catalogue. It is usually sitting at the point where throughput is most restricted, where variation is most costly, where workarounds accumulate and where the customer's promise is most vulnerable.

The road to value does not begin by exploring use cases. It begins by excavating the value stream.

This is what SAP's Autonomous Enterprise makes more visible; the possibility of seeing enterprise activity as flow rather than disconnected transactions. It is also what Anthropic's Dreams makes more intelligent. The possibility that agents may carry forward experience rather than reason only in the moment.

A Careful Prediction

Prediction is always dangerous. Technology markets move unevenly, vendors overstate capability and enterprises often adopt new tools through old mental models. But directionally, this destination now feels increasingly difficult to avoid.

As AI moves from task assistance to agents, from agents to memory and from memory to consolidation, the enterprise problem changes. The advantage will not simply belong to organisations that automate fastest. It will belong to those that recognise earliest that intelligence must be designed around value, flow and kept promises.

This is not just a first-mover advantage. It is a learning advantage. Organisations that move early will begin learning where their value streams leak, where outcomes become vulnerable and where intelligent restraint matters. That accumulated learning may prove difficult for slower organisations to recover.

The practical implication is simple. Enterprises should stop treating AI adoption as a catalogue exercise. They should begin by identifying their commitments that matter most, the value streams that keep them, the constraints that put them at risk and the points where intelligent governance would stop leakage and release the greatest value.

A Seminal Moment

The quality management movement spent decades building the intellectual foundations for what is now becoming technically possible. The founders of that world understood with beautiful clarity that organisations are not collections of processes. They are systems with a purpose, and that purpose is expressed in the promises they make and keep.

They gave us the concepts. They gave us the methods. What they could not give us was the technology to instantiate the control layer in real time, at scale, across the whole value stream.

That technology is now arriving. Not fully formed, not yet integrated and certainly not yet fully mature, but clearly visible in the direction that enterprise platforms and frontier AI research are now travelling.

For the first time since the value stream became a serious management idea, we can begin to imagine it as governable. Not merely described, not simply improved, but actually governed by a persistent intelligence layer capable of holding value flow in mind, detecting where the promise is at risk and acting with judgement rather than speed alone.

This is not another AI announcement. It is the beginning of a different enterprise question.

Not: where can we apply AI?

But: where is value leaking from the system, and what would it mean to govern that leakage with intelligence?

Value, properly understood, is a kept promise. The framework to navigate that moment is what From Efficiency to Intelligence was written to provide.

A Seminal Moment in Enterprise AI

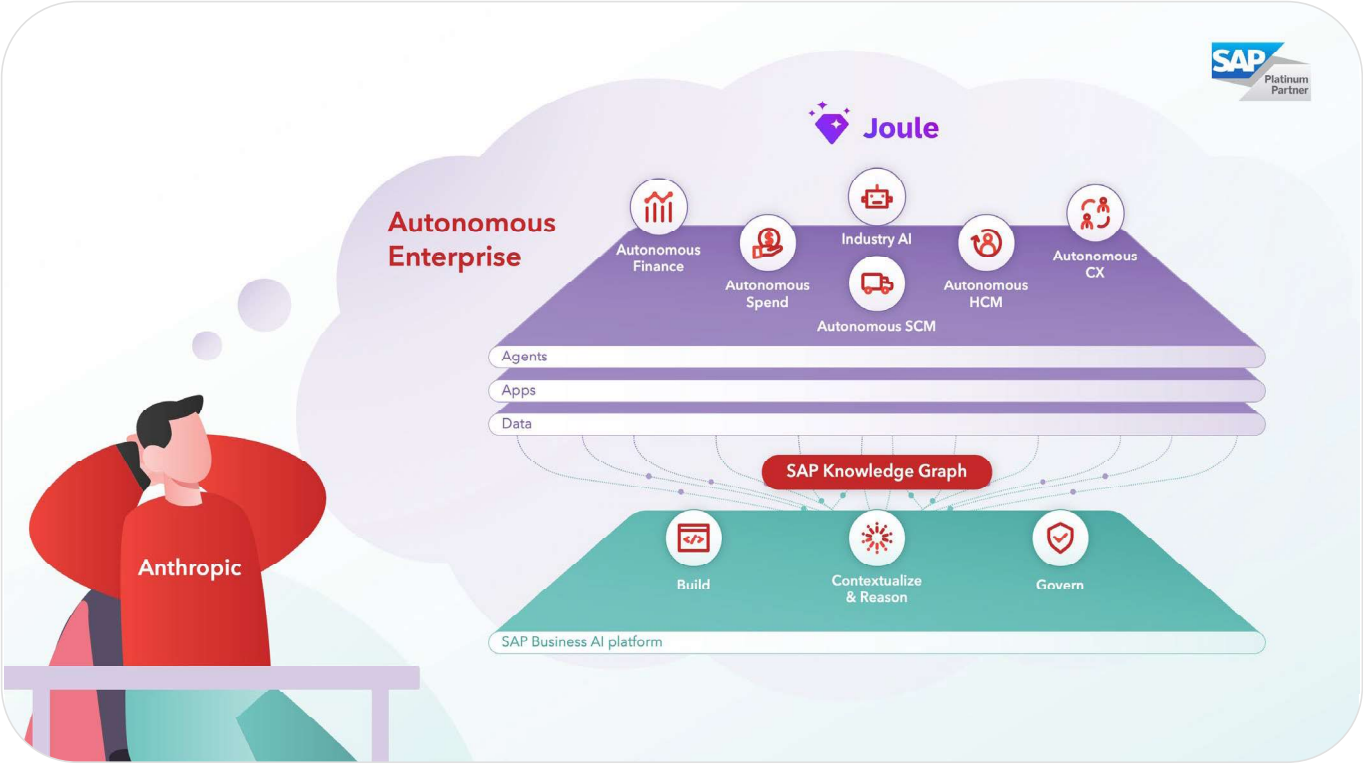

Mark Bolton

Head of AI, Data and Analytics for SAP

mark.bolton@delaware.co.uk

Connect on LinkedIn

Get in touch

delaware

we commit. we deliver.

delaware.co.uk